

Junos® OS

Monitoring, Sampling, and Collection Services Interfaces User Guide

Published
2022-12-15

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 Monitoring, Sampling, and Collection Services Interfaces User Guide
Copyright © 2022 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About This Guide | xxxi

1

Flow Monitoring and Flow Collection Services

Understanding Flow Monitoring | 2

Flow Monitoring Terms and Acronyms | 2

- active flow monitoring | 3
- Adaptive Services PIC | 3
- cflowd | 3
- content destination | 3
- control source | 3
- dynamic flow capture | 3
- DTCP (Dynamic Tasking Control Protocol) | 3
- ES PIC | 4
- flow collector interface | 4
- Monitoring Services PIC | 4
- Monitoring Services II PIC | 4
- Monitoring Services III PIC | 4
- MultiServices 100 PIC | 4
- MultiServices 400 PIC | 4
- MultiServices 500 PIC | 4
- passive flow monitoring | 4

Configuring Flow Monitoring | 5

Flow Monitoring Output Formats | 11

Flow Monitoring Version 5 Format Output Fields | 11

Flow Monitoring Version 8 Format Output Fields | 16

Flow Monitoring Version 9 Format Output Fields | 26

Monitoring Traffic Using Active Flow Monitoring | 41

Configuring Active Flow Monitoring | 42

Active Flow Monitoring System Requirements | 45

Active Flow Monitoring Applications	46
Active Flow Monitoring PIC Specifications	49
Active Flow Monitoring Overview	53
Active Flow Monitoring Overview	54
Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System	58
Example: Configuring Flow Monitoring on an MX Series Router with MS-MIC and MS-MPC	62
Hardware and Software Requirements	68
Junos Traffic Vision Support on MS-MIC and MS-MPC	69
Verification	70
Configuring Services Interface Redundancy with Flow Monitoring	72
Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250	74
Configuring Flow Offloading on MX Series Routers	84
Configuring Active Flow Monitoring on PTX Series Packet Transport Routers	85
Configuring Actively Monitored Interfaces on M, MX and T Series Routers	88
Collecting Flow Records	88
Configuring M, MX and T Series Routers for Discard Accounting with an Accounting Group	89
Configuring M, MX and T Series Routers for Discard Accounting with a Sampling Group	90
Configuring M, MX and T Series Routers for Discard Accounting with a Template	92
Defining a Firewall Filter on M, MX and T Series Routers to Select Traffic for Active Flow Monitoring	93
Processing IPv4 traffic on an M, MX or T Series Router Using Monitoring services, Adaptive services or Multiservices Interfaces	95
Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers	96
Replicating Version 9 Flow Aggregation From M, MX and T Series Routers to Multiple Flow Servers	97
Configuring Routing Engine-Based Sampling on M, MX and T Series Routers for Export to Multiple Flow Servers	98
Example: Copying Traffic to a PIC While an M, MX or T Series Router Forwards the Packet to the Original Destination	99
Requirements	100

Overview and Topology | **100**

Configuration | **101**

Verification | **122**

Configuring an Aggregate Export Timer on M, MX and T Series Routers for Version 8 Records | **126**

Example: Sampling Configuration for M, MX and T Series Routers | **127**

Associating Sampling Instances for Active Flow Monitoring with a Specific FPC, MPC, or DPC | **132**

Example: Sampling Instance Configuration | **133**

Example: Sampling and Discard Accounting Configuration on M, MX and T Series Routers | **141**

Monitoring Traffic Using Passive Flow Monitoring | 149

Passive Flow Monitoring Overview | **150**

Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers | **152**

Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers | **154**

Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers | **156**

Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers | **157**

Configuring Passive Flow Monitoring | **166**

Example: Passive Flow Monitoring Configuration on M, MX and T Series Routers | **168**

Configuring a Routing Table Group on an M, MX or T Series Router to Add Interface Routes into the Forwarding Instance | **187**

Using IPSec and an ES PIC on an M, MX or T Series Router to Send Encrypted Traffic to a Packet Analyzer | **188**

Applying a Firewall Filter Output Interface on an M, MX or T Series Router to Port-mirror Traffic to PICs or Flow Collection Services | **189**

Monitoring Traffic on a Router with a VRF Instance and a Monitoring Group | **190**

Specifying a Firewall Filter on an M, MX or T Series Router to Select Traffic to Monitor | **191**

Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers | **192**

Establishing a VRF Instance on an M, MX or T Series Router for Monitored Traffic | **196**

Configuring a Monitoring Group on an M, MX or T Series Router to Send Traffic to the Flow Server | **196**

Configuring Policy Options on M, MX or T Series Routers | 198

Stripping MPLS Labels on ATM, Ethernet-Based and SONET/SDH Router Interfaces | 199

Using an M, MX or T Series Router Flow Collector Interface to Process and Export Multiple Flow Records | 201

Example: Configuring a Flow Collector Interface on an M, MX or T Series Router | 208

Processing and Exporting Multiple Records Using Flow Collection | 225

Flow Collection Overview | 225

Configuring Flow Collection | 226

Example: Configuring Flow Collection | 231

Sending cflowd Records to Flow Collector Interfaces | 239

Configuring Flow Collection Mode and Interfaces on Router Services PICs on M and T Series Routers | 239

Logging Flow Monitoring Records with Version 9 and IPFIX Templates for NAT Events | 241

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241

Configure Active Flow Monitoring Logs for NAT44/NAT64 | 254

Overview | 254

Requirements | 254

Configuration | 254

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256

Exporting Syslog Messages to an External Host Without Flow Monitoring Formats Using an MX Series Router or NFX250 | 258

Exporting Version 9 Flow Data Records to a Log Collector Overview Using an MX Series Router or NFX250 | 259

Understanding Exporting IPFIX Flow Data Records to a Log Collector Using an MX Series Router or NFX250 | 260

Mapping Between Field Values for Version 9 Flow Templates and Logs Exported From an MX-Series Router or NFX250 | 262

Mapping Between Field Values for IPFIX Flow Templates and Logs Exported From an MX Series Router or NFX250 | 265

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275

Requirements | 275

Generation of Log Messages Using Flow Templates for NAT Operations on MS-MPCs, MS-MICs, and MX-SPC3s | 276

Configuration | 276

Verification | 280

2

Flow Capture Services

Dynamically Capturing Packet Flows Using Junos Capture Vision | 286

Understanding Junos Capture Vision | 286

Configuring Junos Capture Vision | 289

Example: Configuring Junos Capture Vision on M and T Series Routers | 297

Monitoring a Capture Group Using SNMP or Show Services Commands | 301

Detecting Threats and Intercepting Flows Using Junos Packet Vision | 302

Understanding Junos Packet Vision | 302

Configuring Junos Packet Vision on MX, M and T Series Routers | 303

Examples: Configuring Junos Packet Vision on M, T, and MX Series Routers | 306

Sending Packets to a Mediation Device on MX, M and T Series Routers | 309

Example: Configuring IPv6 Support for FlowTapLite on an M120 Router With Enhanced III FPCs | 310

Requirements | 311

Overview and Topology | 312

Configuration | 313

Verification | 317

Using Flow-Tap to Monitor Packet Flow | 321

Understanding Flow-Tap Architecture | 321

Configuring a Flow-Tap Interface on MX, M and T Series Routers | 325

Configuring Flow-Tap Security Properties on MX, M and T Series Routers | 326

Flow-Tap Application Restrictions | 327

Example: Flow-Tap Configuration on T and M Series Routers | 327

Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs | 329

Inline Monitoring Services and Inband Network Telemetry

Inline Monitoring Services | 334

Inline Monitoring Services Configuration | 334

Understanding Inline Monitoring Services | 334

Configuring Inline Monitoring Services | 342

Flow-Based Telemetry | 349

Flow-Based Telemetry (EX4100, EX4100-F, and EX4400 Series) | 349

FBT Overview | 349

Configure FBT (EX4100, EX4100-F, and EX4400 Series) | 356

Flow-Based Telemetry for VXLANs (QFX5120) | 361

FBT for VXLANs Overview | 361

Configure FBT for VXLANs (QFX5120) | 366

Inband Flow Analyzer 2.0 | 370

Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring | 370

Inband Flow Analyzer 2.0 | 370

Configure Inband Flow Analyzer 2.0 | 384

Configure IFA Initiator Node | 388

Configure IFA Transit Node | 391

Configure IFA Terminating Node | 391

View Inband Flow Analyzer Statistics | 393

Example - Configure Inband Flow Analyzer 2.0 for Traffic Monitoring | 394

Juniper Resiliency Interface | 409

Juniper Resiliency Interface | 409

Understand Juniper Resiliency Interface | 409

Configure JRI for Operating System and Routing Exceptions | 412

Configure JRI for Forwarding Exceptions | 413

Sampling and Discard Accounting Services

Sampling Data Using Traffic Sampling and Discard Accounting | 420

Configuring Traffic Sampling on MX, M and T Series Routers | 420

Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches | 433

Configuring Discard Accounting | 435

Sampling Data Using Inline Sampling | 437

Understand Inline Active Flow Monitoring | 437

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 527

Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers | 537

Configuring Inline Active Flow Monitoring on PTX Series Routers | 540

Platform and Feature Support | 540

How to Configure Inline Active Flow Monitoring on PTX Series Routers | 543

Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers | 550

MPLS-over-UDP Flow Monitoring Overview | 550

Configuring Inline Active Flow Monitoring of MPLS-over-UDP Flows | 553

Configuring the Template to Specify Output Properties | 554

Configuring the Sampling Instance | 555

Assigning the Sampling Instance to an FPC | 557

Configuring a Firewall Filter | 557

Assigning the Firewall Filter to the Monitored Interface | 557

Inline Active Flow Monitoring on IRB Interfaces | 558

Overview | 558

Understand Inline Active Flow Monitoring on IRB interfaces | 559

Configure Inline Active Flow Monitoring on IRB Interfaces on PTX Series Routers | 561

Configure the Template to Specify Output Properties | 561

Configure the Sampling Instance | 562

Assign the Sampling Instance to an FPC | 564

Configure a Firewall Filter | 564

Associate a Layer 3 Interface with the VLAN to Route Traffic | 565

Assign the Firewall Filter to the Monitored Interface | 566

Example: Configuring Inline Active Flow Monitoring on MX Series and T4000 Routers | 567

Software and Hardware Requirements | 574

Overview | 575

Sampling Data Using Flow Aggregation | 576

Understanding Flow Aggregation | 576

Enabling Flow Aggregation | 577

Configuring Flow Aggregation on MX, M and T Series Routers and NFX250 to Use Version 5 or Version 8 cflowd | 578

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583

Configuring Flow Aggregation on PTX Series Routers to Use Version 9 Flow Templates | 596

Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603

Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers | 615

Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows | 620

Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows | 625

Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers | 633

Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers | 637

Logging cflowd Flows on M and T Series Routers Before Export | 640

Configuring Next-Hop Address Learning on MX Series and PTX Series Routers for Destinations Accessible Over Multiple Paths | 641

5

Real-Time Performance Monitoring and Video Monitoring Services

Monitoring Traffic Using Real-Time Performance Monitoring | 645

Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651

Configuring RPM Receiver Servers | 662

Limiting the Number of Concurrent RPM Probes on M, MX, T and PTX Routers and EX Series Switches | 663

Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches | 663

Analyzing Network Efficiency in IPv6 Networks on MX Series Routers Using RPM Probes | 668

Configuring BGP Neighbor Discovery Through RPM | 671

Examples: Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers With RPM | **674**

Trace RPM Operations | **676**

RPM Trace Operations Overview | **676**

Configure the Trace Operations | **677**

Configure the RPM Log File Name | **678**

Configure the Number and Size of RPM Log Files | **678**

Configure Access to the Log File | **679**

Configure a Regular Expression for Lines to Be Logged | **679**

Examples: Configuring Real-Time Performance Monitoring on MX, M, T and PTX Series Routers | **680**

Enabling RPM on MX, M and T Series Routers and SRX Firewalls for the Services SDK | **685**

Understand Two-Way Active Measurement Protocol | **686**

Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches | **694**

Understand TWAMP Configuration | **695**

Configure a TWAMP Server | **699**

Configure a TWAMP Client | **702**

Example: Configuring TWAMP Client and Server on MX Series Routers | **705**

Requirements | **705**

Overview | **706**

Configuration for TWAMP client | **706**

Configuration for TWAMP server | **709**

Verification | **712**

Understanding TWAMP Auto-Restart | **713**

Configuring TWAMP Client and TWAMP Server to Reconnect Automatically After TWAMP Server Unavailability | **716**

Managing License Server for Throughput Data Export | 724

License Server Management for Throughput Data Export on MX Series Routers for NAT, Firewall, and Inline Flow Monitoring Services | **724**

Guidelines for Configuring an MX Series Router to Transmit Per-Service Throughput to an External Log Collector | **726**

Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking | 728

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | **728**

Understanding RFC2544-Based Benchmarking Tests for E-LAN and E-Line Services on MX Series Routers | **733**

Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | **737**

Configuring an RFC 2544-Based Benchmarking Test | **739**

Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network | **741**

Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire | **743**

Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain | **745**

Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers | **747**

Example: Configure an RFC 2544-Based Benchmarking Test on an MX104 Router for Layer 3 IPv4 Services | **749**

Requirements | **749**

Overview | **750**

Configuration | **750**

Verify the Results of the Benchmarking Test for Layer 3 IPv4 Services | **762**

Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for UNI Direction of Ethernet Pseudowires | **763**

Requirements | **763**

Overview | **763**

Configuration | **764**

Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service | **774**

Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for NNI Direction of Ethernet Pseudowires | **775**

Requirements | **775**

Overview | **776**

Configuration | **777**

Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service | **787**

Example: Configuring RFC2544-Based Benchmarking Tests on an MX104 Router for Layer 2 E-LAN Services in Bridge Domains | **788**

Requirements | **788**

Overview | **788**

Configuration | **790**

Verifying the Results of the Benchmarking Tests for Layer 2 Services (E-LAN) in Bridge Domains | **810**

Example: Configuring Benchmarking Tests to Measure SLA Parameters for E-LAN Services on an MX104 Router Using VPLS | **823**

Requirements | **824**

Overview | **824**

Configuration | **825**

Verifying the Results of the Benchmarking Test for Layer 2 ELAN Services Using VPLS | **853**

Configuring RFC 2544-Based Benchmarking Tests on ACX Series | 855

RFC 2544-Based Benchmarking Tests for ACX Routers Overview | **855**

Layer 2 and Layer 3 RFC 2544-Based Benchmarking Test Overview | **859**

Configuring RFC 2544-Based Benchmarking Tests | **864**

Test Profile and Test Name Overview | **865**

Configure a Test Profile for an RFC 2544-Based Benchmarking Test | **871**

Configure a Test Name for an RFC 2544-Based Benchmarking Test For an Initiator | **874**

Configure a Test Name for an RFC 2544-Based Benchmarking Test For a Reflector | **878**

Start and Stop the RFC 2544-Based Benchmarking Test | **881**

Copying an RFC 2544-Based Benchmarking Test Result | **881**

Configuring Ethernet Loopback for RFC 2544-Based Benchmarking Tests | **882**

RFC 2544-Based Benchmarking Test States | **885**

Example: Configure an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services | **887**

Requirements | **887**

Overview | **887**

Configuration | **888**

Verify the Results of the Benchmarking Test for Layer 3 IPv4 Services | **900**

Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires | **901**

Requirements | **902**

Overview | **902**

Configuration | **903**

Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service | 912

Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires | 913

Requirements | 914

Overview | 914

Configuration | 915

Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service | 925

Configuring a Service Package to be Used in Conjunction with PTP | 926

Tracking Streaming Media Traffic Using Inline Video Monitoring | 927

Understanding Inline Video Monitoring on MX Series Routers | 927

Configuring Inline Video Monitoring on MX Series Routers | 934

Configuring Media Delivery Indexing Criteria | 934

Configuring Interface Flow Criteria | 937

Configuring the Number of Flows That Can Be Measured | 945

Inline Video Monitoring Syslog Messages on MX Series Routers | 946

Generation of SNMP Traps and Alarms for Inline Video Monitoring on MX Series Routers | 947

SNMP Traps for Inline Video Monitoring Statistics on MX Series Routers | 951

Processing SNMP GET Requests for MDI Metrics on MX Series Routers | 952

6

Configuration Statements and Operational Commands

Configuration Statements | 955

accounting | 965

address (Interfaces) | 967

address (Services Dynamic Flow Capture) | 968

aggregate-export-interval | 970

aggregation | 971

alarms | 973

alarm-mode | 975

allowed-destinations | **977**

analyzer-address | **978**

analyzer-id | **979**

archive-sites | **980**

authentication-mode | **982**

authentication-key-chain (TWAMP) | **983**

autonomous-system-type | **985**

bandwidth-kbps | **987**

bgp | **988**

bridge-template | **990**

capture-group | **991**

category | **993**

cflowd (Discard Accounting) | **995**

cflowd (Flow Monitoring) | **997**

client | **998**

client-delegate-probes | **1002**

client-list | **1003**

collector | **1005**

collector (Inline Monitoring) | **1006**

collector (Flow Monitoring Logs for NAT) | **1009**

collector (Flow Template Profiles for NAT) | **1010**

collector-group (Flow Template Profiles for NAT) | **1012**

collector-group (Flow Monitoring Logs for NAT) | **1014**

content-destination | **1016**

control-connection (Junos OS) | **1017**

control-connection (Junos OS Evolved) | **1019**

control-source | **1024**

controller | **1025**

core-dump | **1028**

data-fill | **1029**

data-fill-with zeros | **1031**

data-format | **1032**

data-record-fields | **1033**

data-size | **1037**

delay-factor | **1039**

delegate-probes | **1041**

destination (Interfaces) | **1043**

destination-address (Flow Monitoring Logs for NAT) | **1044**

destination-interface | **1046**

destination-ipv4-address | **1048**

destination-mac-address | **1050**

destination-port | **1051**

destination-port (Flow Monitoring Logs for NAT) | **1054**

destination-udp-port | **1055**

destinations | **1057**

direction | **1058**

disable (Forwarding Options) | **1060**

disable-signature-check | **1062**

dscp (flow-server) | **1063**

dscp-code-points (RPM) | **1065**

dscp-code-points (RFC 2544 Benchmarking) | **1067**

dump-on-flow-control | **1070**

duplicates-dropped-periodicity | 1071

dynamic-flow-capture | 1072

em-hw-profile | 1074

engine-id (Forwarding Options) | 1076

engine-type | 1077

exception-reporting | 1079

exceptions | 1080

export-format | 1082

family (Monitoring) | 1083

family | 1085

family (Sampling) | 1087

features | 1090

file (Sampling) | 1092

file (Trace Options) | 1094

file-specification (File Format) | 1095

file-specification (Interface Mapping) | 1096

filename | 1098

filename-prefix | 1099

files | 1100

filter | 1102

flex-flow-sizing | 1103

flow-active-timeout | 1105

flow-collector | 1107

flow-control-options | 1109

flow-export-destination | 1111

flow-export-rate | 1113

flow-export-timer | **1114**

flow-inactive-timeout | **1116**

flow-key (Flow Monitoring) | **1117**

flow-monitoring | **1119**

flow-monitoring (Inline Monitoring Services) | **1122**

flow-server | **1125**

flow-table-size | **1127**

flow-table-size (Chassis) | **1129**

flow-tap | **1130**

forwarding-class (RFC 2544 Benchmarking) | **1132**

forwarding-class (Sampling) | **1134**

ftp (Flow Collector Files) | **1135**

ftp (Transfer Log Files) | **1138**

g-duplicates-dropped-periodicity | **1139**

g-max-duplicates | **1140**

generate-snmp-traps | **1142**

halt-on-prefix-down (RFC 2544 Benchmarking) | **1143**

hard-limit | **1145**

hard-limit-target | **1146**

hardware-timestamp | **1147**

history-size | **1148**

host-outbound media-interface | **1150**

icmp | **1151**

in-service | **1153**

inactivity-timeout (Services RPM) | **1155**

inet6-options (Services) | **1156**

inband-flow-telemetry | **1157**

inline-jflow | **1162**

inline-monitoring | **1163**

instance | **1165**

input (Sampling) | **1168**

input-interface-index | **1169**

input-packet-rate-threshold | **1170**

instance (Sampling) | **1171**

interface (Accounting or Sampling) | **1174**

interfaces | **1175**

interface (Services Flow Tap) | **1177**

interface-map | **1178**

interfaces (Services Dynamic Flow Capture) | **1179**

interfaces (Video Monitoring) | **1181**

inet6-options (Services) | **1185**

ipfix-sw-mode | **1186**

ip-swap | **1187**

ipv4-flow-table-size | **1189**

ipv4-template | **1191**

ipv6-flow-table-size | **1192**

ipv6-extended-attrb | **1194**

ipv6-template | **1195**

ivlan-cfi (RFC 2544 Benchmarking) | **1196**

ivlan-id (RFC 2544 Benchmarking) | **1198**

ivlan-priority (RFC 2544 Benchmarking) | **1199**

jflow-log (Interfaces) | **1200**

jflow-log (Services) | 1202

label-position | 1204

license-server | 1205

light | 1207

local-dump | 1209

logical-system | 1210

managed | 1211

match | 1214

max-connection-duration | 1215

max-duplicates | 1216

max-packets-per-second | 1218

maximum-age | 1219

maximum-connections | 1221

maximum-connections-per-client | 1222

maximum-packet-length | 1224

maximum-sessions | 1226

maximum-sessions-per-connection | 1227

media-loss-rate | 1229

media-rate-variation | 1230

message-rate-limit (Flow Monitoring Logs for NAT) | 1232

minimum-priority | 1233

mode | 1235

monitoring (Forwarding Options) | 1236

monitoring (Services) | 1238

moving-average-size | 1241

mpls-flow-table-size | 1243

mpls-ipv4-template | 1245

mpls-ipvx-template | 1246

mpls-template | 1248

multiservice-options | 1250

name-format | 1251

next-hop (Forwarding Options) | 1253

next-hop (RPM) | 1255

next-hop-group (Forwarding Options) | 1256

nexthop-learning | 1258

no-remote-trace (Trace Options) | 1260

no-syslog | 1261

no-syslog-generation | 1262

notification-targets | 1264

observation-domain-id | 1265

offload-type | 1267

one-way-hardware-timestamp | 1268

option-refresh-rate | 1270

options-template-id | 1271

outer-tag-protocol-id (RFC 2544 Benchmarking) | 1273

output (Accounting) | 1275

output (Monitoring) | 1276

output (Sampling) | 1278

output-interface-index | 1280

ovlan-cfi (RFC 2544 Benchmarking) | 1282

ovlan-id (RFC 2544 Benchmarking) | 1283

ovlan-priority (RFC 2544 Benchmarking) | 1284

owner | **1286**

packet-loss-priority (RFC 2544 Benchmarking) | **1288**

packet-size (RFC 2544 Benchmarking) | **1289**

passive-monitor-mode | **1291**

password (Flow Collector File Servers) | **1292**

password (Transfer Log File Servers) | **1293**

peer-as-billing-template | **1294**

persistent-results | **1296**

pfe | **1297**

pic-memory-threshold | **1300**

pop-all-labels | **1301**

port (Flow Monitoring) | **1303**

port (RPM) | **1304**

port (TWAMP) | **1306**

post-cli-implicit-firewall | **1307**

pre-rewrite-tos | **1309**

primary-data-record-fields | **1310**

probe | **1313**

probe-count | **1316**

probe-interval | **1317**

probe-limit | **1319**

probe-server | **1320**

probe-type | **1323**

profiles (RFC 2544 Benchmarking) | **1324**

rate (Interface Services) | **1326**

rate (Forwarding Options) | **1327**

receive-failure-threshold (RFC 2544 Benchmarking) | **1329**

receive-options-packets | **1330**

receive-ttl-exceeded | **1331**

reflect-etype | **1333**

reflect-mode | **1334**

refresh-rate (Flow Monitoring Logs for NAT) | **1337**

required-depth | **1338**

resiliency | **1340**

retry (Services Flow Collector) | **1343**

retry-delay | **1344**

rfc2544 | **1345**

rfc2544-benchmarking | **1347**

routing-instance (RPM) | **1350**

routing-instance (cflowd) | **1352**

routing-instance-list (TWAMP) | **1354**

routing-instances | **1355**

rpm (Interfaces) | **1357**

rpm (Services) | **1358**

rpm-scale | **1372**

rpm-tracking | **1375**

run-length | **1377**

sample-once | **1379**

sampling (Forwarding Options) | **1380**

sampling (Interfaces) | **1384**

sampling-instance | **1385**

server (Junos OS) | **1387**

server (Junos OS Evolved) | **1388**

server-inactivity-timeout | **1391**

service-port | **1392**

service-type | **1393**

services-options | **1395**

shared-key | **1397**

size | **1398**

skip-arp-iteration (RFC 2544 Benchmarking) | **1400**

slamon-services | **1401**

soft-limit | **1403**

soft-limit-clear | **1404**

source-address (Forwarding Options) | **1405**

source-address (RPM) | **1407**

source-address (TWAMP) | **1408**

source-addresses | **1410**

source-id | **1412**

source-ip (Flow Monitoring Logs for NAT) | **1413**

source-ipv4-address (RFC 2544 Benchmarking) | **1415**

source-mac-address | **1416**

source-udp-port (RFC 2544 Benchmarking) | **1418**

stamp | **1420**

step-percent (RFC 2544 Benchmarking) | **1421**

store | **1423**

storm-control | **1426**

syslog | **1427**

target-address | **1428**

tcp | **1431**

tcp-keepcnt | **1432**

tcp-keepidle | **1434**

tcp-keepintvl | **1435**

template (Flow Monitoring IPFIX Version) | **1436**

template (Flow Monitoring Version 9) | **1438**

template (Forwarding Options) | **1440**

template (Forwarding Options Version IPFIX) | **1441**

template (Inline Monitoring) | **1442**

template-id | **1446**

template-profile (Flow Monitoring Logs for NAT) | **1448**

template-refresh-rate | **1449**

template-type (Flow Monitoring Logs for NAT) | **1451**

templates | **1452**

test | **1456**

tests | **1459**

test-count | **1462**

test-finish-wait-duration (RFC 2544 Benchmarking) | **1463**

test-interface (RFC 2544 Benchmarking) | **1465**

test-interval | **1467**

test-iterator-duration (RFC 2544 Benchmarking) | **1469**

test-iterator-pass-threshold (RFC 2544 Benchmarking) | **1470**

test-name | **1471**

test-profile (RFC 2544 Benchmarking) | **1474**

test-session (Junos OS) | **1476**

test-session (Junos OS Evolved) | **1477**

test-type (RFC 2544 Benchmarking) | 1483

thresholds (Junos OS) | 1485

thresholds (Junos OS Evolved) | 1487

timestamp-format (RFC 2544 Benchmarking) | 1489

traceoptions (Dynamic Flow Capture) | 1490

traceoptions (Forwarding Options) | 1492

traceoptions (Inline Monitoring) | 1493

traceoptions (Resiliency) | 1496

traceoptions (RPM) | 1499

transfer | 1502

transfer-log-archive | 1503

transmit-failure-threshold (RFC 2544 Benchmarking) | 1504

traps | 1506

ttl | 1509

ttl (RPM probe) | 1511

tunnel-observation | 1513

twamp | 1515

twamp-server | 1521

trio-flow-offload | 1522

udp | 1524

udp-tcp-port-swap | 1526

unit | 1527

use-extended-flow-memory | 1529

username (Services) | 1530

variant | 1532

version | 1533

version (Flow Monitoring Logs for NAT) | 1534

version9 (Forwarding Options) | 1536

version9 (Flow Monitoring) | 1537

version-ipfix (Forwarding Options) | 1539

version-ipfix (Services) | 1540

video-monitoring | 1542

vpls-flow-table-size | 1546

vpls-template | 1548

world-readable | 1549

Operational Commands | 1551

clear passive-monitoring statistics | 1553

clear services accounting statistics inline-jflow | 1555

clear services dynamic-flow-capture | 1556

clear services flow-collector statistics | 1558

clear inband-flow-telemetry stats | 1560

clear services inline-monitoring statistics | 1561

clear services monitoring rfc2544 | 1562

clear services rpm rfc2544-benchmarking | 1564

clear services monitoring twamp server control-connection | 1565

clear services rpm twamp server connection | 1566

clear services service-sets statistics jflow-log | 1567

clear services video-monitoring mdi errors fpc-slot | 1569

clear services video-monitoring mdi statistics fpc-slot | 1571

request services flow-collector change-destination primary interface | 1572

request services flow-collector change-destination secondary interface | 1573

request services flow-collector test-file-transfer | 1575

request services monitoring twamp client | **1577**

request services rpm twamp | **1578**

show forwarding-options next-hop-group | **1580**

show forwarding-options port-mirroring | **1584**

show interfaces (Dynamic Flow Capture) | **1587**

show interfaces (Flow Collector) | **1594**

show interfaces (Flow Monitoring) | **1603**

show passive-monitoring error | **1611**

show passive-monitoring flow | **1614**

show passive-monitoring memory | **1618**

show passive-monitoring status | **1620**

show passive-monitoring usage | **1622**

show route rpm-tracking | **1625**

show services accounting aggregation | **1630**

show services accounting aggregation template | **1636**

show services accounting errors | **1638**

show services accounting flow | **1645**

show services accounting flow-detail | **1654**

show services accounting memory | **1661**

show services accounting packet-size-distribution | **1664**

show services accounting status | **1666**

show services accounting usage | **1671**

show services dynamic-flow-capture content-destination | **1674**

show services dynamic-flow-capture control-source | **1676**

show services dynamic-flow-capture statistics | **1680**

show services flow-collector file interface | **1684**

show services flow-collector input interface | **1688**

show services flow-collector interface | **1690**

show services inband-flow-telemetry | **1701**

show services inline-monitoring feature-profile-mapping fpc-slot | **1704**

show services inline-monitoring statistics fpc-slot | **1707**

show services monitoring rfc2544 | **1710**

show services monitoring rfc2544 | **1715**

show services monitoring rpm history-results | **1719**

show services monitoring rpm probe-results | **1727**

show services monitoring twamp client control-info | **1740**

show services monitoring twamp client history-results | **1743**

show services monitoring twamp client probe-results | **1749**

show services monitoring twamp client test-info | **1759**

show services monitoring twamp server control-info | **1762**

show services monitoring twamp server test-info | **1764**

show services rpm active-servers | **1768**

show services rpm history-results | **1769**

show services rpm probe-results | **1775**

show services rpm rfc2544-benchmarking | **1791**

show services rpm rfc2544-benchmarking test-id | **1800**

show services rpm twamp client connection | **1826**

show services rpm twamp client history-results | **1828**

show services rpm twamp client probe-results | **1833**

show services rpm twamp client session | **1844**

show services rpm twamp server connection | **1847**

show services rpm twamp server session | **1849**

show services service-sets statistics jflow-log | **1852**

show services video-monitoring mdi errors fpc-slot | **1862**

show services video-monitoring mdi flows fpc-slot | **1865**

show services video-monitoring mdi stats fpc-slot | **1872**

test services monitoring rfc2544 | **1876**

test services rpm rfc2544-benchmarking test | **1880**

About This Guide

Use this guide to configure traffic flow monitoring, packet flow capture, inline monitoring, traffic sampling for accounting or discard, real-time performance monitoring (RPM and TWAMP), RFC 2544 performance benchmarking, and inline video monitoring.

1

PART

Flow Monitoring and Flow Collection Services

Understanding Flow Monitoring | 2

Monitoring Traffic Using Active Flow Monitoring | 41

Monitoring Traffic Using Passive Flow Monitoring | 149

Processing and Exporting Multiple Records Using Flow Collection | 225

Logging Flow Monitoring Records with Version 9 and IPFIX Templates for NAT Events | 241

CHAPTER 1

Understanding Flow Monitoring

IN THIS CHAPTER

- Flow Monitoring Terms and Acronyms | 2
- Configuring Flow Monitoring | 5
- Flow Monitoring Output Formats | 11
- Flow Monitoring Version 5 Format Output Fields | 11
- Flow Monitoring Version 8 Format Output Fields | 16
- Flow Monitoring Version 9 Format Output Fields | 26

Flow Monitoring Terms and Acronyms

IN THIS SECTION

- active flow monitoring | 3
- Adaptive Services PIC | 3
- cflowd | 3
- content destination | 3
- control source | 3
- dynamic flow capture | 3
- DTCP (Dynamic Tasking Control Protocol) | 3
- ES PIC | 4
- flow collector interface | 4
- Monitoring Services PIC | 4
- Monitoring Services II PIC | 4
- Monitoring Services III PIC | 4
- MultiServices 100 PIC | 4

- MultiServices 400 PIC | 4
- MultiServices 500 PIC | 4
- passive flow monitoring | 4

active flow monitoring

Technique to lawfully intercept and observe specified data network traffic on an active router participating in the network.

Adaptive Services PIC

Advanced PIC that handles active flow monitoring, Network Address Translation (NAT), stateful firewall, and intrusion detection functions. For more information on the Adaptive Services PIC, see the *Junos Services Interfaces Configuration Guide*.

cflowd

Version 5 and version 8 flow monitoring process that captures flow information from network traffic and exports this data into summary tables. Once captured, flow data can be analyzed as needed. For more information about cflowd, see <http://www.caida.org>.

content destination

A recipient of monitored packets sent by a DTCP or dynamic flow capture-enabled monitoring station.

control source

A dynamic flow capture client that wants to monitor electronic data or voice transfer over the network. The control source sends filter requests to the dynamic flow capture-enabled monitoring station by using DTCP.

dynamic flow capture

Technique that allows DTCP-enabled control sources to send specified filtering criteria in real time to a monitoring station. The monitoring station passively monitors the specified traffic flows on demand and sends the captured packets to content destinations.

DTCP (Dynamic Tasking Control Protocol)

Protocol used to specify filtering criteria in a dynamic flow capture environment.

ES PIC

PIC that handles encryption and security services (such as IP Security [IPSec]).

flow collector interface

Converted Monitoring Services II PIC that processes multiple flow records into compressed ASCII data files and exports these files to an FTP server.

Monitoring Services PIC

Original PIC that handles passive and active flow monitoring functions.

Monitoring Services II PIC

Advanced PIC that handles passive flow monitoring functions.

Monitoring Services III PIC

Advanced PIC that handles dynamic flow capture functions.

MultiServices 100 PIC

Also referred to as MultiServices PIC Type 1. Advanced PIC that handles active flow capture functions.

MultiServices 400 PIC

Also referred to as MultiServices PIC Type 2. Advanced PIC that handles active flow capture functions.

MultiServices 500 PIC

Also referred to as MultiServices PIC Type 3. Advanced PIC that handles active flow capture functions.

passive flow monitoring

Technique to lawfully intercept and observe specified data network traffic on a passive flow monitoring station not participating in the network.

Configuring Flow Monitoring

IN THIS SECTION

- [Configuring Flow-Monitoring Interfaces | 5](#)
- [Configuring Flow-Monitoring Properties | 7](#)
- [Example: Configuring Flow Monitoring | 9](#)

The flow-monitoring application performs traffic flow monitoring and enables lawful interception of traffic between two routers or switches. Traffic flows can either be passively monitored by an offline router or switch or actively monitored by a router participating in the network.

Configuring Flow-Monitoring Interfaces

To enable flow monitoring on the Monitoring Services PIC, include the `mo-fpc/pic/port` statement at the `[edit interfaces]` hierarchy level:

```
mo-fpc/pic/port {
  unit logical-unit-number {
    family inet {
      address address {
        destination address;
      }
      filter {
        group filter-group-number;
        input filter-name;
        output filter-name;
      }
      sampling {
        [ input output ];
      }
    }
  }
  multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
    flow-control-options {
```



```

        down-on-flow-control;
        dump-on-flow-control;
        reset-on-flow-control;
    }
}
}

```

Specify the physical and logical location of the flow-monitoring interface. You cannot use unit 0, because it is already used by internal processes. Specify the source and destination addresses. The filter statement allows you to associate an input or output filter or a filter group that you have already configured for this purpose. The sampling statement specifies the traffic direction: input, output, or both.

The multiservice-options statement allows you to configure properties related to flow-monitoring interfaces:

- Include the core-dump statement to enable storage of core files in **/var/tmp**.
- Include the syslog statement to enable storage of system logging information in **/var/log**.

NOTE: Boot images for monitoring services interfaces are specified at the [edit chassis images pic] hierarchy level. You must include the following configuration to make the flow monitoring feature operable:

```

[edit system]
ntp {
    boot-server ntp.example.net;
    server 172.17.28.5;
}
processes {
    ntp enable;
}

```

- Include the flow-control-options statement to configure flow control.

NOTE: Starting with Junos OS Release 15.1, the multiservices PIC management daemon core file is generated when a prolonged flow control failure occurs and when you configure the setting to generate a core dump during prolonged flow control (by using the dump-on-flow-control option with the flow-control-options statement). The watchdog functionality continues to generate a kernel core file in such scenarios. In Junos OS Release 14.2 and earlier, an

eJunos kernel core file is generated when a prolonged flow control failure occurs and when you configure the setting to generate a core dump during prolonged flow control.

Configuring Flow-Monitoring Properties

To configure flow-monitoring properties, include the `monitoring` statement at the [edit forwarding-options] hierarchy level:

```
monitoring name {
  family inet {
    output {
      cflowd hostname port port-number;
      export-format format;
      flow-active-timeout seconds;
      flow-export-destination {
        collector-pic;
      }
      flow-inactive-timeout seconds;
      interface interface-name {
        engine-id number;
        engine-type number;
        input-interface-index number;
        output-interface-index number;
        source-address address;
      }
    }
  }
}
```

A monitoring instance is a named entity that specifies collector information under the `monitoring name` statement. The following sections describe the properties you can configure:

Directing Traffic to Flow-Monitoring Interfaces

To direct traffic to a flow-monitoring interface, include the `interface` statement at the [edit forwarding-options monitoring *name* output] hierarchy level. By default, the Junos OS automatically assigns values for the `engine-id` and `engine-type` statements:

- `engine-id`—Monitoring interface location.
- `engine-type`—Platform-specific monitoring interface type.

The source-address statement specifies the traffic source for transmission of cflowd information; you must configure it manually. If you provide a different source-address statement for each monitoring services output interface, you can track which interface processes a particular cflowd record.

By default, the input-interface-index value is the SNMP index of the input interface. You can override the default by including a specific value. The input-interface-index and output-interface-index values are exported in fields present in the cflowd version 5 flow format.

Exporting Flows

To direct traffic to a flow collection interface, include the flow-export-destination statement. For more information about flow collection, see ["Active Flow Monitoring Overview" on page 54](#).

To configure the cflowd version number, include the export-format statement at the [edit forwarding-options monitoring *name* output] hierarchy level. By default, version 5 is used. Version 8 enables the router software to aggregate the flow information using broader criteria and reduce cflowd traffic. Version 8 aggregation is performed periodically (every few seconds) on active flows and when flows are allowed to expire. Because the aggregation is performed periodically, active timeout events are ignored.

For more information on cflowd properties, see ["Enabling Flow Aggregation" on page 577](#).

Configuring Time Periods When Flow Monitoring Is Active and Inactive

To configure time periods for active flow monitoring and intervals of inactivity, include the flow-active-timeout and flow-inactive-timeout statements at the [edit forwarding-options monitoring *name* output] hierarchy level:

- The flow-active-timeout statement specifies the time interval between flow exports for active flows. If the interval between the time the last packet was received and the time the flow was last exported exceeds the configured value, the flow is exported.

This timer is needed to provide periodic updates when a flow has a long duration. The active timeout setting enables the router to retain the start time for the flow as a constant and send out periodic cflowd reports. This in turn allows the collector to register the start time and determine that a flow has survived for a duration longer than the configured active timeout.

NOTE: In active flow monitoring, the cflowd records are exported after a time period that is a multiple of 60 seconds and greater than or equal to the configured active timeout value. For example, if the active timeout value is 90 seconds, the cflowd records are exported at 120-second intervals. If the active timeout value is 150 seconds, the cflowd records are exported at 180-second intervals, and so forth.

- The `flow-inactive-timeout` statement specifies the interval of inactivity for a flow that triggers the flow export. If the interval between the current time and the time that the last packet for this flow was received exceeds the configured inactive timeout value, the flow is allowed to expire.

If the flow stops transmitting for longer than the configured inactive timeout value, the router or switch purges it from the flow table and exports the cflowd record. As a result, the flow is forgotten as far as the PIC is concerned and if the same 5-tuple appears again, it is assigned a new start time and considered a new flow.

Both timers are necessary. The active timeout setting is needed to provide information for flows that constantly transmit packets for a long duration. The inactive timeout setting enables the router or switch to purge flows that have become inactive and that can waste tracking resources.

NOTE: The router must contain an Adaptive Services, Multiservices, or Monitoring Services PIC for the `flow-active-timeout` and `flow-inactive-timeout` statements to take effect.

Example: Configuring Flow Monitoring

The following is an example of flow-monitoring properties configured to support input SONET/SDH interfaces, output monitoring services interfaces, and export to cflowd for flow analysis. To complete the configuration, you also need to configure the interfaces and set up a virtual private network (VPN) routing and forwarding (VRF) instance. For information on cflowd, see ["Enabling Flow Aggregation" on page 577](#).

```
[edit forwarding-options]
monitoring group1 {
  family inet {
    output {
      cflowd 192.168.245.2 port 2055;
      export-format cflowd-version-5;
      flow-active-timeout 60;
      flow-inactive-timeout 30;
      interface mo-4/0/0.1 {
        engine-id 1;
        engine-type 1;
        input-interface-index 44;
        output-interface-index 54;
        source-address 192.168.245.1;
      }
      interface mo-4/1/0.1 {
        engine-id 2;
```

```
        engine-type 1;
        input-interface-index 45;
        output-interface-index 55;
        source-address 192.168.245.1;
    }
    interface mo-4/2/0.1 {
        engine-id 3;
        engine-type 1;
        input-interface-index 46;
        output-interface-index 56;
        source-address 192.168.245.1;
    }
    interface mo-4/3/0.1 {
        engine-id 4;
        engine-type 1;
        input-interface-index 47;
        output-interface-index 57;
        source-address 192.168.245.1;
    }
}
}
```

Release History Table

Release	Description
15.1	Starting with Junos OS Release 15.1, the multiservices PIC management daemon core file is generated when a prolonged flow control failure occurs and when you configure the setting to generate a core dump during prolonged flow control (by using the dump-on-flow-control option with the flow-control-options statement).

RELATED DOCUMENTATION

Active Flow Monitoring Overview	 54
Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers	 637
Configuring Services Interface Redundancy with Flow Monitoring	 72
Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System	 58

Flow Monitoring Output Formats

When you implement passive flow monitoring and active flow monitoring, you should be familiar with flow monitoring formats and fields. Version 5 and version 8 export data into specified fields. Version 9 exports data into templates.

The flow monitoring station monitors the traffic flow and exports the data in flow format to an external server. The Junos OS collects information about the following fields:

- Source and destination IP address
- Total number of bytes and packets sent
- Start and end times of the data flow
- Source and destination port numbers
- TCP flags
- IP protocol and IP type of service
- Originating AS of source and destination address
- Source and destination address prefix mask lengths
- Next-hop router's IP address
- MPLS label (version 9 only)
- ICMP (version 9 only)

Detailed descriptions of the formats are available as follows:

- ["Flow Monitoring Version 5 Format Output Fields" on page 11](#)
- ["Flow Monitoring Version 8 Format Output Fields" on page 16](#)
- ["Flow Monitoring Version 9 Format Output Fields" on page 26](#)

Flow Monitoring Version 5 Format Output Fields

A detailed explanation of version 5 packet formats and fields is shown in the following figures and tables:

- [Figure 1 on page 12](#)

- [Table 1 on page 12](#)
- [Figure 2 on page 13](#)
- [Table 2 on page 14](#)

Figure 1: Version 5 Packet Header Format




Table 1: Export Version 5 Packet Header Fields

Field	Description	Comments
Version	5	-
Count	The number of records in the Protocol Data Unit (PDU) or packet	-
sysUptime	Current time elapsed, in milliseconds, since the router started	-
UNIX seconds	Current seconds since 0000 UTC 1970	NTP synchronized time; the clock on each services PIC is autonomous (200-400 msec jitter) across PICs in a chassis
UNIX nanoseconds	Residual nanoseconds since 0000 UTC 1970	See Comments above for UNIX seconds

Table 1: Export Version 5 Packet Header Fields *(Continued)*

Field	Description	Comments
Flow sequence number	Sequence number of total flows received	–
Engine type	User-configured 8-bit value	Also known as VIP type on other vendors' equipment
Engine ID	User-configured 8-bit value	–

Figure 2: Version 5 Flow-Export Flow Header Format

Byte 3	Byte 2	Byte 1	Byte 0
Source IP address			
Destination IP address			
Next-hop IP address			
Input ifIndex		Output ifIndex	
Packets			
Bytes			
Start time of flow			
End time of flow			
Source port		Destination port	
Padding	TCP flags	IP protocol	TOS
Source AS		Destination AS	
Source mask length	Dest. mask length	Padding	

0003133

g003133

Table 2: Export Version 5 Flow-Export Flow Header Fields

Field	Description	Comments
Source IP address	Source IP address of the flow	–
Destination IP address	Destination IP address of the flow	–
Next-hop IP address	IP address of the router where flows are forwarded	–
Input ifIndex	SNMP index value for the input interface where the router receives flows	<p>Junos OS Release 5.7 and later—Dynamically inserted, but overridden by manual configuration</p> <p>Junos OS Release 5.5—Manually set</p> <p>Junos OS Release 5.4—Set to zero</p>
Output ifIndex	SNMP index value for the output interface where the router forwards flows	<p>Junos OS Release 5.7 and later—Dynamically inserted, but overridden by manual configuration</p> <p>Junos OS Release 5.5—Manually set</p> <p>Junos OS Release 5.4—Set to zero</p>
Packets	Total number of packets received in a flow	–
Bytes	Total number of bytes received in a flow	–
Start time of flow	System up time, in seconds, at the start of the flow	System up time for the services PIC accepting flows
End time of flow	System up time, in seconds, at the end of the flow	System up time for the services PIC accepting flows
Source port	Source application port	–

Table 2: Export Version 5 Flow-Export Flow Header Fields (*Continued*)

Field	Description	Comments
Destination port	Destination application port	The ICMP type is placed in the high-order byte and the ICMP type code is placed in the low-order byte of this field
TCP flags	TCP flags set in the flow	–
IP protocol	IP protocol number	–
TOS	IP type of service	–
Source AS	AS number of the source address	Junos OS Release 5.7 and later—Dynamically inserted if AS information is available
Destination AS	AS number of the destination address	Junos OS Release 5.7 and later—Dynamically inserted if AS information is available
Source mask length	Source address network mask length	–
Dest. mask length	Destination address network mask length	–
Padding	Bytes available to ensure a minimum packet length	–

Useful formulas for flow monitoring are:

- start flow timestamp absolute = $unixTime \times 1000 - (sysUptime - \text{start flow timestamp})$
- end flow timestamp absolute = $unixTime \times 1000 - (sysUptime - \text{end flow timestamp})$

NOTE: In the 2-byte destination port field of the export version 5 flow-export flow format, the following information can be derived:

- High-order byte—ICMP type
- Low-order byte—ICMP type code

For example, if the ICMP type is 3 (00000011 in binary) and the ICMP type code is network unreachable (Type Code 0, or 00000000 in binary), the resulting destination port field value is 00000011 00000000 (768 in decimal).

For more information on ICMP type and type code, see RFC 792 at <http://www.ietf.org>.

Flow Monitoring Version 8 Format Output Fields

A detailed explanation of version 8 packet formats and fields is shown as follows:

- [Figure 3 on page 17](#)
- [Table 3 on page 17](#)
- [Figure 4 on page 19](#)
- [Table 4 on page 19](#)
- [Figure 5 on page 20](#)
- [Table 5 on page 20](#)
- [Figure 6 on page 22](#)
- [Table 6 on page 22](#)
- [Figure 7 on page 24](#)
- [Table 7 on page 24](#)
- [Figure 8 on page 25](#)

- [Table 8 on page 25](#)

Figure 3: Version 8 Template Flow Format

Byte 3	Byte 2	Byte 1	Byte 0
Version		Count	
sysUptime			
UNIX seconds			
UNIX nanoseconds			
Flow Sequence Number			
Engine type	Engine ID	Aggregation method	Aggregation version
Reserved			

g003076

Byte 3	Byte 2	Byte 1	Byte 0
Version		Count	
sysUptime			
UNIX seconds			
UNIX nanoseconds			
Flow Sequence Number			
Engine type	Engine ID	Aggregation method	Aggregation version
Reserved			

g003076

Table 3: Version 8 Flow Template Fields

Field	Description
Version	8
Count	The number of records in the protocol data unit (PDU) or packet

Table 3: Version 8 Flow Template Fields *(Continued)*

Field	Description
sysUptime	Current time elapsed, in milliseconds, since the router started
UNIX seconds	Current seconds since 0000 UTC 1970
UNIX nanoseconds	Residual nanoseconds since 0000 UTC 1970
Flow sequence number	Sequence counter of total flows received
Engine type	Type of flow switching engine
Engine ID	ID number of the flow switching engine
Aggregation method	Aggregation method used
Aggregation version	Version of the aggregation export
Reserved	Empty field reserved for future usage

Figure 4: Version 8 AS Aggregation Flow Entry Format

Byte 3	Byte 2	Byte 1	Byte 0
Flows			
Packets			
Bytes			
Start Time of Flow			
End Time of Flow			
Source AS		Destination AS	
Input interface		Output interface	

9003077

Table 4: Version 8 AS Aggregation Flow Entry Fields

Field	Description
Flows	Total number of flows
Packets	Total number of packets received in a flow
Bytes	Total number of bytes received in a flow
Start time of flow	System up time, in seconds, at the start of the flow
End time of flow	System up time, in seconds, at the end of the flow
Source AS	AS number of the source address
Destination AS	AS number of the destination address
Input interface	SNMP index value for the input interface where the router receives flows

Table 4: Version 8 AS Aggregation Flow Entry Fields *(Continued)*

Field	Description
Output interface	SNMP index value for the output interface where the router forwards flows

Figure 5: Version 8 Protocol/Port Aggregation Flow Entry Format

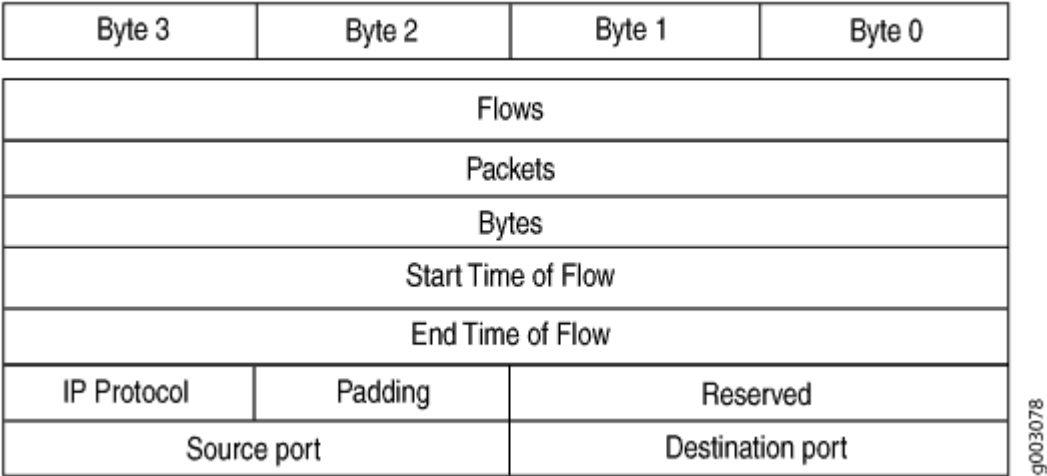


Table 5: Version 8 Protocol/Port Aggregation Flow Entry Fields

Field	Description
Flows	Total number of flows
Packets	Total number of packets received in a flow
Bytes	Total number of bytes received in a flow
Start time of flow	System up time, in seconds, at the start of the flow
End time of flow	System up time, in seconds, at the end of the flow

Table 5: Version 8 Protocol/Port Aggregation Flow Entry Fields (*Continued*)

Field	Description
IP protocol	IP protocol number
Padding	Bytes available to ensure a minimum packet length
Reserved	Empty field reserved for future usage
Source port	Source application port
Destination port	Destination application port

Figure 6: Version 8 Prefix Aggregation Flow Entry Format

Byte 3		Byte 2		Byte 1		Byte 0	
Flows							
Packets							
Bytes							
Start Time of Flow							
End Time of Flow							
Source prefix							
Destination prefix							
Source Mask Length		Dest. Mask Length		Reserved			
Source AS				Destination AS			
Input interface				Output interface			

g003079

Byte 3		Byte 2		Byte 1		Byte 0	
Flows							
Packets							
Bytes							
Start Time of Flow							
End Time of Flow							
Source prefix							
Destination prefix							
Source Mask Length		Dest. Mask Length		Reserved			
Source AS				Destination AS			
Input interface				Output interface			

g003079

Table 6: Version 8 Prefix Aggregation Flow Entry Fields

Field	Description
Flows	Total number of flows

Table 6: Version 8 Prefix Aggregation Flow Entry Fields (*Continued*)

Field	Description
Packets	Total number of packets received in a flow
Bytes	Total number of bytes received in a flow
Start time of flow	System up time, in seconds, at the start of the flow
End time of flow	System up time, in seconds, at the end of the flow
Source prefix	Source IP address prefix
Destination prefix	Destination IP address prefix
Source mask length	Source address network mask length
Dest. mask length	Destination address network mask length
Reserved	Empty field reserved for future usage
Source AS	AS number of the source address
Destination AS	AS number of the destination address
Input interface	SNMP index value for the input interface where the router receives flows
Output interface	SNMP index value for the output interface where the router forwards flows

Figure 7: Version 8 Source Prefix Aggregation Flow Entry Format

Byte 3		Byte 2		Byte 1		Byte 0	
Flows							
Packets							
Bytes							
Start Time of Flow							
End Time of Flow							
Source prefix							
Source Mask Length		Padding		Source AS			
Input interface				Reserved			

9003080

Table 7: Version 8 Source Prefix Aggregation Flow Entry Fields

Field	Description
Flows	Total number of flows
Packets	Total number of packets received in a flow
Bytes	Total number of bytes received in a flow
Start time of flow	System up time, in seconds, at the start of the flow
End time of flow	System up time, in seconds, at the end of the flow
Source prefix	Source IP address prefix
Source mask length	Source address network mask length
Padding	Bytes available to ensure a minimum packet length
Source AS	AS number of the source address

Table 7: Version 8 Source Prefix Aggregation Flow Entry Fields *(Continued)*

Field	Description
Input interface	SNMP index value for the input interface where the router receives flows
Reserved	Empty field reserved for future usage

Figure 8: Version 8 Destination Prefix Aggregation Flow Entry Format




Table 8: Version 8 Destination Prefix Aggregation Flow Entry Fields

Field	Description
Flows	Total number of flows
Packets	Total number of packets received in a flow
Bytes	Total number of bytes received in a flow
Start time of flow	System up time, in seconds, at the start of the flow

Table 8: Version 8 Destination Prefix Aggregation Flow Entry Fields (*Continued*)

Field	Description
End time of flow	System up time, in seconds, at the end of the flow
Destination prefix	Destination IP address prefix
Dest. mask length	Destination address network mask length
Padding	Bytes available to ensure a minimum packet length
Destination AS	AS number of the destination address
Output interface	SNMP index value for the output interface where the router forwards flows
Reserved	Empty field reserved for future usage

For more information about version 5 and version 8 packet formats and fields, see <http://www.caida.org>.

Flow Monitoring Version 9 Format Output Fields

IN THIS SECTION

- [IPFIX \(Version 10\) IPv4 Fields | 38](#)

A detailed explanation of active flow monitoring version 9 packet formats and fields is shown as follows:

- [Table 9 on page 27](#)
- [Figure 9 on page 30](#)
- [Table 10 on page 31](#)

- [Figure 11 on page 35](#)
- [Table 10 on page 31](#)
- [Figure 12 on page 36](#)
- [Table 14 on page 36](#)
- [Figure 13 on page 37](#)
- [Table 15 on page 38](#)

The Junos OS supports the version 9 template formats:

Table 9: Flow Monitoring Version 9 Template Formats

Template	Fields
IPv4	<div>Flow selectors:</div> <ul style="list-style-type: none">• Source and destination IP address• Source and destination address prefix mask lengths• Source and destination port numbers• IP protocol and IP type of service• ICMP type <div>Flow nonselectors:</div> <ul style="list-style-type: none">• TCP flags• Input and output SNMP• Input bytes• Input packets• Start time• End time

Table 9: Flow Monitoring Version 9 Template Formats *(Continued)*

Template	Fields
MPLS	<p>Flow selectors:</p> <ul style="list-style-type: none"> • MPLS label 1 • MPLS label 2 • MPLS label 3 <p>Flow nonselectors:</p> <ul style="list-style-type: none"> • Input and output SNMP • Input bytes • Input packets • Start time • End time
MPLS_IPv4	<p>Flow selectors:</p> <ul style="list-style-type: none"> • MPLS label 1 • MPLS label 2 • MPLS label 3 • MPLS top-level FEC address <p>Flow nonselectors:</p> <ul style="list-style-type: none"> • Input and output SNMP • Input bytes • Input packets • Start time • End time

Table 9: Flow Monitoring Version 9 Template Formats *(Continued)*

Template	Fields
IPv6	<p>Flow selectors:</p> <ul style="list-style-type: none"> • IP protocol and IP type of service • Source and destination port numbers • Input SNMP • Source and destination IPv6 address • ICMP type <p>Flow nonselectors:</p> <ul style="list-style-type: none"> • Input bytes • Input packets • TCP flags • Output SNMP • Source and destination autonomous system • Last and first switched • IPv6 source and destination mask • IP protocol version • IPv6 next hop

Table 9: Flow Monitoring Version 9 Template Formats *(Continued)*

Template	Fields
Peer AS billing	<div>Flow selectors:<ul style="list-style-type: none">IPv4 class of serviceIngress interface informationBGP peer destination AS numberBGP IPv4 next hop address</div> <div>Flow nonselectors<ul style="list-style-type: none">Input and output SNMPInput bytesInput packetsFirst switchLast switched</div> <div>NOTE: Peer AS billing traffic is not supported for active flow monitoring version 9 configuration on PTX5000 routers tethered to CSE2000.</div>

Figure 9: Version 9 Flow Header Format

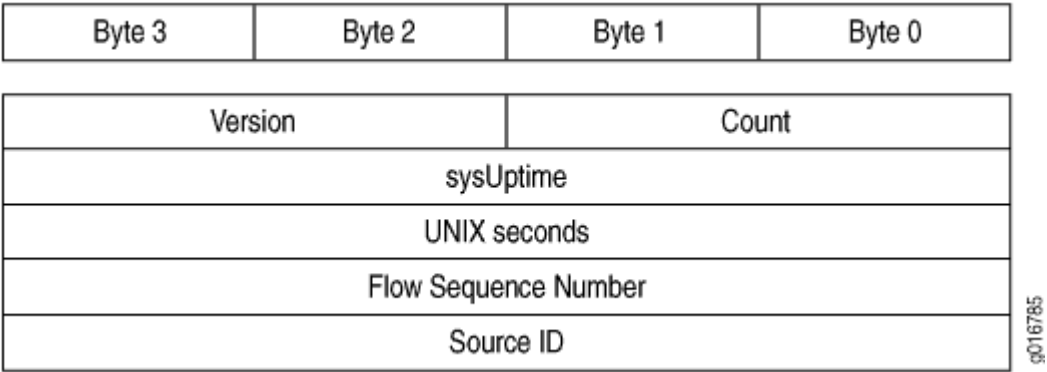


Table 10: Version 9 Flow Header Fields

Field	Description
Version	9
Count	Total number of records in the protocol data unit (PDU) or packet. This number includes all of the options FlowSet records, template FlowSet records, and data FlowSet records.
sysUptime	Current time elapsed, in milliseconds, since the router started.
UNIX seconds	Current seconds since 0000 UTC 1970.
Flow sequence number	Sequence counter of total flows received.
Source ID	32-bit value that identifies the data exporter. Version 9 uses the integrated field diagnostics (IFD) SNMP index of the PIC or device that is exporting the data flow. This field is equivalent to engine type and engine ID fields found in versions 5 and 8.

Figure 10: Version 9 Template FlowSet Format

Byte 3	Byte 2	Byte 1	Byte 0
Flowset ID = 0		Length	
Template ID 256		Field Count	
Field Type 1		Field Length 1	
Field Type 2		Field Length 2	
...		...	
Field Type N		Field Type N	
Template ID 257		Field Count	
Field Type 1		Field Length 1	

g016786

Table 11: Version 9 Template FlowSet Fields

Field	Description
FlowSet ID	FlowSet type. FlowSet ID 0 is reserved for the Template FlowSet.
Length	FlowSet length. Individual template FlowSets might contain multiple template records, which means that the length of template FlowSets varies.
Template ID	Unique template ID assigned to each newly generated template. Templates numbered 256 and higher define data formats. Templates numbered 0 through 255 define FlowSet IDs.
Field Count	Fields in the template record. This field allows the collector to determine the end of the current template record and the start of the next.
Field Type	Field type. These are defined in Table 12 on page 33 .

Table 11: Version 9 Template FlowSet Fields (*Continued*)

Field	Description
Field Length	Length, in bytes, of the corresponding field type.

Table 12: Field Type Definitions Supported in Junos OS

Field Type	Description
1	IN_BYTES: The number of bytes associated with an IP flow. By default, the length is 4 bytes.
2	IN_PKTS: The number of packets associated with an IP flow. By default, the length is 4 packets.
4	PROTOCOL: The IP protocol byte.
5	TOS: The type-of-service byte setting of an incoming packet.
6	TCP_FLAGS: The cumulative TCP flags associated with a flow.
7	L4_SRC_PORT: The TCP/UDP source port.
8	IPv4_SRC_ADDR: The IPv4 source address.
9	SRC_MASK: The number of contiguous bits in the source subnet mask.
10	INPUT_SNMP: The IFD SNMP input interface index. By default, the length is 2.
11	L4_DST_PORT: The TCP/UDP destination port number.
12	IPv4_DST_ADDR: The IPv4 destination address.
13	DST_MASK: The number of contiguous bits in the destination subnet mask.

Table 12: Field Type Definitions Supported in Junos OS (Continued)

Field Type	Description
14	OUTPUT_SNMP: The IFD SNMP output interface index. By default, the length is 2.
16	SRC_AS: The source autonomous system number. This is always set to zero.
17	DST_AS: The destination autonomous system number. This is always set to zero.
18	BGP_IPV4_NEXT_HOP: The BGP IPv4 next-hop address.
21	LAST_SWITCHED: The uptime of the device (in milliseconds) at which the last packet of the flow was switched.
22	FIRST_SWITCHED: The uptime of the device (in milliseconds) at which the first packet of the flow was switched.
29	IPv6_SRC_MASK: The length of the IPv6 source mask, in contiguous bits.
30	IPv6_DST_MASK: The length of the IPv6 destination mask, in contiguous bits.
32	ICMP_TYPE: The ICMP type.
34	SAMPLING_INTERVAL: The rate at which packets are sampled. As an example, a rate of 100 means that one packet is sampled for every 100 packets in the data flow.
35	SAMPLING_ALGORITHM: The type of algorithm being used. 0x01 indicates deterministic sampling and 0x02 indicates random sampling.
47	MPLS_TOP_LABEL_IP_ADDRESS: The MPLS top- label address.
60	IP_PROTOCOL_VERSION: The IP protocol version being used.
62	IPv6_NEXT_HOP: The IPv6 address of the next-hop router.

Table 12: Field Type Definitions Supported in Junos OS *(Continued)*

Field Type	Description
70	MPLS_LABEL_1: The first MPLS label in the stack.
71	MPLS_LABEL_2: The second MPLS label in the stack.
72	MPLS_LABEL_3: The third MPLS label in the stack.
128	DST_PEER_AS: The destination of the BGP peer AS.

Figure 11: Version 9 Data FlowSet Format

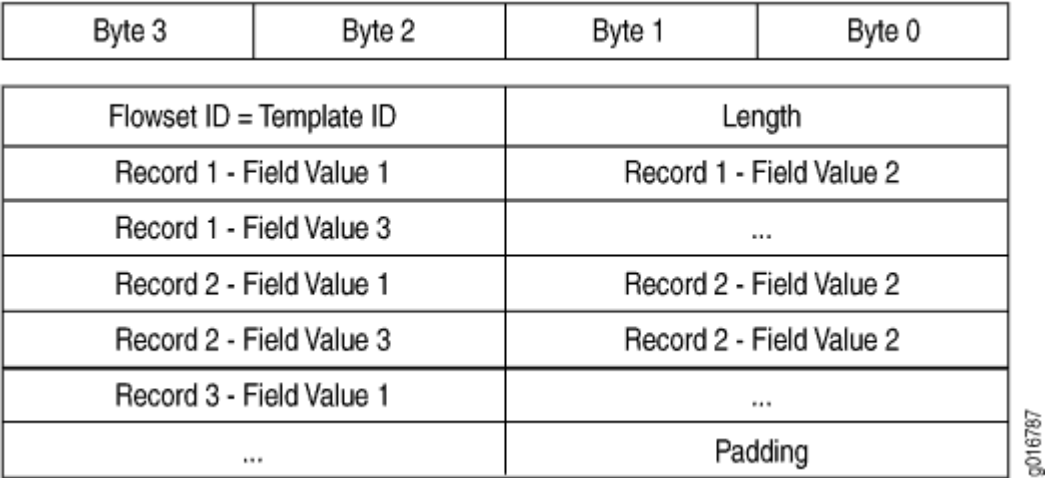


Table 13: Version 9 Data FlowSet Format

Field	Description
FlowSet ID = Template ID	Data FlowSet that associated with a FlowSet ID. The FlowSet ID maps to a previously generated template ID. The flow collector must use the FlowSet ID to find the corresponding template record and decode the flow records from the FlowSet.

Table 13: Version 9 Data FlowSet Format *(Continued)*

Field	Description
Length	FlowSet length. Data FlowSets are fixed in length.
Record Number - Field Value Number	Flow data records, each containing a set of field values. The template record identified by the FlowSet ID dictates the type and length of the field values.
Padding	Bytes (in zeros) that the exporter inserts so that the subsequent FlowSet starts at a 4-byte aligned boundary.

Figure 12: Version 9 Options Template Format

Byte 3	Byte 2	Byte 1	Byte 0
Flowset ID = 1		Length	
Template ID		Option Scope Length	
Option Length		Scope 1 Field Type	
Scope 1 Field Length		...	
Scope N Field Length		Option 1 Field Type	
Option 1 Field Length		...	
Option M Field Length		Padding	

g016788

Table 14: Version 9 Options Template Format

Field	Description
FlowSet ID	FlowSet type. FlowSet ID 1 is reserved for the options template.
Length	FlowSet length. Option template FlowSets are fixed in length.

Table 14: Version 9 Options Template Format *(Continued)*

Field	Description
Template ID	Template ID of the options template. Options template values are greater than 255.
Option Scope Length	Length, in bytes, of any scope field definition that is part of the options template record.
Scope 1 Field Type	Relevant process. The Junos OS supports the system process (1).
Scope 1 Field Length	Length, in bytes, of the option field.
Padding	Bytes the exporter inserts so that the subsequent FlowSet starts at a 4-byte aligned boundary.

Figure 13: Active Flow Monitoring Version 9 Options Data Record Format

Byte 3	Byte 2	Byte 1	Byte 0
Flowset ID = Template ID		Length	
Record 1 - Scope 1 Value		Record 1 - Option Field 1 Value	
Record 1 - Option Field 2 Value		...	
Record 2 - Option Field 2 Value		...	
Record 3 - Scope 1 Value		Record 3 - Option Field 1 Value	
...		Padding	

g016789

Table 15: Active Flow Monitoring Version 9 Options Data Record Format

Field	Description
FlowSet ID = Template ID	ID that precedes each options data flow record. The FlowSet ID maps to a previously generated template ID. The collector must use the FlowSet ID to find the corresponding template record and decode the options data flow records from the FlowSet.
Length	FlowSet length. Option FlowSets are fixed in length.
Number of Flow Data Records	Remainder of the options data FlowSet is a collection of flow data records, each containing a set of field values. The template record identified by the FlowSet ID dictates the type and length of the field values.
Padding	Bytes (in zeros) the exporter inserts so that the subsequent FlowSet starts at a 4-byte aligned boundary.

IPFIX (Version 10) IPv4 Fields

Field Name	Flow Key	Element ID	Length in Bytes
IPV4_SADDR	Y	8	4
IPV4_DADDR	Y	12	4
IPV4_TOS	Y	5	1
IPV4_PROTO	Y	4	1
TCP_UDP_SPORT	Y	7	2
TCP_UDP_DPORT	Y	11	2

(Continued)

Field Name	Flow Key	Element ID	Length in Bytes
IMCP_TYPE_CODE_IPV4	Y	32	2
IIF	Y	10	4
VLAN_ID	Configurable	58	2
IPV4_SMASK	N	9	1
IPV4_DMASK	N	13	1
SRC_AS	N	16	4
DST_AS	N	17	4
IPV4_NEXTHOP	N	15	4
TCP_FLAGS	N	6	1
OIF	N	14	4
FLOW_BYTES	N	1	8
FLOW_PACKETS	N	2	8
MIN_TTL	N	52	1
MAX_TTL	N	53	1
START_TIME	N	152	8

(Continued)

Field Name	Flow Key	Element ID	Length in Bytes
END_TIME	N	153	8
FIRST_SWITCHED	N	22	4
LAST_SWITCHED	N	21	4
FLOW_END_REASON	N	136	1
IP_PROTOCOL_VERSION	N	60	1
BGP_NEXTHOP_ID	N	18	4
FLOW_DIRECTION	Configurable	61	1
DOT_1Q_VLAN_ID	N	243	2
Dot_1Q_CUSTOMER_VLAN_ID	N	245	2
IP IDENTIFIER	N	54	4

Monitoring Traffic Using Active Flow Monitoring

IN THIS CHAPTER

- [Configuring Active Flow Monitoring | 42](#)
- [Active Flow Monitoring System Requirements | 45](#)
- [Active Flow Monitoring Applications | 46](#)
- [Active Flow Monitoring PIC Specifications | 49](#)
- [Active Flow Monitoring Overview | 53](#)
- [Active Flow Monitoring Overview | 54](#)
- [Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System | 58](#)
- [Example: Configuring Flow Monitoring on an MX Series Router with MS-MIC and MS-MPC | 62](#)
- [Configuring Services Interface Redundancy with Flow Monitoring | 72](#)
- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)
- [Configuring Flow Offloading on MX Series Routers | 84](#)
- [Configuring Active Flow Monitoring on PTX Series Packet Transport Routers | 85](#)
- [Configuring Actively Monitored Interfaces on M, MX and T Series Routers | 88](#)
- [Collecting Flow Records | 88](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with an Accounting Group | 89](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with a Sampling Group | 90](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with a Template | 92](#)
- [Defining a Firewall Filter on M, MX and T Series Routers to Select Traffic for Active Flow Monitoring | 93](#)
- [Processing IPv4 traffic on an M, MX or T Series Router Using Monitoring services, Adaptive services or Multiservices Interfaces | 95](#)
- [Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers | 96](#)
- [Replicating Version 9 Flow Aggregation From M, MX and T Series Routers to Multiple Flow Servers | 97](#)
- [Configuring Routing Engine-Based Sampling on M, MX and T Series Routers for Export to Multiple Flow Servers | 98](#)
- [Example: Copying Traffic to a PIC While an M, MX or T Series Router Forwards the Packet to the Original Destination | 99](#)
- [Configuring an Aggregate Export Timer on M, MX and T Series Routers for Version 8 Records | 126](#)

- [Example: Sampling Configuration for M, MX and T Series Routers | 127](#)
- [Associating Sampling Instances for Active Flow Monitoring with a Specific FPC, MPC, or DPC | 132](#)
- [Example: Sampling Instance Configuration | 133](#)
- [Example: Sampling and Discard Accounting Configuration on M, MX and T Series Routers | 141](#)

Configuring Active Flow Monitoring

In active flow monitoring, the router participates in both the monitoring application and in the normal routing functionality of the network. Although the Monitoring Services PIC was designed initially for use as an offline passive flow monitoring tool, it can also be used in an active flow monitoring topology.

[Table 16 on page 42](#) shows which Juniper Networks PICs and corresponding routers support active flow monitoring. For more information on Juniper Networks PICs, see the PIC guide that corresponds to your router.

Table 16: Passive and Active Flow Monitoring PIC Support

PIC Type and Service	M5/M10	M7i/M10i	M20	M40e	M120	M160	T Series/M320	TX Matrix
Monitoring Services PIC: active flow monitoring	Yes (version 8 only)	Yes	Yes	Yes	No	Yes (version 8 only)	No	No
Monitoring Services II PIC: flow collection services	No	No	No	Yes	No	Yes (version 8 only)	No	No
Adaptive Services PIC: active flow monitoring	Yes (version 8 only)	Yes	Yes	Yes	No	Yes (version 8 only)	No	No

Table 16: Passive and Active Flow Monitoring PIC Support *(Continued)*

PIC Type and Service	M5/M10	M7i/M10i	M20	M40e	M120	M160	T Series/M320	TX Matrix
Adaptive Services II PIC: active flow monitoring	Yes (version 8 only)	Yes	Yes	Yes	Yes	Yes (version 8 only)	Yes	Yes
Adaptive Services II PIC: flow-tap services	No	Yes	Yes	Yes	Yes	No	Yes	No
MultiServices 100 PIC: active flow monitoring	No	Yes	No	Yes	No	No	Yes	Yes
MultiServices 400 PIC: active flow monitoring	No	No	No	Yes	Yes	No	Yes	Yes
MultiServices 500 PIC: active flow monitoring	No	No	No	Yes	Yes	No	Yes	Yes
Junos OS-enabled active flow monitoring	No	No	No	No	No	No	No	No

Specified packets can be filtered and sent to the monitoring interface. For the Monitoring Services PIC, the interface name contains the **mo-** prefix. For the Adaptive Services PICs and MultiServices PICs, the interface name contains the **sp-** prefix.

NOTE: If you upgrade from the Monitoring Services PIC to the Adaptive Services PIC or MultiServices PIC for active flow monitoring, you must modify the interface name of your monitoring interface from **mo-fpc/pic/port** to **sp-fpc/pic/port**.

The major active flow monitoring actions you can configure at the `[edit forwarding-options]` hierarchy level are as follows:

- Sampling, with the `[edit forwarding-options sampling]` hierarchy. This option extracts limited information (such as the source and destination IP address) from a copy of some of the packets in a flow, while the original packets are forwarded to the intended destination. This option is extended to define active sampling on a per Packet Forwarding Engine basis by defining a sampling instance that specifies a name for the sampling parameters and binding the instance to the particular Packet Forwarding Engine.
- Templates, with the `[edit forwarding-options sampling]` and `[edit services monitoring]` hierarchies. With active flow monitoring support for version 5, version 8, and the customizing version 9, you can use templates to organize the data gathered from sampling.
- Discard accounting, with the `[edit forwarding-options accounting]` hierarchy. This option quarantines unwanted packets, creates flow monitoring records that describe the packets, and discards the packets instead of forwarding them.
- *Port mirroring*, with the `[edit forwarding-options port-mirroring]` hierarchy. This option makes one full copy of all packets in a flow and delivers the copy to a single destination.
- Multiple port mirroring, with the `[edit forwarding-options next-hop-group]` hierarchy. This option allows multiple copies of selected traffic to be delivered to multiple destinations. (Multiple port mirroring requires a Tunnel Services PIC.)
- Flow-tap services processing, with the `[edit services flow-tap]` hierarchy. This option sends copies of packets that match dynamic filter criteria to one or more content destinations.

Unlike passive flow monitoring, you do not need to configure a monitoring group. Instead, you can send filtered packets to a monitoring services or adaptive services interface (**mo-** or **sp-**) by using sampling or discard accounting. Optionally, you can configure port mirroring or multiple port mirroring to direct packets to additional interfaces.

These active flow monitoring options provide a wide variety of actions that can be performed on network traffic flows. However, the following restrictions apply:

- The router can perform either sampling *or* port mirroring at any one time.
- The router can perform either forwarding *or* discard accounting at any one time.

Because the Monitoring Services PIC, Adaptive Services PIC, and MultiServices PIC allow only one action to be performed at any one time, the following configuration options are available:

- Sampling and forwarding
- Sampling and discard accounting
- Port mirroring and forwarding

- Port mirroring and discard accounting
- Sampling and port mirroring on different sets of traffic

To configure active flow monitoring, complete these steps:

- ["Defining a Firewall Filter on M, MX and T Series Routers to Select Traffic for Active Flow Monitoring" on page 93](#)
- ["Configuring Actively Monitored Interfaces on M, MX and T Series Routers " on page 88](#)
- ["Processing IPv4 traffic on an M, MX or T Series Router Using Monitoring services, Adaptive services or Multiservices Interfaces" on page 95](#)
- ["Collecting Flow Records" on page 88](#)
- [Rerouting Packets on an M, MX or T Series Router with Port Mirroring](#)
- [Option: Configuring Port Mirroring with Filter-Based Forwarding and a Monitoring Group](#)
- [Sending Port-Mirrored Traffic from an M, MX or T Series Router to Multiple Export Interfaces by Using Next-Hop Groups](#)
- ["Sending Packets to a Mediation Device on MX, M and T Series Routers " on page 309](#)

Active Flow Monitoring System Requirements

To implement active flow monitoring, your system must meet these minimum requirements:

- Junos 10.4 or later for peer AS billing support on flow monitoring version 9
- Junos 9.3R2 or later for IPv6 support on flow monitoring version 9
- Junos 9.3R2 or later for multiple flows for flow monitoring version 9
- Junos OS Release 9.0 or later for version 9 flow aggregation to multiple flow servers
- Junos OS Release 8.5 or later for active flow monitoring support on MultiServices 500 PICs
- Junos OS Release 8.3 or later for flow monitoring version 9 support, MPLS support, and active flow monitoring support on MultiServices 100 and 400 PICs
- Junos OS Release 8.2 or later for M120 router support and for flow monitoring version 5 and 8 support on MultiServices 100 and 400 PICs
- Junos OS Release 8.1 or later for the flow-tap services application on Adaptive Services II PICs installed in M7i, M10i, M20, M40e, M320, and T Series routers

- Junos OS Release 7.4 or later for port mirroring of IPv6 packets
- Junos OS Release 7.3 or later for active flow monitoring on Adaptive Services II PICs installed in TX Matrix platforms
- Junos OS Release 7.0 or later for active flow monitoring on Adaptive Services II PICs installed in T Series and M320 routers
- Junos OS Release 6.0 or later for the Adaptive Services PIC
- Junos OS Release 5.7 or later for the automatic insertion of AS numbers and SNMP index values for input and output interfaces into records, port mirroring to multiple ports, and discard accounting
- Junos OS Release 5.6 or later for the Monitoring Services PIC
- M5, M7i, M10, M10i, M20, M40e, M120, M160, M320, or T Series router with an Internet Processor II ASIC or later
- Type 1 enhanced FPCs
- Two M Series or T Series PICs of your choice: One to receive incoming traffic and one to forward outgoing traffic (the second PIC or PIM is not necessary for discard accounting)
- Export PICs to connect to the collector or packet analyzer
- Tunnel Services PIC (required for multiple port mirroring or **mo-** interface load balancing)
- Flow collector version 5, 8, or 9
- ES PIC and packet analyzers (optional)

RELATED DOCUMENTATION

[Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers](#) | 152

[Active Flow Monitoring PIC Specifications](#) | 49

Active Flow Monitoring Applications

Flow monitoring can be used for many different reasons such as network planning, accounting, usage-based network billing, security, and monitoring for Denial-of-Service attacks.

Some examples of the types of things you can use flow monitoring for are:

- Tracking what kind of traffic is entering or exiting an ISP or corporate network.


- Tracking traffic flows between BGP autonomous systems.
- Tracking traffic flows between enterprise network regions.
- Taking a snapshot of the existing quality-of-service (QoS) policy results prior to making changes in QoS policy in case you need to roll back changes later in the process.
- Verifying that load balancing techniques are performing as intended.
- Capturing a base line of current network performance prior to making changes intended to improve performance so that you know if the changes are helping.
- Discovering if network users at an enterprise are using bandwidth for work-related activities or for non work-related activities.

Examples of how flow monitoring helps with network administration include the following:

- A large service provider uses active flow monitoring on its core uplinks as a way to collect data on the protocols in use, packet sizes, and flow durations to better understand the usage of its Internet service offering. This helps the provider understand where network growth is coming from.
- Service providers bill customers for the data sent or bandwidth used by sending captured flow data to third-party billing software.
- At a large enterprise, VoIP users at a remote site complained of poor voice quality. The flow monitoring reports showed that the VoIP traffic did not have the correct type of service settings.
- Users on an enterprise network, reported network slowdowns. The flow monitoring reports showed that one user's PC was generating a large portion of the network traffic. The PC was infected with malware.
- A growing enterprise planned to deploy new business management software and needed to know what type of network bandwidth demand the new software would create. During the software trial period, flow monitoring reports were used to identify the expected increase in traffic.

Thus, while flow monitoring is traditionally associated with traffic analysis, it also has a role in accounting and security.

Figure 14: Active Flow Monitoring



RELATED DOCUMENTATION

[Flow Monitoring Overview](#)

[Active Flow Monitoring Overview](#) | 53

Active Flow Monitoring PIC Specifications

For Monitoring Services PIC specifications, see [Table 17 on page 49](#) and [Table 18 on page 49](#). For Adaptive Services PIC specifications, see [Table 19 on page 50](#). For MultiServices PIC specifications, see [Table 20 on page 51](#) and [Table 21 on page 52](#).

Table 17: Monitoring Services PIC Specifications

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	DB-9 diagnostic serial console port
Status LED	One tricolor: <ul style="list-style-type: none"> • Off—The PIC is offline; it is safe to remove it from the chassis. • Green—The PIC is operating normally. • Amber—The PIC is initializing. • Red—The PIC has an error or failure; no further harm can be done by removing it from the chassis.
Application LED	One tricolor: <ul style="list-style-type: none"> • Off—The service is not running. • Green—The service is running under acceptable load. • Amber—The service is overloaded.

Table 18: Monitoring Services II PIC Specifications

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	N/A

Table 18: Monitoring Services II PIC Specifications (Continued)

Specification	Description
Status LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The PIC is offline; it is safe to remove it from the chassis. • Green—The PIC is operating normally. • Amber—The PIC is initializing. • Red—The PIC has an error or failure; no further harm can be done by removing it from the chassis.
Application LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The flow collector is not running. • Green—The flow collector is running under acceptable load. • Amber—The flow collector is overloaded.

Table 19: Adaptive Services PIC Specifications

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	N/A
Status LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The PIC is offline; it is safe to remove it from the chassis. • Green—The PIC is operating normally. • Amber—The PIC is initializing. • Red—The PIC has an error or failure; no further harm can be done by removing it from the chassis.

Table 19: Adaptive Services PIC Specifications (Continued)

Specification	Description
Application LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The flow collector is not running. • Green—The flow collector is running under acceptable load. • Amber—The flow collector is overloaded.

Table 20: MultiServices 100 PIC

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	N/A
Status LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The PIC is offline; it is safe to remove it from the chassis. • Green—The PIC is operating normally. • Amber—The PIC is initializing. • Red—The PIC has an error or failure; no further harm can be done by removing it from the chassis.
Application LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The service is not running. • Green—The service is running under acceptable load. • Amber—The service is overloaded.

Table 21: MultiServices 400 PIC

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	N/A
Status LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The PIC is offline; it is safe to remove it from the chassis. • Green—The PIC is operating normally. • Amber—The PIC is initializing. • Red—The PIC has an error or failure; no further harm can be done by removing it from the chassis.
Application LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The service is not running. • Green—The service is running under acceptable load. • Amber—The service is overloaded.

Table 22: MultiServices 500 PIC

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	N/A

Table 22: MultiServices 500 PIC (Continued)

Specification	Description
Status LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The PIC is offline; it is safe to remove it from the chassis. • Green—The PIC is operating normally. • Amber—The PIC is initializing. • Red—The PIC has an error or failure; no further harm can be done by removing it from the chassis.
Application LED	<p>One tricolor:</p> <ul style="list-style-type: none"> • Off—The service is not running. • Green—The service is running under acceptable load. • Amber—The service is overloaded.

RELATED DOCUMENTATION

[Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers | 152](#)

[Active Flow Monitoring System Requirements | 45](#)

Active Flow Monitoring Overview

Flow monitoring versions 5, 8, and 9 support active flow monitoring. For active flow monitoring, the monitoring station participates in the network as an active router. The major actions the router can perform during active flow monitoring are as follows:

- Sampling—The router selects and analyzes only a portion of the traffic.
- Sampling with templates—The router selects, analyzes, and arranges a portion of the traffic into templates.
- Sampling per sampling instance—The router selects, analyzes, and arranges a portion of the traffic according to the configuration and binding of a sampling instance.

- *Port mirroring*—The router copies entire packets and sends the copies to another interface.
- Multiple port mirroring—The router sends multiple copies of monitored packets to multiple export interfaces with the `next-hop-group` statement at the `[edit forwarding-options]` hierarchy level.
- Discard accounting—The router accounts for selected traffic before discarding it. Such traffic is not forwarded out of the router. Instead, the traffic is quarantined and deleted.
- Flow-tap processing—The router processes requests for active flow monitoring dynamically by using the Dynamic Tasking Control Protocol (DTCP).

RELATED DOCUMENTATION

[Flow Monitoring Overview](#)

[Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers](#) | 156

Active Flow Monitoring Overview

Using a Juniper Networks M Series Multiservice Edge or T Series Core router or EX9200, a selection of PICs (including the Monitoring Services PIC, Adaptive Services [AS] PIC, Multiservices PIC, or Multiservices DPC) and other networking hardware, you can monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to do the following:

- Gather and export detailed information about IP version 4 (IPv4) traffic flows between source and destination nodes in your network.
- Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
- Perform discard accounting on an incoming traffic flow.
- Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
- Direct filtered traffic to different packet analyzers and present the data in its original format (port mirror).

NOTE: Monitoring Services PICs, AS PICs, and Multiservices PICs must be mounted on an Enhanced Flexible PIC Concentrator (FPC) in an M Series or T Series router.

Multiservices DPCs installed in Juniper Networks MX Series 3D Universal Edge routers support the same functionality, with the exception of the passive monitoring and flow-tap features.

Although the Monitoring Services PIC was designed initially for use as an offline passive flow monitoring tool, it can also be used in an active flow monitoring topology. In contrast, the AS or Multiservices PIC is designed exclusively for active flow monitoring. To use either the Monitoring Services PIC, AS PIC, or Multiservices PIC for active flow monitoring, you must install the PIC in an M Series or T Series router. The router participates in both the monitoring application and in the normal routing functionality of the network.

Starting with Junos OS Release 11.4, support for active monitoring is extended to logical systems running on T Series and MX Series routers. A logical system is a partition created from a physical router that performs independent routing tasks. Several logical systems in a single router with their own interfaces, policies, instances, and routing tables can perform functions handled by several different routers. A shared services PIC handles flows from all the logical systems. Only version 9 flows, IPv4, and MPLS templates are supported. See ["Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System "](#) on page 58 for a sample configuration that enables active monitoring on a logical system.

Specified packets can be filtered and sent to the monitoring interface. For the Monitoring Services PIC, the interface name contains the `mo-` prefix. For the AS or Multiservices PIC, the interface name contains the `sp-` prefix.

NOTE: If you upgrade from the Monitoring Services PIC to the Adaptive Services or Multiservices PIC for active flow monitoring, you must change the name of your monitoring interface from `mo-fpc/pic/port` to `sp-fpc/pic/port`.

The major active flow monitoring actions you can configure at the `[edit forwarding-options]` hierarchy level are as follows:

- Sampling, with the `[edit forwarding-options sampling]` hierarchy. This option sends a copy of the traffic stream to an AS or Monitoring Services PIC, which extracts limited information (such as the source and destination IP address) from some of the packets in a flow. The original packets are forwarded to the intended destination as usual.
- Discard accounting, with the `[edit forwarding-options accounting]` hierarchy. This option quarantines unwanted packets, creates cflowd records that describe the packets, and discards the packets instead of forwarding them.

- *Port mirroring*, with the [edit forwarding-options port-mirroring] hierarchy. This option makes one full copy of all packets in a flow and delivers the copy to a single destination. The original packets are forwarded to the intended destination.
- Multiple port mirroring, with the [edit forwarding-options next-hop-group] hierarchy. This option allows multiple copies of selected traffic to be delivered to multiple destinations. (Multiple port mirroring requires a Tunnel Services PIC.)

Unlike passive flow monitoring, you do not need to configure a monitoring group. Instead, you can send filtered packets to a monitoring services or adaptive services interface (mo- or sp-) by using sampling or discard accounting. Optionally, you can configure port mirroring or multiple port mirroring to direct packets to additional interfaces.

These active flow monitoring options provide a wide variety of actions that can be performed on network traffic flows. However, the following restrictions apply:


- The router or switch can perform sampling *or* port mirroring at any one time.
- The router or switch can perform forwarding *or* discard accounting at any one time.

Because the Monitoring Services, AS, and Multiservices PICs allow only one action to be performed at any one time, the following configuration options are available:

- Sampling and forwarding
- Sampling and discard accounting
- Port mirroring and forwarding
- Port mirroring and discard accounting
- Sampling and port mirroring on different sets of traffic

Figure 15 on page 57 shows a sample topology.

Figure 15: Active Monitoring Configuration Topology



In Figure 15 on page 57, traffic from Router 1 arrives on the monitoring router's Gigabit Ethernet ge-2/3/0 interface. The exit interface on the monitoring router leading to destination Router 2 is ge-3/0/0, but this can be any interface type (such as SONET, Gigabit Ethernet, and so on). The export interface leading to the cflowd server is fe-1/0/0.

To enable active monitoring, configure a *firewall filter* on the interface ge-2/3/0 with the following match conditions:

- Traffic matching certain firewall conditions is sent to the Monitoring Services PIC using filter-based forwarding. This traffic is quarantined and not forwarded to other routers.

- All other traffic is port-mirrored to the Monitoring Services PIC. Port mirroring copies each packet and sends the copies to the port-mirroring next hop (in this case, a Monitoring Services PIC). The original packets are forwarded out of the router as usual.

RELATED DOCUMENTATION

[Configuring Flow Monitoring | 5](#)

[Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers | 637](#)

[Configuring Services Interface Redundancy with Flow Monitoring | 72](#)

[Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System | 58](#)

Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System

This example shows a sample configuration that allows you to configure active monitoring on a logical M-series, MX-series, T-series, or PTX Series system.

The following section shows the configuration on the primary router:

```
[edit forwarding-options]
sampling {
  instance inst1 {
    input {
      rate 1;
    }
    family inet;
    output {
      flow-server 198.51.100.2 {
        port 2055;
        version9 {
          template {
            ipv4;
          }
        }
      }
    }
  }
}
```

```

        interface sp-0/1/0 {
            source-address 10.11.12.13;
        }
    }
}
family mpls;
output {
    flow-server 198.51.100.2 {
        port 2055;
        version9 {
            template {
                mpls;
            }
        }
    }
}
interface sp-0/1/0 {
    source-address 10.11.12.13;
}
}
}
services {
    flow-monitoring {
        version9 {
            template ipv4 {
                flow-active-timeout 60;
                flow-inactive-timeout 60;
                ipv4-template;
                template-refresh-rate {
                    packets 1000;
                    seconds 10;
                }
                option-refresh-rate {
                    packets 1000;
                    seconds 10;
                }
            }
            template mpls {
                mpls-template;
            }
        }
    }
}

```

```

    }
}

```

The configuration for the logical router uses the input parameters and the output interface for sampling from the primary router. Each logical router should have separate template definitions for the flow-server configuration. The following section shows the configuration on the logical router:

```

logical-systems {
  ls-1 {
    firewall {
      family inet {
        filter test-sample {
          term term-1 {
            then {
              sample;
              accept;
            }
          }
        }
      }
    }
    interfaces {
      ge-0/0/1 {
        unit 0 {
          family inet {
            filter {
              input test-sample;
              output test-sample;
            }
          }
        }
      }
    }
    forwarding-options {
      sampling {
        instance sample-inst1 {
          family inet;
          output {
            flow-server 198.51.100.2 {
              port 2055;
              version9 {
                template {

```

```

        ipv4-ls1;
    }
}
}
}
}
}
family mpls;
    output {
        flow-server 198.51.100.2 {
            port 2055;
            version9 {
                template {
                    mpls-ls1;
                }
            }
        }
    }
}
}
}
services {
    flow-monitoring {
        version9 {
            template ipv4-ls1 {
                flow-active-timeout 60;
                flow-inactive-timeout 60;
                ipv4-template;
                template-refresh-rate {
                    packets 1000;
                    seconds 10;
                }
                option-refresh-rate {
                    packets 1000;
                    seconds 10;
                }
            }
            template mpls-ls1 {
                mpls-template;
            }
        }
    }
}
}

```



```
}
}
```

RELATED DOCUMENTATION

[Active Flow Monitoring Overview | 54](#)

[Configuring Flow Monitoring | 5](#)

[Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers | 637](#)

[Configuring Services Interface Redundancy with Flow Monitoring | 72](#)

Example: Configuring Flow Monitoring on an MX Series Router with MS-MIC and MS-MPC

IN THIS SECTION

- [Hardware and Software Requirements | 68](#)
- [Junos Traffic Vision Support on MS-MIC and MS-MPC | 69](#)
- [Verification | 70](#)

This example shows how you can configure Junos Traffic Vision for flow monitoring on an MX Series Router with MS-MIC and MS-MPC, and contains the following sections:

Configuring Flow Monitoring on MS-MIC

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.

NOTE: You can follow the same procedure and use the same configuration for configuring flow monitoring on MS-MPC.

Enabling the Services Interface Card

```
set interfaces ms-2/0/0 unit 0 family inet
```

Configuring the Template and Timers

```
set services flow-monitoring version9 template template1
set services flow-monitoring version9 template template1 flow-active-timeout 120
set services flow-monitoring version9 template template1 flow-inactive-timeout 60
set services flow-monitoring version9 template template1 ipv4-template
set services flow-monitoring version9 template template1 template-refresh-rate packets 100
set services flow-monitoring version9 template template1 template-refresh-rate seconds 600
set services flow-monitoring version9 template template1 option-refresh-rate packets 100
set services flow-monitoring version9 template template1 option-refresh-rate seconds 600
```

Configuring Service Set Properties

```
set services service-set ss1 jflow-rules sampling
set services service-set ss1 sampling-service service-interface ms-2/0/0.0
```

Configuring Forwarding Options and Flow Server Settings

```
set forwarding-options sampling input rate 10
set forwarding-options sampling input run-length 18
set forwarding-options sampling family inet output flow-server 10.44.4.3 port 1055
set forwarding-options sampling family inet output flow-server 10.44.4.3 version9 template
template1
set forwarding-options sampling family inet output interface ms-2/0/0.0 source-address
203.0.113.1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

NOTE: The MS interface must be configured with the family type that the collector will be reachable by. If the collector for the sampling traffic is reachable via IPv4, you must set the family inet under the MS interface even if you are only sampling IPv6 and MPLS traffic, for example.

1. Configure the services interface.

```
[edit interfaces]
user@router1# set interfaces ms-2/0/0 unit 0 family inet
user@router1# set interfaces ms-2/0/0 unit 1 family inet6
user@router1# set interfaces ms-2/0/0 unit 2 family mpls
```

2. Configure the template properties and the export policy timers.

```
[edit services]
user@router1# set flow-monitoring version9 template template1
user@router1# set flow-monitoring version9 template template1 flow-active-timeout 120
user@router1# set flow-monitoring version9 template template1 flow-inactive-timeout 60
user@router1# set flow-monitoring version9 template template1 ipv4-template
user@router1# set flow-monitoring version9 template template1 template-refresh-rate packets
100
user@router1# set flow-monitoring version9 template template1 template-refresh-rate seconds
600
user@router1# set flow-monitoring version9 template template1 option-refresh-rate packets 100
user@router1# set flow-monitoring version9 template template1 option-refresh-rate seconds 600
```

Table 23: Quick Reference to Key Configuration Statements at This Hierarchy Level

Configuration Statement	Description
flow-active-timeout	Configures the interval (in seconds) after which an active flow is exported. Range is 10 through 600 seconds, and the default value is 60 seconds.

Table 23: Quick Reference to Key Configuration Statements at This Hierarchy Level *(Continued)*

Configuration Statement	Description
flow-inactive-timeout	<p>Configures the interval (in seconds) of inactivity after which a flow is marked inactive.</p> <p>Range is 10 through 600 seconds, and the default value is 60 seconds.</p>
<i>ipv4-template / ipv6-template / mpls-template / mpls-ipv4-template</i>	Specifies the type of traffic for which the template is used for.
template-refresh-rate	<p>Specifies the template refresh rate either as number of packets (range is 1 through 480,000 and the default value is 4800) or in seconds (the range is 10 through 600 and the default is 60).</p> <p>Because the communication between the flow generator and the flow collector is a one-way communication, the flow generator has to regularly send updates about template definitions to the flow collector. The value configured for this statement controls the frequency of such updates.</p>
option-refresh-rate	Specifies the option refresh rate either as number of packets (range is 1 through 480,000 and the default value is 4800) or in seconds (the range is 10 through 600 and the default is 60).

3. Configure service set properties.

```
[edit services]
user@router1# set service-set ss1 jflow-rules sampling
user@router1# set service-set ss1 sampling-service service-interface ms-2/0/0.0
```

Table 24: Quick Reference to Configuration Statements at This Hierarchy Level

Configuration Statement	Description
sampling	Configures the service set to handle sampling/flow monitoring activities.

Table 24: Quick Reference to Configuration Statements at This Hierarchy Level (*Continued*)

Configuration Statement	Description
service-interface	<p>Specifies the service interface associated with the service set.</p> <p>The interface configured here should match the interface configured at the [edit forwarding-options sampling family inet output]. Also, note that the interface should not be associated with any other service set.</p>

4. Configure forwarding options and flow-server properties.

```
[edit forwarding-options]
user@router1# set sampling input rate 10
user@router1# set sampling input run-length 18
user@router1# set sampling family inet output flow-server 10.44.4.3 port 1055
user@router1# set sampling family inet output flow-server 10.44.4.3 version9 template
template1
user@router1# set sampling family inet output interface ms-2/0/0.0 source-address 203.0.113.1
```

NOTE: You can specify the sampling parameters either at the global level (as shown in this example) or at the FPC level by defining a sampling instance. To define a sampling instance, include the instance statement at the [edit forwarding-options sampling] hierarchy level, and the sampling-instance statement at the [edit chassis fpc *number*] hierarchy level to associate the sampling instance with an FPC. Under the [edit forwarding-options sampling instance *instance*] hierarchy level, you must also include the input and output configurations explained in this step.

Table 25: Quick Reference to Key Configuration Statements at this Hierarchy Level

Configuration Statement	Description
rate	<p>The ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.</p> <p>The range is 1 through 16000000(16M).</p>

Table 25: Quick Reference to Key Configuration Statements at this Hierarchy Level *(Continued)*

Configuration Statement	Description
run-length	The number of samples following the initial trigger event. This enables you to sample packets following those already being sampled. The range is 0 through 20, and the default is 0.
flow-server	A host system to collect sampled flows using the version 9 format.
source-address	An IPv4 address to be used as the source address of the exported packet.

Result

From the configuration mode, confirm your configuration by entering the `show chassis fpc 2`, `show interfaces`, and `show forwarding-options` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@router1# show interfaces
ms-2/0/0 {
  unit 0 {
    family inet;
  }
}
```

```
user@router1# show services
flow-monitoring {
  version9 {
    template template1 {
      flow-active-timeout 120;
      flow-inactive-timeout 60;
      template-refresh-rate {
        packets 100;
        seconds 600;
      }
      option-refresh-rate {
        packets 100;
```

```

        seconds 600;
    }
    ipv4-template;
}
}
}
service-set ss1 {
    jflow-rules {
        sampling;
    }
    sampling-service {
        service-interface ms-2/0/0.0
    }
}
}

```

```

user@router1# show forwarding-options
sampling {
    input {
        rate 10;
        run-length 18;
    }
    family inet {
        output {
            flow-server 10.44.4.3 {
                port 1055;
                version9 {
                    template {
                        template1;
                    }
                }
            }
            interface ms-2/0/0.0 {
                source-address 203.0.113.1;
            }
        }
    }
}
}

```

Hardware and Software Requirements

This example requires an MX Series router that has:

- Junos OS Release 13.2 running on it.
- An MS-MIC installed in it.

Junos Traffic Vision Support on MS-MIC and MS-MPC

Junos Traffic Vision (previously known as Jflow) is the accounting service that is available on the MS-MIC and MS-MPC. Junos Traffic Vision enables users to keep track of the packets received on the MS-MIC or MS-MPC and to generate flow records that contain information such as the source address of the packet, the destination address of the packet, packets and byte counts, and so on. Junos Traffic Vision implementation does not interrupt the traffic, instead it makes a copy of the incoming packet and sends that copy to the service interface card for analyzing the information and maintaining the record.

Starting with Release 13.2, the Junos OS extension-provider packages come preinstalled on a multiservices MIC and MPC (MS-MIC and MS-MPC). The adaptive-services configuration at the [edit chassis fpc *number* pic *number*] hierarchy level is preconfigured on these cards.

Before you configure Junos Traffic Vision on an MS-MIC or an MS-MPC, you must create a firewall filter that has `sample` configured as action, and apply that to the interface on which you want to monitor the traffic. The flow-collector in Junos Traffic Vision implementations is a device for collecting the flow records. The flow collector is typically deployed outside the network.

NOTE: For more information about configuring firewall filters, see the Junos OS *Firewall Filters Configuration Guide*.

On MS-MIC and MS-MPC, Junos OS supports Junos Traffic Vision Version 9 (v9). Junos Traffic Vision v9 supports sampling of IPv4, IPv6, and MPLS traffic. A services interface card is essential for the v9 implementation, and hence this is often known as PIC-based monitoring.

You can configure the maximum time for which the flow records are stored on the services interface card. The active timeout and inactive timeout values, configured while defining the template, control the export of flow records to the collector. An MS-MIC can store a maximum of 14 million flow records, whereas an MS-MPC can store upto 30 million flows per NPU.

NOTE: In Junos Traffic Vision configurations using the Junos OS extension-provider package, modifying the following statements after flow monitoring has been initiated causes all existing flows to expire:

- At the [edit forwarding-options sampling instance *instance-name* family (inet |inet6 |mpls) output] and [edit forwarding-options sampling family (inet |inet6 |mpls) output] hierarchy levels:

- flow-server *ip-address*
- flow-server port *port-number*
- flow-server template *template*
- At the [edit services flow-monitoring version9 template *template-name* mpls-ipv4-template] and [edit services flow-monitoring version9 template *template-name* mpls-template] hierarchy levels:
 - label-position

Because these changes can disrupt the ongoing flow monitoring, we recommend that you do not change these values after flow monitoring has been initiated on a device. The changes made to these configuration statements when flow monitoring is going on, apply only to the newly created flows.

Also, note that these changes do not disrupt flow monitoring on devices running Jflow configuration using the Junos OS Layer 2 services package. However, even in the case of Layer 2 service package-based configuration, the changes are applied only to the newly created flows. The existing flows continue to use the initial settings.

NOTE: When Junos Traffic Vision is configured on the MS-MIC and MS-MPC, the next-hop address and outgoing interfaces are incorrectly displayed in the IPv4 and IPv6 flow records when the destination of the sampled flow is reachable through multiple paths.

Verification

IN THIS SECTION

- [Verifying the Junos Traffic Vision Configuration | 71](#)
- [Viewing the Flow Details | 71](#)
- [Viewing Details of Errors That Occurred on the Services Interface | 72](#)

Confirm that the configuration is working properly.

Verifying the Junos Traffic Vision Configuration

Purpose

Verify that Junos Traffic Vision is enabled on the router.

Action

From operational mode, enter the `show services accounting status` command.

```
user@router1> show services accounting status
Service Accounting interface: ms-2/0/0
Export format: 9, Route record count: 2093
IFL to SNMP index count: 35, AS count: 2
Configuration set: Yes, Route record set: Yes, IFL SNMP map set: Yes
```

Meaning

Shows the service interface on which monitoring is configured, and also provides information about the export format used (version 9 in this case).

Viewing the Flow Details

Purpose

View the flow details on the interface configured for flow monitoring.

Action

From operational mode, enter the `show services accounting flow` command.

```
user@router1> show services accounting flow
Flow information
Service Accounting interface: ms-2/0/0, Local interface index: 229
Flow packets: 220693, Flow bytes: 24276230
Flow packets 10-second rate: 99, Flow bytes 10-second rate: 10998
Active flows: 10, Total flows: 12
Flows exported: 199, Flows packets exported: 718
Flows inactive timed out: 2, Flows active timed out: 199
```

Viewing Details of Errors That Occurred on the Services Interface

Purpose

View details of errors, if any, on the interface that is configured for flow monitoring.

Action

From operational mode, enter the `show services accounting errors` command.

```
user@router1> show services accounting errors
Error information
  Service Accounting interface: ms-2/0/0
  Service sets dropped: 0, Active timeout failures: 0
  Export packet failures: 0, Flow creation failures: 0
  Memory overload: No
```

RELATED DOCUMENTATION

Multiservices MIC and Multiservices MPC (MS-MIC and MS-MPC) Overview

Example: Configuring Junos VPN Site Secure on MS-MIC and MS-MPC

Configuring Services Interface Redundancy with Flow Monitoring

Active monitoring services configurations on AS, Multiservices PICs, and Multiservices DPCs support redundancy. To configure redundancy, you specify a redundancy services PIC (`rsp`) interface in which the primary AS or Multiservices PIC is active and a secondary PIC is on standby. If the primary PIC fails, the secondary PIC becomes active, and all service processing is transferred to it. If the primary PIC is restored, it remains on standby and does not preempt the secondary PIC; you need to manually restore the services to the primary PIC. To determine which PIC is currently active, issue the `show interfaces redundancy` command.

NOTE: On flow-monitoring configurations, the only service option supported is *warm standby*, in which one backup PIC supports multiple working PICs. Recovery times are not guaranteed, because the configuration must be completely restored on the backup PIC after a failure is detected. However, configuration is preserved and available on the new active PIC.

As with the other services that support warm standby, you can issue the `request interfaces (revert | switchover)` command to switch manually between the primary and secondary flow monitoring interfaces.

For more information, see *Configuring AS or Multiservices PIC Redundancy*. For information on operational mode commands, see the [CLI Explorer](#).

A sample configuration follows.

```
interface {
  rsp0 {
    redundancy-options {
      primary sp-0/0/0;
      secondary sp-1/3/0;
    }
    unit 0 {
      family inet;
    }
  }
}
interface {
  ge-0/2/0 {
    unit 0 {
      family inet {
        filter {
          input as_sample;
        }
      }
      address 10.58.255.49/28;
    }
  }
}
forwarding-options {
  sampling {
    instance instance1 { # named instances of sampling parameters
      input {
        rate 1;
        run-length 0;
        max-packets-per-second 65535;
      }
      family inet {
```


Support for active flow monitoring with IPFIX templates on QFX10002 switches was added in Junos OS Release 17.2R1. Starting in Junos OS Release 20.3R1 on QFX10002-60C switches, you can configure inline active flow monitoring for IPv4 and IPv6 traffic. Both IPFIX and version 9 templates are supported.

On routers with MS-PICs or MS-DPCs, IPv4 and IPv6 fragments are processed accurately. The flow monitoring application creates two flows for every fragmented flow. The first fragment that has the complete Layer 4 information forms the first flow with 5-tuple data and subsequently, all the fragmented packets related to this flow form another flow with the Layer 4 fields set to zero.

The following considerations apply to the inline flow-monitoring instance configuration:

- Sampling run-length and clip-size are not supported.
- For inline configurations, collectors are not reachable via `fxp0`.
- Inline flow monitoring does not support `cflowd`. Therefore, inline flow monitoring does not support the local dump option, which is available only with `cflowd`.
- Inline active flow monitoring is not supported when you enable Next Gen Services on an MX Series router.
- The number of collectors that are supported depends on the device:
 - In Junos OS Release 16.2 and in Junos OS Release 16.1R3 and earlier, you can configure only one collector under a family for inline active flow monitoring. Starting with Junos OS Release 16.1R4 and 17.2R1, you can configure up to four collectors under a family for inline active flow monitoring. Starting with Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring. The Packet Forwarding Engine (PFE) can export the flow record, flow record template, option data, and option data template packet to all configured collectors. To configure a collector under a family for inline active flow monitoring, configure the `flow-server` statement at the edit forwarding-options sampling-instance *instance-name* family (inet | inet6) output hierarchy level. To specify up to four collectors, include up to four `flow-server` statements.
 - For inline configurations on all other devices, each family can support only one collector.

Inline active flow monitoring is available in four hierarchy levels:

- `[edit chassis]` —At this level, you associate the sampling instance with the FPC on which the media interface is present (except on the MX80 and MX104—see ["Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers" on page 537](#)). If you are configuring sampling of IPv4

flows, IPv6 flows or VPLS flows (Junos OS only), you can configure the flow hash table size for each family, as described below.

- [edit firewall]—At this level, you configure a firewall filter for the family of traffic to be sampled. You must attach this filter to the interface on which you want to sample the traffic.
- [edit forwarding-options]—At this level, you configure a sampling instance and associate the template with the sampling instance. At this level, you also configure the flow-server IP address and port number as well as the flow export rate.
- [edit services flow-monitoring] —At this level, you configure the template properties for inline flow monitoring.

Before you configure inline active flow monitoring, you should ensure that you have adequately-sized hash tables for IPv4, IPv6, MPLS, and VPLS flow sampling. (VPLS flow sampling is Junos OS only). These tables can use one to fifteen 256K areas. Starting with Junos OS Release 16.1R1 and 15.1F2, the IPv4 table is assigned a default value of 1024. Prior to Junos OS Release 16.1 and 15.1F2, the IPv4 table is assigned a default value of fifteen 256K areas. The IPv6 table is assigned a default value of 1024, and the VPLS table is assigned a default value of 1024. When anticipated traffic volume requires larger tables, allocate larger tables.

To allocate flow hash tables:

1. Go to the [edit chassis fpc 0 inline-services flow-table-size] hierarchy level for inline services on the FPC that processes the monitored flows.

```
[edit]
user@host# edit chassis fpc 0 inline-services flow-table-size
```

2. Specify the required sizes for the sampling hash tables.

```
[edit chassis fpc 0 inline-services flow-table-size]
user@host# set bridge-flow-table-size units
user@host# set ipv4-flow-table-size units
user@host# set ipv6-flow-table-size units
user@host# set mpls-flow-table-size units
user@host# set vpls-flow-table-size units
```

NOTE: Starting in Junos OS Release 18.2R1, the `bridge-flow-table-size` option is available and the `vpls-flow-table-size` option is deprecated; use the `bridge-flow-table-size` option instead. The `bridge-flow-table-size` option supports both VPLS and bridge records.

NOTE: The total number of units used for IPv4, IPv6, MPLS, and VPLS cannot exceed 15. Also, starting in Junos OS Release 16.1R1 and 15.1F2, changing the flow hash table size *not* automatically reboot the FPC (for earlier releases changing the flow hash table size triggers the FPC to reboot).

To configure inline active flow monitoring on MX Series routers (except for MX80 and MX104 routers), EX Series switches, and T4000 routers with Type 5 FPC:

1. Enable inline active flow monitoring and specify the source address for the traffic.

```
[edit forwarding-options sampling instance instance-name family (bridge | inet | inet6 |
mpls | vpls ) output]
user@host# set inline-jflow source address address
```

2. Specify the template to use with the sampling instance.

```
[edit forwarding-options sampling instance instance-name family (bridge | inet | inet6 |
mpls | vpls ) output flow-server hostname]
user@host# set (version9 | version-ipfix) template template-name
```

3. Configure a template to specify output properties.

```
[edit services flow-monitoring]
user@host# set (version-ipfix | version9) template template-name
```

4. (Optional) Configure the interval after which an active flow is exported.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-active-timeout seconds
```


5. (Optional) Configure the interval of activity that marks a flow as inactive.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-inactive-timeout seconds
```

6. (Optional) Configure the template refresh rate in either number of packets or number of seconds.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set template-refresh-rate (packets packets | seconds seconds)
```

7. (Optional) Configure the refresh rate in either number of packets or number of seconds.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set option-refresh-rate (packets packets | seconds seconds)
```

8. Specify the type of record that the template is used for.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set (bridge-template | ipv4-template | ipv6-template | mpls-ipv4-template | mpls-
template | peer-as-billing-template | vpls-template)
```

The vpls-template option is only for IPFIX templates.

Starting in Junos OS Release 18.2R1, the bridge-template option is available and the vpls-template option is deprecated; use the bridge-template option instead. The bridge-template option (Junos OS only) supports both VPLS and bridge records and is for both IPFIX and version9 templates.

Starting in Junos OS Release 18.4R1, the mpls-ipv4-template option is deprecated for inline flow monitoring. To configure MPLS records starting in Junos OS Release 18.4R1, use the mpls-template option and the tunnel-observation option. This is described in step "9" on page 78.

9. Starting in Junos OS Release 18.4R1 for the MX Series, if you are configuring any type of MPLS flow records, perform the following:

- a. Specify the MPLS template.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set mpls-template
```

- b. Configure the type of MPLS flow records to create.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set tunnel-observation [ipv4 | ipv6]
```

The tunnel-observation values enable the creation of the following types of flow records:

- ipv4—MPLS-IPv4 flows
- ipv6—MPLS-IPv6 flows

You can configure multiple values for tunnel-observation.

For an MPLS traffic type that does *not* match any of the tunnel-observation values, plain MPLS flow records are created. For example, if you only configure ipv4, then MPLS-IPv6 traffic results in plain MPLS flow records.

If you do not configure tunnel-observation, plain MPLS flow records are created.

- c. If you are running inline flow monitoring on a Lookup (LU) card, enable sideband mode to create MPLS-IPv6 flow records.

```
[edit chassis fpc slot-number inline-services]
user@host# set use-extended-flow-memory
```

If you are running inline flow monitoring on an LU card and do not enable sideband mode, then MPLS-IPv6 traffic results in plain MPLS flow records.

10. (Optional) Include the flow direction value in the template.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-key flow-direction
```

The reported data field contains 0x00 (ingress) or 0x01 (egress). If you do not include the flow-key flow-direction statement, the flow direction data field contains the invalid value 0xFF.

11. (Optional) Include VLAN IDs in both the ingress and egress directions in the flow key.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-key vlan-id
```

This statement is not required for ingress and egress VLAN ID reporting on interfaces.

12. Associate the sampling instance with the FPC on which you want to implement inline active flow monitoring.

For MX240, MX480, MX960, MX2010, MX2020, use the following command:

```
[edit ]
user@host# set chassis fpc fpc-number sampling-instance instance-name
```

- a. Confirm the configuration by running the following show command:

```
user@host# show chassis
```

```
fpc 0 {
    sampling-instance sample-ins1;
}
```

For MX5, MX10, MX40, and MX80, use the following command:

```
[edit ]
user@host# set chassis tfeb slot 0 sampling-instance instance-name
```

- a. Confirm the configuration by running the following show command:

```
user@host# show chassis
```

```
tfeb {
    slot 0 {
        sampling-instance sample-ins1;
    }
}
```

For MX104, use the following command:

```
[edit ]
user@host# set chassis afeb slot 0 sampling-instance instance-name
```

- a. Confirm the configuration by running the following show command:

```
user@host# show chassis
```

```
afeb {
  slot 0 {
    sampling-instance sample-ins1;
  }
}
```

This example shows the sampling configuration for an instance that supports inline active flow monitoring on family inet:

```
[edit]
user@host> show forwarding-options
sampling {
  instance {
    sample-ins1 {
      input {
        rate 1;
      }
      family inet {
        output {
          flow-server 192.0.2.2 {
            port 2055;
            version-ipfix {
              template {
                ipv4;
              }
            }
          }
        }
        inline-jflow {
          source-address 10.11.12.13;
```

```

    }
  }
}
}
}

```

Here is the output format configuration:

```

[edit]
user@host> show services flow-monitoring
services {
  flow-monitoring {
    version-ipfix {
      template ipv4 {
        flow-active-timeout 60;
        flow-inactive-timeout 60;
        ipv4-template;
        template-refresh-rate {
          packets 1000;
          seconds 10;
        }
        option-refresh-rate {
          packets 1000;
          seconds 10;
        }
      }
    }
  }
}

```

The following example shows the output format configuration for chassis fpc0:

```

[edit]
user@host> show services flow-monitoring
sampling-instance instance-1; {
  inline-services {
    flow-table-size {
      ipv4-flow-table-size 8;
      ipv6-flow-table-size 7;
    }
  }
}

```

```

    }
}

```

Release History Table

Release	Description
21.1R1-Evo	Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring.
20.4R1-Evo	Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
20.3R1-Evo	Starting with Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
19.3R2	Inline active flow monitoring is not supported when you enable Next Gen Services on an MX Series router.
18.4R1	Starting in Junos OS Release 18.4R1, the <code>mpls-ipv4-template</code> option is deprecated for inline flow monitoring. To configure MPLS records starting in Junos OS Release 18.4R1, use the <code>mpls-template</code> option and the <code>tunnel-observation</code> option.
18.2R1	Starting in Junos OS Release 18.2R1, the <code>bridge-flow-table-size</code> option is available and the <code>vpls-flow-table-size</code> option is deprecated; use the <code>bridge-flow-table-size</code> option instead.
18.2R1	Starting in Junos OS Release 18.2R1, the <code>bridge-template</code> option is available and the <code>vpls-template</code> option is deprecated; use the <code>bridge-template</code> option instead.
16.1R4	In Junos OS Release 16.2 and in Junos OS Release 16.1R3 and earlier, you can configure only one collector under a family for inline active flow monitoring. Starting with Junos OS Release 16.1R4 and 17.2R1, you can configure up to four collectors under a family for inline active flow monitoring.
16.1R1	Also, starting in Junos OS Release 16.1R1 and 15.1F2, changing the flow hash table size does <i>not</i> automatically reboot the FPC.
16.1R1	Starting with Junos OS Release 16.1R1 and 15.1F2, the IPv4 table is assigned a default value of 1024.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers | 537](#)

[Example: Configuring Inline Active Flow Monitoring on MX Series and T4000 Routers | 567](#)

[inline-jflow | 1162](#)

Configuring Flow Offloading on MX Series Routers

The Junos OS enables you to configure flow offloading for PICS on MX Series routers using Modular Port Concentrator (MPCs) with Modular Interface Cards (MICs). Flows are offloaded to Fast Update Filters (FUFs) on the Packet Forwarding Engine. Offloading produces the greatest benefits when applied to long-lasting or high-bandwidth flows.

The maximum number of active offloads is 200,000 per PIC. When offloaded flows are deleted, more flows can be offloaded.

To configure flow offloading:

- At the [edit interfaces *interface-name* services-options] hierarchy level, enter the trio-flow-offload minimum-bytes *minimum-bytes* statement.

```
user@host# edit services interface-name
[edit services interface-name services-options]
user@host# set trio-flow-offload minimum-bytes minimum-bytes
```

In the following example, flows are offloaded when they consist of no less than 1024 bytes:

```
user@host# edit services ms-0/1/0
[edit services ms-0/1/0 services-options]
user@host# set trio-flow-offload minimum-bytes 1024
```

RELATED DOCUMENTATION

[trio-flow-offload](#)

Configuring Active Flow Monitoring on PTX Series Packet Transport Routers

You can use flow monitoring to help with network administration. Active flow monitoring on PTX Series routers allows you to collect sampled packets, then the router does GRE encapsulation of the packets and sends them to a remote server for flow processing. The GRE encapsulation includes an interface index and GRE key field. The GRE encapsulation removes MPLS tags. You configure one or more port-mirroring instances to define which traffic to sample and configure a server to receive the GRE encapsulated packets. You configure a firewall filter on interfaces where you want to capture flows. You can configure as many as 48 port-mirroring instances.

To configure the router to do GRE encapsulation of sampled packets and send them to a remote server for flow processing:

1. Configure one or more server profiles that specify a host where GRE encapsulated sampled packets are sent, and optionally, a source address to include in the header of each sampled packet.
 - a. Specify a name for each server profile and an IP address of the host where sampled packets are sent:

```
[edit services hosted-services]
user@host# set server-profile server-profile-name server-address ipv4-address
```

- b. (Optional) For each server profile, specify a source address to include in the header of each sampled packet:

```
[edit services hosted-services server-profile server-profile-name]
user@host# set client-address ipv4-address
```

NOTE: The default client address is 0.0.0.0. You must specify an IPv4 address as the client address. You can also specify the loopback address or management interface address as the client address.

2. Configure one or more port-mirroring instances.
 - a. Specify a name for each port-mirroring instance:

```
[edit forwarding-options port-mirroring]
user@host# set instance instance-name
```


NOTE: You can configure a maximum of 48 port-mirroring instances.

- b. Specify a protocol family for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name]
user@host# set family (inet | inet6 )
```

3. To set the ratio of the number of packets to sample, specify a value from 1 through 65,535 for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name input]
user@host# set rate number
```

NOTE: You must specify a value for the rate statement. The default value is zero, which effectively disables sampling. If, for example, you specify a rate value of 4, every fourth packet (1 packet out of 4) is sampled.

4. (Optional) Specify the number of samples to collect after the initial trigger event for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name input]
user@host# set run-length number
```

NOTE: The default value is zero. You can specify a number up to 20.

5. To designate a host where sampled traffic is sent, specify the name of server profile configured at the [edit services hosted-services] hierarchy level for each port-mirroring instance:

```
[edit forwarding-options port-mirroring instance instance-name family ( inet | inet6) output]
user@host# set server-profile server-profile-name
```

6. Configure one or more firewall filters.

- a. For each firewall filter, specify a protocol family, filter name, and match conditions:

```
[edit firewall]
user@host# set filter family (inet | inet6) filter filter-name term term-name from match-conditions
```

- b. For each firewall filter you configure, specify the name of a port-mirroring instance you configured at the [edit forwarding-options] hierarchy level as a nonterminating action so that the traffic that matches that instance is sampled:

```
[edit firewall family (inet | inet6) filter filter-name term term-name]
user@host# set then port-mirroring instance instance-name
```

7. Apply each firewall filter to an interface to evaluate incoming traffic:

```
[edit interfaces interface-name unit logical-unit-number]
user@host# set family (inet | inet6) filter input firewall-filter-name
```

NOTE: Active flow monitoring is supported only on incoming traffic. You cannot apply firewall filters to evaluate outgoing traffic.

8. Configure the remote server, where GRE encapsulated packets are sent, to perform flow processing.

RELATED DOCUMENTATION

Configuring Port Mirroring

hosted-services

port-mirroring

server-profile (Active Flow Monitoring)

Firewall Filter Nonterminating Actions

Configuring Actively Monitored Interfaces on M, MX and T Series Routers

Configure the input interfaces and apply the firewall filter that you defined earlier. Unlike passive flow monitoring, the input interfaces for active flow monitoring are not restricted, so you can select most standard network interfaces (such as ATM1 or Ethernet-based interfaces) as the input.

If you configure active flow monitoring with sampling, you can configure an interface filter in place of a firewall filter with the `sampling` statement at the `[edit interfaces interface-name-fpc/pic/port unit unit-number family inet]` hierarchy level.

```
[edit]
interfaces {
  so-2/2/0 {
    unit 0 {
      family inet {
        filter {
          input active_filter;
        }
        address 10.36.11.2/32 {
          destination 10.36.11.1;
        }
        sampling {
          (input | output | [input output]);
        }
      }
    }
  }
}
```

Collecting Flow Records

Traffic flows can be exported in flow monitoring version 5, 8, and 9 formats for active flow monitoring. The default export format for flow monitoring records is version 5. To change the export format to flow monitoring version 8, include the `version 8` statement at either the `[edit forwarding-options accounting name output flow-server flow-server-address]` or the `[edit forwarding-options sampling output flow-server flow-server-address]` hierarchy level. To change the export format to flow monitoring version 9, include the `version9 template template-name` statement at the `[edit forwarding-options sampling output`

flow-server *flow-server-address*] hierarchy level. For more information on flow record formats, see ["Flow Monitoring Output Formats" on page 11](#).

To capture flow data generated by the Monitoring Services PIC, Adaptive Services PIC, or MultiServices PIC and export it to a flow server, you can use one of the following active flow monitoring methods:

- ["Configuring M, MX and T Series Routers for Discard Accounting with a Sampling Group" on page 90](#)
- ["Configuring M, MX and T Series Routers for Discard Accounting with an Accounting Group" on page 89](#)
- ["Configuring M, MX and T Series Routers for Discard Accounting with a Template" on page 92](#)
- ["Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers" on page 96](#)
- ["Replicating Version 9 Flow Aggregation From M, MX and T Series Routers to Multiple Flow Servers" on page 97](#)
- ["Configuring Routing Engine-Based Sampling on M, MX and T Series Routers for Export to Multiple Flow Servers" on page 98](#)
- ["Configuring an Aggregate Export Timer on M, MX and T Series Routers for Version 8 Records" on page 126](#)

Configuring M, MX and T Series Routers for Discard Accounting with an Accounting Group

To perform discard accounting on specified traffic, you can collect flow records with the accounting statement at the [edit forwarding-options] hierarchy level. Like sampling, your topology must be simple (for example, one input interface and one export interface).

Again, you can collect flow records by specifying input and output interfaces. You can configure the input interface to perform discard accounting by applying a firewall filter that contains the then discard accounting statement. This match condition directs the filtered traffic to be converted into flow records and exported for analysis by the monitoring services or adaptive services interface. The original packets are then sent to the discard process. For the output, remember to specify the IP address and port of your flow server and the services interface you plan to use for processing flow records.

You must configure a source address, but the **engine-id** and **engine-type** output interface statements are added automatically. You can override these values manually to track different flows with a single flow

collector. SNMP input and output interface index information is captured in flow records by default when you configure discard accounting.

```
[edit]
forwarding-options {
  accounting counter1 {
    output {
      flow-inactive-timeout 65;
      flow-active-timeout 65;
      flow-server 10.60.2.1 {
        port 2055;
        version 8;
        aggregation {
          protocol-port;
          source-destination-prefix;
        }
      }
    }
    interface sp-2/0/0 {
      engine-id 1;
      engine-type 11;
      source-address 10.60.2.2;
    }
  }
}
```

Configuring M, MX and T Series Routers for Discard Accounting with a Sampling Group

If your needs for active flow monitoring are simple, you can collect flow records with a sampling group. Sampling does not require you to configure a monitoring group (as required in passive flow monitoring) because you can configure flow server information in the **sampling** hierarchy. When you wish to sample traffic, include the `sampling` statement at the `[edit forwarding-options]` hierarchy level.

The typical sampling configuration has one input interface and one export interface. The input interface is activated by the `then sample` statement in a firewall filter term. This match condition directs traffic to the sampling process. Alternatively, you can use an interface-based filter in place of a firewall filter if you include the `sampling` statement at the `[edit interfaces interface-name-fpc/pic/port unit unit-number family inet]` hierarchy level.

There are two types of sampling available: PIC-based sampling and Routing Engine-based sampling. PIC-based sampling occurs when a monitoring services or adaptive services interface is the target for the output of the sampling process. To enable PIC-based sampling, include the `interface` statement at the `[edit forwarding-options sampling output]` hierarchy level and specify a monitoring services or adaptive services interface as the output interface. If an output interface is not specified in the sampling configuration, sampling is performed by the Routing Engine.

To specify a flow server in a sampling configuration, include the `flow-server` statement at the `[edit forwarding-options sampling output]` hierarchy level. You must specify the IP address, port number, and flow monitoring version of the destination flow server. Routing Engine-based sampling supports flow aggregation of up to eight flow servers (version 5 servers and version 8 only) at a time. The export packets are replicated to all flow servers configured to receive them. In contrast, PIC-based sampling allows you to specify just one version 5 flow server and one version 8 server simultaneously. Flow servers operating simultaneously must have different IP addresses.

As part of the output interface statements, you must configure a source address. In contrast, the interface-level statements of **engine-id** and **engine-type** are both added automatically. However, you can override these values with manually configured statements to track different flows with a single flow collector, as needed. When you configure sampling, SNMP input and output interface index information is captured in flow records by default.

```
[edit]
forwarding-options {
  sampling {
    input {
      rate 1;
    }
    family inet {
      output {
        flow-inactive-timeout 15;
        flow-server 10.60.2.1 {
          port 2055;
          version 5;
        }
        interface sp-2/0/0 {
          engine-id 5;
          engine-type 55;
          source-address 10.60.2.2;
        }
      }
    }
  }
}
```

```
}
}
```

Configuring M, MX and T Series Routers for Discard Accounting with a Template

Flow monitoring version 9, which is based on RFC 3954, provides a way to organize flow data into templates. Version 9 also provides a way to actively monitor IPv4, IPv6, MPLS, and peer AS billing traffic. Version 9 is not supported on the AS-I PIC.

To activate templates in flow monitoring, you must configure a template and include that template in the version 9 flow monitoring configuration. Version 9 does not work in conjunction with versions 5 and 8.

To configure a version 9 template, include the **template** *template-name* statement at the [edit services flow-monitoring version9] hierarchy level. The Junos OS supports five different templates: **ipv4-template**, **ipv6-template**, **mpls-template**, **mpls-ipv4-template**, and **peer-as-billing-template**. To view the fields selected in each of these templates, see ["Flow Monitoring Version 9 Format Output Fields" on page 26](#).

```
[edit services]
flow-monitoring {
  version9 { # Specifies flow monitoring version 9.
    template mpls { # Specifies template you are configuring.
      template-refresh-rate {
        packets 6000; # The default is 4800 packets and the range is 1-480000 packets.
        seconds 90; # The default is 60 seconds and the range is 1-600 seconds.
      }
      option--refresh-rate {
        packets 3000; # The default is 4800 packets and the range is 1-480000 # packets.
        seconds 30; # The default is 60 seconds and the range is 1-600.
      }
      flow-active-timeout 60; # The default is 60 seconds and the range is # 10-600.
      flow-inactive-timeout 30; # The default is 60 seconds and the range 10-600.
      template-refresh-rate seconds 10; # The default is 60 seconds and the # range is 10-600
      mpls-template {
        label-positions [1 | 2 | 3]; # Specifies label position for the MPLS template.
      }
    }
  }
}
```

You can export to multiple templates at a time to a maximum of eight flow servers for AS PICs and one flow server for all other PICs. To assign a template to a flow output, include the template *template-name* statement at the [edit forwarding options sampling output flow-server version9] hierarchy level:

```
[edit]
forwarding-options {
  sampling {
    input {
      family mpls {
        rate 1;
        run-length 1;
      }
    }
    output {
      flow-server 10.60.2.1 { # The IP address and port of the flow server.
        port 2055;
        source-address 192.0.2.1;
        version9 { # Records are sent to the flow server using version 9 format.
          template { # Indicates a template will organize records.
            mpls; # Records are sent to the MPLS template.
          }
        }
      }
    }
  }
}
```

Defining a Firewall Filter on M, MX and T Series Routers to Select Traffic for Active Flow Monitoring

The first step in active flow monitoring is to configure the match conditions for acceptable traffic or quarantined traffic. Common match actions for active flow monitoring include **sample**, **discard**, **accounting**, **port-mirror**, and **accept**. To configure, include the desired action statements and a counter as part of the then statement in a firewall filter and apply the filter to an interface.

In sampling, the router reviews a portion of the traffic and sends reports about this sample to the flow monitoring server. Discard accounting traffic is counted and monitored, but not forwarded out of the router. Port-mirrored traffic is copied and sent to another interface. Accepted traffic is forwarded to the intended destination.

Most of these match combinations are valid. However, you can either port-mirror or sample with the same traffic at the same time, but not perform more than one action simultaneously on the same packets.

```
[edit]
firewall {
  family inet {
    filter active_filter {
      term quarantined_traffic {
        from {
          source-address {
            10.36.1.2/32;
          }
        }
        then {
          count quarantined-counter;
          sample;
          discard accounting;
        }
      }
      term copy_and_forward_the_rest {
        then {
          port-mirror;
          accept;
        }
      }
    }
  }
}
```

Processing IPv4 traffic on an M, MX or T Series Router Using Monitoring services, Adaptive services or Multiservices Interfaces

You configure the monitoring services, adaptive services, or multiservices interfaces with the `family inet` statement so they can process IPv4 traffic. However, you must remember that a monitoring services interface uses an **mo-** prefix and adaptive services and multiservices interfaces use an **sp-** prefix.

```
[edit]
interfaces {
  sp-2/0/0 {
    unit 0 {
      family inet {
        address 10.36.100.1/32 {
          destination 10.36.100.2;
        }
      }
    }
  }
}
```

Active flow monitoring records leave the router through an export interface to reach the flow monitoring server.

```
[edit]
interfaces {
  fe-1/0/0 {
    unit 0 {
      family inet {
        address 10.60.2.2/30;
      }
    }
  }
}
```

Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers

Routing Engine-based sampling supports up to eight flow servers for both flow monitoring version 5 and version 8 configurations. The total number of flow servers is limited to eight, regardless of how many are configured for version 5 or version 8.

When you configure version 5 or version 8 sampling, the export packets are replicated to all flow servers configured to receive them. If two flow servers are configured to receive version 5 records, both flow servers will receive records for a specified flow.

NOTE: With Routing-Engine-based sampling, if multiple flow servers are configured with version 8 export format, all of them must use the same aggregation type (for example, all flow servers receiving version 8 export could be configured for source-destination aggregation type).

The following configuration example allows replication of export packets to two flow servers.

```
[edit]
forwarding-options {
  sampling {
    input {
      rate 1;
    }
  }
  output {
    flow-server 10.10.3.2 {
      port 2055;
      version 5;
      source-address 192.168.164.119;
    }
    flow-server 172.17.20.62 {
      port 2055;
      version 5;
      source-address 192.168.164.119;
    }
  }
}
```

Replicating Version 9 Flow Aggregation From M, MX and T Series Routers to Multiple Flow Servers

With this feature, you can configure up to eight flow servers to receive packets for a version 9 flow monitoring template. Once a flow server is configured to receive this data, it will also receive the following periodic version 9 flow monitoring updates:

- Options data
- Template definition

With Routing Engine-based sampling, if multiple collectors are configured with version 8 export format, all of them must use the same aggregation-type.

The option and template definition refresh period is configured on a per-template basis at the **[edit services flow-monitoring]** hierarchy level.

The following configuration example allows replication of version 9 export packets to two flow servers.

```
forwarding-options {
  sampling {
    input {
      family inet {
        rate 1;
      }
    }
    output {
      flow-server 10.10.3.2 {
        port 2055;
        version9 {
          template {
            ipv4;
          }
        }
      }
      flow-server 172.17.20.62 {
        port 2055;
        version9 {
          template {
            ipv4;
          }
        }
      }
    }
  }
}
```

```

        flow-inactive-timeout 30;
        flow-active-timeout 60;
        interface sp-4/0/0 {
            source-address 10.10.3.4;
        }
    }
}
}

```

RELATED DOCUMENTATION

[Active Flow Monitoring Overview | 54](#)

[Active Flow Monitoring Overview | 53](#)

[Active Flow Monitoring Applications | 46](#)

[Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers | 96](#)

Configuring Routing Engine-Based Sampling on M, MX and T Series Routers for Export to Multiple Flow Servers

Routing Engine-based sampling supports up to eight flow servers for both version 5 and version 8 configurations. The total number of collectors is limited to eight, regardless of how many are configured for version 5 or version 8. When you configure sampling, the export packets are replicated to all collectors configured to receive them. If two collectors are configured to receive version 5 records, both collectors will receive records for a specified flow.

The following configuration example allows replication of export packets to two collectors.

```

forwarding-options {
    sampling {
        input {
            family inet {
                rate 1;
            }
        }
        output {
            cflowd 10.10.3.2 {
                port 2055;
            }
        }
    }
}

```

```

        version 5;
        source-address 192.168.164.119;
    }
    cflowd 172.17.20.62 {
        port 2055;
        version 5;
        source-address 192.168.164.119;
    }
}
}
}
}

```

Example: Copying Traffic to a PIC While an M, MX or T Series Router Forwards the Packet to the Original Destination

IN THIS SECTION

- [Requirements | 100](#)
- [Overview and Topology | 100](#)
- [Configuration | 101](#)
- [Verification | 122](#)

Traffic sampling enables you to copy traffic to a Physical Interface Card (PIC) while the router forwards the packet to its original destination. This example describes how to configure a router to perform sampling on the Routing Engine using the **samplerd** process. For this method, you configure a filter (input or output) with a matching term that contains the `then sample` statement. In addition, for VPN routing and forwarding (VRF) Routing Engine-based sampling, you configure a VRF routing instance that maps to an interface. Each VRF instance corresponds with a forwarding table. Routes on the interface go into the corresponding forwarding table.

For VRF Routing Engine-based sampling, the kernel queries the correct VRF route table based on the ingress interface index for the received packet. For interfaces configured in VRF, the sampled packets contain the correct input and output interface SNMP index, the source and destination AS numbers, and the source and destination mask.

NOTE: With Junos OS Release 10.1, VRF Routing Engine-based sampling is performed only on IPv4 traffic. You cannot use Routing Engine-based sampling on IPv6 traffic or on MPLS label-switched paths.

This example describes how to configure and verify VRF Routing Engine-based sampling on one router in a four-router topology.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 10.1 or later
- M Series, MX Series, or T Series router

Before you configure VRF Routing Engine-Based sampling on your router, be sure you have an active connection between the routers on which you configure sampling. In addition, you need to have an understanding of VRF to configure the interfaces and routing instances that form the basis of the sampling configuration; and an understanding of the BGP, MPLS, and OSPF protocols to configure the other routers in the network to bring up the sampling configuration.

Overview and Topology


IN THIS SECTION

- [Topology | 101](#)

The scenario in this example illustrates VRF Routing Engine-based sampling configured on the PE1 router in a four-router network. The CE routers use BGP as the routing protocol to communicate with the PE routers. MPLS LSPs pass traffic between the PE routers. Packets from the CE1 router are sampled on the PE1 router. Regular traffic is forwarded to the original destination (the CE2 router).

Topology

Figure 16: Routing Engine-Based Sampling Network Topology



Configuration

IN THIS SECTION

- [Configuring the CE1 Router | 101](#)
- [Configuring the PE1 Router | 104](#)
- [Configuring the PE2 Router | 112](#)
- [Configuring the CE2 Router | 119](#)

In this configuration example, the VRF Routing Engine-based sampling is configured on the PE1 router that samples the traffic that goes through the interface and routes configured in the VRF. The configurations on the other three routers are included to show the sampling configuration on the PE1 router working in the context of a network.

To configure VRF Routing Engine-based sampling for the network example, perform these tasks:

Configuring the CE1 Router

Step-by-Step Procedure

In this step, you configure interfaces, routing options, protocols, and policy options for the CE1 router. To configure the CE1 router:

1. Configure one interface with two IP addresses. One address is for traffic to the PE1 router; the other address is to check that traffic is flowing to the CE2 router:

```
[edit interfaces]
user@router-ce1# set ge-1/3/2 unit 0 family inet address 192.0.2.1/24
user@router-ce1# set ge-1/3/2 unit 0 family inet address 198.51.100.2/8
```

2. Configure the autonomous system to establish a connection between BGP peers:

```
[edit routing-options]
user@router-ce1# set autonomous-system 95000
```

3. Configure BGP as the routing protocol between the CE router and the PE router:

```
[edit protocols]
user@router-ce1# set bgp group to_r1 type external
user@router-ce1# set bgp group to_r1 export my_lo0_addr
user@router-ce1# set bgp group to_r1 peer-as 200
user@router-ce1# set bgp group to_r1 neighbor 192.0.2.2
```

4. Configure the policies that ensure that the CE routers exchange routing information. In this example, Router CE1 exchanges routing information with Router CE2:

```
[edit policy-options]
user@router-ce1# set policy-statement my_lo0_addr term one from protocol direct
user@router-ce1# set policy-statement my_lo0_addr term one from route-filter 10.255.15.32/32
exact
user@router-ce1# set policy-statement my_lo0_addr term one then accept
user@router-ce1# set policy-statement my_lo0_addr term four from protocol direct
user@router-ce1# set policy-statement my_lo0_addr term four from route-filter 203.0.113.0/8
exact
user@router-ce1# set policy-statement my_lo0_addr term four then accept
```

Results

The output below shows the configuration of the CE1 router:

```
[edit]
user@router-ce1# show
[...Output Truncated...]
interfaces {
  ge-1/3/2 {
    unit 0 {
      family inet {
        address 192.0.2.1/24;
        address 198.51.100.2/8;
      }
    }
  }
}
routing-options {
  autonomous-system 95000;
}
protocols {
  bgp {
    group to_r1 {
      type external;
      export my_lo0_addr;
      peer-as 200;
      neighbor 192.0.2.2;
    }
  }
}
policy-options {
  policy-statement my_lo0_addr {
    term one {
      from {
        protocol direct;
        route-filter 10.255.15.32/32 exact;
      }
      then accept;
    }
    term four {
      from {
        protocol direct;
```

```

        route-filter 203.0.113.0/8 exact;
    }
    then accept;
}
}
}
}

```

Configuring the PE1 Router

Step-by-Step Procedure

In this step, you configure a filter with a matching term that contains the `then sample` statement and apply the filter to the ingress interface. You also configure a VRF routing instance with import and export policies. In addition, you configure interfaces, forwarding options, routing options, protocols, and policy options for the PE1 router. To configure the PE1 router:

1. Create the **fw** firewall filter that is applied to the logical interface being sampled:

```

[edit firewall]
user@router-pe1# set family inet filter fw term 1 from protocol tcp
user@router-pe1# set family inet filter fw term 1 from port bgp
user@router-pe1# set family inet filter fw term 1 then accept
user@router-pe1# set family inet filter fw term 2 then sample

```

2. Configure two interfaces, one interface that connects to the CE1 router (**ge-2/0/2**), and another that connects to the PE2 router (**ge-2/0/0**):

```

[edit interfaces]
user@router-pe1# set ge-2/0/2 unit 0 family inet address 192.0.2.2/24
user@router-pe1# set ge-2/0/0 unit 0 family inet address 192.168.20.1/24
user@router-pe1# set ge-2/0/0 unit 0 family mpls

```

3. Enable MPLS on the interface that connects to the PE2 router (**ge-2/0/0**):

```

[edit interfaces]
user@router-pe1# set ge-2/0/0 unit 0 family mpls

```

4. On the interface that connects to the CE1 router, apply the **fw** filter that was configured in the firewall configuration:

```
[edit interfaces]
user@router-pe1# set ge-2/0/2 unit 0 family inet filter input fw
user@router-pe1# set ge-2/0/2 unit 0 family inet filter output fw
```

5. Configure the management (**fxp0**) and loopback (**lo0**) interfaces:

```
[edit interfaces]
user@router-pe1# set fxp0 unit 0 family inet address 192.168.69.153/21
user@router-pe1# set lo0 unit 0 family inet address 127.0.0.1/32
```

6. Configure the **sampld** log file in the **/var/log** directory to record traffic sampling:

```
[edit forwarding-options]
user@router-pe1# set sampling traceoptions file sampld
user@router-pe1# set sampling traceoptions file world-readable
user@router-pe1# set sampling traceoptions flag all
```

7. Specify the sampling rate and threshold value for traffic sampling:

```
[edit forwarding-options]
user@router-pe1# set sampling input rate 1
user@router-pe1# set sampling input run-length 0
user@router-pe1# set sampling input max-packets-per-second 20000
```

8. Specify active and inactive flow periods, and the router (**198.51.100.2**) that sends out the monitored information:

```
[edit forwarding-options]
user@router-pe1# set sampling family inet output flow-active-timeout 60
user@router-pe1# set sampling family inet output flow-inactive-timeout 60
user@router-pe1# set sampling family inet output flow-server 198.51.100.2 port 2055
user@router-pe1# set sampling family inet output flow-server 198.51.100.2 local-dump
user@router-pe1# set sampling family inet output flow-server 198.51.100.2 version 500
```

9. Configure the autonomous system to establish a connection between BGP peers:

```
[edit routing-options]
user@router-pe1# set autonomous-system 200
```

10. Configure RSVP to support MPLS label-switched paths (LSPs) between the PE routers:

```
[edit protocols]
user@router-pe1# set rsvp interface all
user@router-pe1# set rsvp interface fxp0.0 disable
```

11. Configure an MPLS LSP from the PE1 router to the PE2 router:

```
[edit protocols]
user@router-pe1# set mpls label-switched-path R1toR2 from 192.168.20.1
user@router-pe1# set mpls label-switched-path R1toR2 to 192.168.20.2
user@router-pe1# set mpls interface all
user@router-pe1# set mpls interface fxp0.0 disable
```

12. Configure an internal BGP group for the PE routers. Include the family inet-vpn unicast statement to enable BGP to carry network layer reachability information (NLRI) parameters and for BGP peers to only carry unicast routes for forwarding:

```
[edit protocols]
user@router-pe1# set bgp group to_r2 type internal
user@router-pe1# set bgp group to_r2 local-address 192.168.20.1
user@router-pe1# set bgp group to_r2 neighbor 192.168.20.2 family inet-vpn unicast
```

13. Configure OSPF as the interior gateway protocol (IGP) and to compute the MPLS LSPs:

```
user@router-pe1# set ospf traffic-engineering
user@router-pe1# set ospf area 0.0.0.0 interface all
user@router-pe1# set ospf area 0.0.0.0 interface fxp0.0 disable
```

14. Create the extended community that is applied in the policy options configuration:

```
[edit policy-options]
user@router-pe1# set community vpna-comm members target:200:100
```

15. Define the **vpna-export** routing policy that is applied in the vrf-export statement in the routing instance configuration. Also, apply the **vpna-comm** community from which routes are learned:

```
[edit policy-options]
user@router-pe1# set policy-statement vpna-export term one from protocol bgp
user@router-pe1# set policy-statement vpna-export term one from protocol direct
user@router-pe1# set policy-statement vpna-export term one then community add vpna-comm
user@router-pe1# set policy-statement vpna-export term one then accept
user@router-pe1# set policy-statement vpna-export term two then reject
```

16. Define the **vpna-import** routing policy that is applied in the vrf-import statement in the routing instance configuration. Also, apply the **vpna-comm** community from which routes are learned:

```
[edit policy-options]
user@router-pe1# set policy-statement vpna-import term one from protocol bgp
user@router-pe1# set policy-statement vpna-import term one from community vpna-comm
user@router-pe1# set policy-statement vpna-import term one then accept
user@router-pe1# set policy-statement vpna-import term two then reject
```

17. Configure a VRF routing instance so that routes received from the provider edge-provider edge (PE-PE) session can be imported into any of the instance's VRF secondary routing tables:

```
[edit routing-instances]
user@router-pe1# set vrf1 instance-type vrf set vrf1 interface ge-2/0/2.0
user@router-pe1# set vrf1 route-distinguisher 10.255.15.51:1
user@router-pe1# set vrf1 vrf-import vpna-import
user@router-pe1# set vrf1 vrf-export vpna-export
user@router-pe1# set vrf1 protocols bgp group customer type external
user@router-pe1# set vrf1 protocols bgp group customer peer-as 95000
user@router-pe1# set vrf1 protocols bgp group customer as-override
user@router-pe1# set vrf1 protocols bgp group customer neighbor 192.168.30.1
user@router-pe1# set vrf1 protocols bgp group customer neighbor 192.0.2.1
```

Results

Check the results of the configuration for the PE1 router:

```
user@router-pe1> show configuration
[...Output Truncated...]
}
interfaces {
  ge-2/0/0 {
    unit 0 {
      family inet {
        address 192.168.20.1/24;
      }
      family mpls;
    }
  }
  ge-2/0/2 {
    unit 0 {
      family inet {
        filter {
          input fw;
          output fw;
        }
        address 192.0.2.2/24;
      }
    }
  }
  fxp0 {
    unit 0 {
      family inet {
        address 192.168.69.153/21;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 127.0.0.1/32;
      }
    }
  }
}
```

```

forwarding-options {
  sampling {
    traceoptions {
      file sampled world-readable;
      flag all;
    }
    input {
      rate 1;
      run-length 0;
      max-packets-per-second 20000;
    }
    family inet {
      output {
        flow-inactive-timeout 60;
        flow-active-timeout 60;
        flow-server 198.51.100.2 {
          port 2055;
          local-dump;
          version 500;
        }
      }
    }
  }
}

routing-options {
[...Output Truncated...]
  autonomous-system 200;
}

protocols {
  rsvp {
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
  mpls {
    label-switched-path R1toR2 {
      from 192.168.20.1;
      to 192.168.20.2;
    }
    interface all;
    interface fxp0.0 {
      disable;
    }
  }
}

```



```

    }
}
bgp {
    group to_r2 {
        type internal;
        local-address 192.168.20.1;
        neighbor 192.168.20.2 {
            family inet-vpn {
                unicast;
            }
        }
    }
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
}
policy-options {
    policy-statement vpna-export {
        term one {
            from protocol [ bgp direct ];
            then {
                community add vpna-comm;
                accept;
            }
        }
        term two {
            then reject;
        }
    }
    policy-statement vpna-import {
        term one {
            from {
                protocol bgp;
                community vpna-comm;
            }
            then accept;
        }
    }
}

```

```

    }
    term two {
        then reject;
    }
}
community vpna-comm members target:200:100;
}
firewall {
    family inet {
        filter fw {
            term 1 {
                from {
                    protocol tcp;
                    port bgp;
                }
                then accept;
            }
            term 2 {
                then sample;
            }
        }
    }
}
routing-instances {
    vrf1 {
        instance-type vrf;
        interface ge-2/0/2.0;
        route-distinguisher 10.255.15.51:1;
        vrf-import vpna-import;
        vrf-export vpna-export;
        protocols {
            bgp {
                group customer {
                    type external;
                    peer-as 95000;
                    as-override;
                    neighbor 192.168.30.1;
                    neighbor 192.0.2.1;
                }
            }
        }
    }
}

```

```
}
}
```

Configuring the PE2 Router

Step-by-Step Procedure

In this step, you configure a filter with a matching term that contains the `then sample` statement and apply the filter to the ingress interface. You also configure a VRF routing instance with import and export policies. In addition, you configure interfaces, forwarding options, routing options, protocols, and policy options for the PE2 router. To configure the PE2 router:

1. Create the **fw** firewall filter that is applied to the logical interface being sampled:

```
[edit firewall]
user@router-pe2# set family inet filter fw term 1 from protocol tcp
user@router-pe2# set family inet filter fw term 1 from port bgp
user@router-pe2# set family inet filter fw term 1 then accept
user@router-pe2# set family inet filter fw term 2 then sample
user@router-pe2# set family inet filter fw term 2 then accept
```

2. Configure two interfaces, one interface that connects to the CE2 router (**ge-3/1/2**), and another that connects to the PE1 router (**ge-3/1/0**):

```
[edit interfaces]
user@router-pe2# set ge-3/1/0 unit 0 family inet address 192.168.20.2/24
user@router-pe2# set ge-3/1/0 unit 0 family mpls
user@router-pe2# set ge-3/1/2 unit 0 family inet address 10.10.10.2/24
```

3. Enable MPLS on the interface that connects to the PE1 router (**ge-3/1/0**):

```
[edit interfaces]
user@router-pe2# set ge-3/1/0 unit 0 family mpls
```

4. On the interface that connects to the CE2 router, apply the **fw** filter that was configured in the firewall configuration:

```
[edit interfaces]
user@router-pe2# set ge-3/1/2 unit 0 family inet filter input fw
user@router-pe2# set ge-3/1/2 unit 0 family inet filter output fw
```

5. Configure the **sampld** log file in the **/var/log** directory to record traffic sampling:

```
[edit forwarding-options]
user@router-pe2# set sampling traceoptions file sampld
user@router-pe2# set sampling traceoptions file world-readable
user@router-pe1# set sampling traceoptions flag all
```

6. Specify the sampling rate and threshold value for traffic sampling:

```
[edit forwarding-options]
user@router-pe2# set sampling input rate 1
user@router-pe2# set sampling input run-length 0
user@router-pe2# set sampling input max-packets-per-second 20000
```

7. Specify active and inactive flow periods, and the router (**198.51.100.2**) that sends out the monitored information:

```
[edit forwarding-options]
user@router-pe2# set sampling family inet output flow-active-timeout 60
user@router-pe2# set sampling family inet output flow-inactive-timeout 60
user@router-pe2# set sampling family inet output flow-server 198.51.100.2 port 2055
user@router-pe2# set sampling family inet output flow-server 198.51.100.2 local-dump
user@router-pe2# set sampling family inet output flow-server 198.51.100.2 version 500
```

8. Configure the autonomous system to establish a connection between BGP peers:

```
[edit routing-options]
user@router-pe2# set autonomous-system 200
```

9. Configure RSVP to support MPLS label-switched paths (LSPs) between the PE routers:

```
[edit protocols]
user@router-pe2# set rsvp interface all
user@router-pe2# set rsvp interface fxp0.0 disable
```

10. Configure an MPLS LSP from the PE2 router to the PE1 router:

```
[edit protocols]
user@router-pe2# set mpls label-switched-path R2toR1 from 192.168.20.2
user@router-pe2# set mpls label-switched-path R2toR1 to 192.168.20.1
user@router-pe2# set mpls interface all
user@router-pe2# set mpls interface fxp0.0 disable
```

11. Configure an internal BGP group for the PE routers. Include the `family inet-vpn unicast` statement to enable BGP to carry network layer reachability information (NLRI) parameters and for BGP peers to only carry unicast routes for forwarding:

```
[edit protocols]
user@router-pe2# set bgp group to_r1 type internal
user@router-pe2# set bgp group to_r1 local-address 192.168.20.2
user@router-pe2# set bgp group to_r1 neighbor 192.168.20.1 family inet-vpn unicast
```

12. Configure OSPF as the interior gateway protocol (IGP) and to compute the MPLS LSPs:

```
[edit protocols]
user@router-pe2# set ospf traffic-engineering
user@router-pe2# set ospf area 0.0.0.0 interface all
user@router-pe2# set ospf area 0.0.0.0 interface fxp0.0 disable
```

13. Create the extended community that is applied in the policy options configuration:

```
[edit policy-options]
user@router-pe2# set community vpna-comm members target:200:100
```

14. Define the **vpna-export** routing policy that is applied in the `vrf-export` statement in the routing instance configuration. Also, apply the **vpna-comm** community from which routes are learned:

```
[edit policy-options]
user@router-pe2# set policy-statement vpna-export term one from protocol bgp
user@router-pe2# set policy-statement vpna-export term one from protocol direct
user@router-pe2# set policy-statement vpna-export term one then community add vpna-comm
user@router-pe2# set policy-statement vpna-export term one then accept
user@router-pe2# set policy-statement vpna-export term two then reject
```

15. Define the **vpna-import** routing policy that is applied in the `vrf-import` statement in the routing instance configuration. Also, apply the **vpna-comm** community from which routes are learned:

```
[edit policy-options]
user@router-pe2# set policy-statement vpna-import term one from protocol bgp
user@router-pe2# set policy-statement vpna-import term one from community vpna-comm
user@router-pe2# set policy-statement vpna-import term one then accept
user@router-pe2# set policy-statement vpna-import term two then reject
```

16. Configure a VRF routing instance so that routes received from the provider edge-provider edge (PE-PE) session can be imported into any of the instance's VRF secondary routing tables:

```
[edit routing-instances]
user@router-pe2# set vrf1 instance-type vrf
user@router-pe2# set vrf1 interface ge-3/1/2.0
user@router-pe2# set vrf1 route-distinguisher 10.255.19.12:1
user@router-pe2# set vrf1 vrf-import vpna-import
user@router-pe2# set vrf1 vrf-export vpna-export
user@router-pe2# set vrf1 protocols bgp group R3-R4 type external
user@router-pe2# set vrf1 protocols bgp group R3-R4 peer-as 65000
user@router-pe2# set vrf1 protocols bgp group R3-R4 as-override
user@router-pe2# set vrf1 protocols bgp group R3-R4 neighbor 10.10.10.1
```

Results

Check the results of the configuration for the PE2 router:

```
user@router-pe2> show configuration
[...Output Truncated...]
}
interfaces {
  ge-3/1/0 {
    unit 0 {
      family inet {
        address 192.168.20.2/24;
      }
      family mpls;
    }
  }
  ge-3/1/2 {
    unit 0 {
      family inet {
        filter {
          input fw;
          output fw;
        }
        address 10.10.10.2/24;
      }
    }
  }
}
forwarding-options {
  sampling {
    traceoptions {
      file sampled world-readable;
      flag all;
    }
    input {
      rate 1;
      run-length 0;
      max-packets-per-second 20000;
    }
    family inet {
      output {
        flow-inactive-timeout 60;
      }
    }
  }
}
```

```

        flow-active-timeout 60;
        flow-server 198.51.100.2 {
            port 2055;
            local-dump;
            version 500;
        }
    }
}

routing-options {
[...Output Truncated...]
    autonomous-system 200;
}

protocols {
    rsvp {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    mpls {
        label-switched-path R2toR1 {
            from 192.168.20.2;
            to 192.168.20.1;
        }
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
    bgp {
        group to_r1 {
            type internal;
            local-address 192.168.20.2;
            neighbor 192.168.20.1 {
                family inet-vpn {
                    unicast;
                }
            }
            neighbor 192.0.2.1;
        }
    }
}

```



```

ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
}
policy-options {
    policy-statement vpna-export {
        term one {
            from protocol [ bgp direct ];
            then {
                community add vpna-comm;
                accept;
            }
        }
        term two {
            then reject;
        }
    }
    policy-statement vpna-import {
        term one {
            from {
                protocol bgp;
                community vpna-comm;
            }
            then accept;
        }
        term two {
            then reject;
        }
    }
    community vpna-comm members target:200:100;
}
firewall {
    family inet {
        filter fw {
            term 1 {
                from {
                    protocol tcp;

```


1. Configure one interface with two IP addresses. One address is for traffic to the PE2 router and the other address is to check that traffic is flowing from the CE1 router:

```
[edit interfaces]
user@router-ce2# set ge-0/1/2 unit 0 family inet address 10.10.10.1/24
user@router-ce2# set ge-0/1/2 unit 0 family inet address 10.4.4.4/16
```

2. Configure the autonomous system to establish a connection between BGP peers:

```
[edit routing-options]
user@router-ce1# set autonomous-system 65000
```

3. Configure BGP as the routing protocol between the CE and the PE routers:

```
[edit protocols]
user@router-ce2# set bgp group R3-R4 type external
user@router-ce2# set bgp group R3-R4 export l3vpn-policy
user@router-ce2# set bgp group R3-R4 peer-as 200
user@router-ce2# set bgp group R3-R4 neighbor 10.10.10.2
```

4. Configure the policies that ensure that the CE routers exchange routing information. In this example, Router CE2 exchanges routing information with Router CE1:

```
[edit policy-options]
user@router-ce2# set policy-statement l3vpn-policy term one from protocol direct
user@router-ce2# set policy-statement l3vpn-policy term one from route-filter
10.255.15.75/32 exact
user@router-ce2# set policy-statement l3vpn-policy term one then accept
user@router-ce2# set policy-statement l3vpn-policy term two from protocol direct
user@router-ce2# set policy-statement l3vpn-policy term two from route-filter 10.4.0.0/16
exact
user@router-ce2# set policy-statement l3vpn-policy term two then accept
```

Results

The output below shows the configuration of the CE2 router:

```
[edit]
user@router-ce2# show
[...Output Truncated...]
interfaces {
  ge-0/1/2 {
    unit 0 {
      family inet {
        address 10.10.10.1/24;
        address 10.4.4.4/16;
      }
    }
  }
}
routing-options {
  autonomous-system 65000;
}
protocols {
  bgp {
    group R3-R4 {
      type external;
      export l3vpn-policy;
      peer-as 200;
      neighbor 10.10.10.2;
    }
  }
}
policy-options {
  policy-statement l3vpn-policy {
    term one {
      from {
        protocol direct;
        route-filter 10.255.15.75/32 exact;
      }
      then accept;
    }
    term two {
      from {
        protocol direct;
```

```

        route-filter 10.4.0.0/16 exact;
    }
    then accept;
}
}
}
}

```

Verification

IN THIS SECTION

- [Verifying the Traffic Flow Between the CE Routers | 122](#)
- [Verifying Sampled Traffic | 123](#)
- [Cross Verifying Sampled Traffic | 124](#)

After you have completed the configuration of the four routers, you can verify that traffic is flowing from the CE1 router to the CE2 router, and you can observe the sampled traffic from two locations. To confirm that the configuration is working properly, perform these tasks:

Verifying the Traffic Flow Between the CE Routers

Purpose

Use the `ping` command to verify traffic between the CE routers.

Action

From the CE1 router, issue the `ping` command to the CE2 router:

```

user@router-ce2> ping 10.4.4.4 source 198.51.100.2
PING 10.4.4.4 (10.4.4.4): 56 data bytes
64 bytes from 10.4.4.4: icmp_seq=0 ttl=64 time=0.861 ms
64 bytes from 10.4.4.4: icmp_seq=1 ttl=64 time=0.869 ms
64 bytes from 10.4.4.4: icmp_seq=2 ttl=64 time=0.786 ms
^C
--- 10.4.4.4 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss

```

```
round-trip min/avg/max/stddev = 0.786/0.839/0.869/0.037 ms
```

Meaning

The output from the ping command shows that the ping command was successful. Traffic is flowing between the CE routers.

Verifying Sampled Traffic

Purpose

You can observe the sampled traffic using the `show log sampled` command from the CLI or from the router shell using the `tail -f /var/log/sampled` command. In addition, you can collect the logs in a flowcollector. The same information appears in the output of both commands and in the flow collector. For information about using a flow collector, see [“Sending cflowd Records to Flow Collector Interfaces” on page 239](#) and [“Example: Configuring a Flow Collector Interface on an M, MX or T Series Router” on page 208](#).

Action

From the PE1 router, use the `show log sampled` command:

```
user@router-pe1> show log sampled
[...Output Truncated...]
Nov 16 23:24:19   Src addr: 198.51.100.2
Nov 16 23:24:19   Dst addr: 10.4.4.4
Nov 16 23:24:19   Nhop addr: 192.168.20.2
Nov 16 23:24:19   Input interface: 503      # SNMP index of the incoming interface on PE1
Nov 16 23:24:19   Output interface: 505     # SNMP index of the outgoing interface on PE1
Nov 16 23:24:19   Pkts in flow: 5
Nov 16 23:24:19   Bytes in flow: 420
Nov 16 23:24:19   Start time of flow: 602411369
Nov 16 23:24:19   End time of flow: 602415369
Nov 16 23:24:19   Src port: 0
Nov 16 23:24:19   Dst port: 2048
Nov 16 23:24:19   TCP flags: 0x0
Nov 16 23:24:19   IP proto num: 1
Nov 16 23:24:19   TOS: 0x0
Nov 16 23:24:19   Src AS: 95000      # The autonomous system of CE1
```

```

Nov 16 23:24:19   Dst AS: 65000,,,,,# The autonomous system of CE2
Nov 16 23:24:19   Src netmask len: 8
Nov 16 23:24:19   Dst netmask len: 16
Nov 16 23:24:19 cflowd header:
Nov 16 23:24:19   Num-records: 1
Nov 16 23:24:19   Version: 500
Nov 16 23:24:19   Flow seq num: 13
Nov 16 23:24:19   Sys Uptime: 602450382 (msecs)
Nov 16 23:24:19   Time-since-epoch: 1258413859 (secs)
Nov 16 23:24:19   Engine id: 0
Nov 16 23:24:19   Engine type: 0
Nov 16 23:24:19   Sample interval: 1
[...Output Truncated...]

```

Meaning

The output from the `show log sampled` command shows the correct SNMP index for the incoming and outgoing interfaces on the PE1 router. Also, the source and destination addresses for the autonomous systems for the two CE routers are correct.

Cross Verifying Sampled Traffic

Purpose

You can also double check that the sampled traffic is the correct traffic by using the `show interface interface-name-fpc/pic/port.unit-number | match SNMP` command and the `show route route-name detail` command.

Action

The following output is a cross check of the output in the ["Verifying Sampled Traffic" on page 123](#) task:

```

user@router-pe1> show interfaces ge-2/0/2.0 | match SNMP
Logical interface ge-2/0/2.0 (Index 76) (SNMP ifIndex 503)
Flags: SNMP-Traps 0x4000000 Encapsulation: ENET2

```

```

user@router-pe1> show route 10.4.4.4 detail

vrf1.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
10.4.0.0/16 (1 entry, 1 announced)

```

```

*BGP      Preference: 170/-101
          Route Distinguisher: 10.255.19.12:1
          Next hop type: Indirect
          Next-hop reference count: 6
          Source: 192.168.20.2
          Next hop type: Router, Next hop index: 659
          Next hop: 192.168.20.2 via ge-2/0/0.0 weight 0x1, selected
          Label operation: Push 299776
          Protocol next hop: 192.168.20.2
          Push 299776
          Indirect next hop: 8e6f780 1048574
          State: <Secondary Active Int Ext>
          Local AS: 200 Peer AS: 200
          Age: 3d 19:49:32 Metric2: 65535
          Task: BGP_200.20.20.20.2+179
          Announcement bits (3): 0-RT 1-BGP RT Background 2-KRT
AS path: 65000 I
          AS path: Recorded
          Communities: target:200:100
          Import Accepted
          VPN Label: 299776
          Localpref: 100
          Router ID: 10.10.10.2
          Primary Routing Table bgp.l3vpn.0

```

Meaning

The output of the `show interfaces ge-2/0/2.0 | match SNMP` command shows that the SNMP ifIndex field has the same value (**503**) as the output for the `show log sampled` command in the ["Verifying Sampled Traffic" on page 123](#) task, indicating that the intended traffic is being sampled.

The output of the `show route 10.4.4.4 detail` command shows that the source address **10.4.4.4**, the source mask (**16**), and the source AS (**65000**) have the same values as the output for the `show log sampled` command in the ["Verifying Sampled Traffic" on page 123](#) task, indicating that the intended traffic is being sampled.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

Configuring an Aggregate Export Timer on M, MX and T Series Routers for Version 8 Records

When you use flow monitoring version 8 records for active flow monitoring, you can configure an aggregate export timer. To configure this timer, include the `aggregate-export-interval` statement at the `[edit forwarding-options sampling output]` hierarchy level. The timer value has a default minimum setting of 90 seconds and a maximum value of 1800 seconds.


```
[edit]
forwarding-options {
  sampling {
    output {
      aggregate-export-interval duration;
    }
  }
}
```

Example: Sampling Configuration for M, MX and T Series Routers

IN THIS SECTION

- Verifying Your Work | 130

Figure 17: Active Flow Monitoring—Sampling Configuration Topology Diagram



In [Figure 17 on page 127](#), traffic from Router 1 arrives on the monitoring router's Gigabit Ethernet `ge-2/3/0` interface. The exit interface on the monitoring router that leads to destination Router 2 is `ge-3/0/0`. In active flow monitoring, both the input interface and exit interface can be any interface type (such as SONET/SDH, Gigabit Ethernet, and so on). The export interface leading to the flow server is `fe-1/0/0`.

Configure a firewall filter to sample, count, and accept all traffic. Apply the filter to the input interface, and configure the exit interface (for traffic forwarding), the adaptive services interface (for flow processing), and the export interface (for exporting flow records).

Configure sampling at the [edit forwarding-options] hierarchy level. Include the IP address and port of the flow server with the flow-server statement and specify the adaptive services interface to be used for flow record processing with the interface statement at the [edit forwarding-options sampling] hierarchy level.

Router 1

```
[edit]
interfaces {
    sp-2/0/0 { # This adaptive services interface creates the flow records.
        unit 0 {
            family inet {
                address 10.5.5.1/32 {
                    destination 10.5.5.2;
                }
            }
        }
    }

    fe-1/0/0 { # This is the interface where records are sent to the flow
server.
        unit 0 {
            family inet {
                address 10.60.2.2/30;
            }
        }
    }

    ge-2/3/0 { # This is the input interface where all traffic enters the
router.
        unit 0 {
            family inet {
                filter {
                    input catch_all; # This is where the firewall filter
is applied.
                }
                address 10.1.1.1/20;
            }
        }
    }

    ge-3/0/0 { # This is the interface where the original traffic is forwarded.
        unit 0 {
```

```

        family inet {
            address 10.2.2.1/24;
        }
    }
}
forwarding-options {
    sampling { # Traffic is sampled and sent to a flow server.
        input {
            rate 1; # Samples 1 out of
x           packets (here, a rate of 1 sample per packet).
        }
    }
    family inet {
        output {
            flow-server 10.60.2.1 { # The IP address and port of the flow server.
                port 2055;
                version 5; # Records are sent to the flow server using version 5 format.
            }
            flow-inactive-timeout 15;
            flow-active-timeout 60;
            interface sp-2/0/0 { # Adding an interface here enables PIC-based sampling.
                engine-id 5; # Engine statements are dynamic, but can be configured.
                engine-type 55;
                source-address 10.60.2.2; # You must configure this statement.
            }
        }
    }
}
firewall {
    family inet {
        filter catch_all { # Apply this filter on the input interface.
            term default {
                then {
                    sample;
                    count counter1;
                    accept;
                }
            }
        }
    }
}

```

Verifying Your Work

To verify that your configuration is correct, use the following commands on the monitoring station that is configured for active flow monitoring:

- `show services accounting errors`
- `show services accounting (flow | flow-detail)`
- `show services accounting memory`
- `show services accounting packet-size-distribution`
- `show services accounting status`
- `show services accounting usage`
- `show services accounting aggregation template template-name name (detail | extensive | terse) (version 9 only)`

Most active flow monitoring operational mode commands contain equivalent output information to the following passive flow monitoring commands:

- `show services accounting errors = show passive-monitoring error`
- `show services accounting flow = show passive-monitoring flow`
- `show services accounting memory = show passive-monitoring memory`
- `show services accounting status = show passive-monitoring status`
- `show services accounting usage = show passive-monitoring usage`

The active flow monitoring commands can be used with most active flow monitoring applications, including sampling, discard accounting, port mirroring, and multiple port mirroring. However, you can use the passive flow monitoring commands only with configurations that contain a monitoring group at the `[edit forwarding-options monitoring]` hierarchy level.

The following shows the output of the `show` commands used with the configuration example:

```
user@router1> show services accounting errors
Service Accounting interface: sp-2/0/0, Local interface index: 542
Service name: (default sampling)
Error information
  Packets dropped (no memory): 0, Packets dropped (not IP): 0
  Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
  Memory allocation failures: 0, Memory free failures: 0
  Memory free list failures: 0
```

Memory overload: No, PPS overload: No, BPS overload: Yes

user@router1> **show services accounting flow-detail limit 10**

Service Accounting interface: sp-2/0/0, Local interface index: 468

Service name: (default sampling)

Protocol	Source Address	Source Port	Destination Address	Destination Port	Packet count	Byte count
udp(17)	10.1.1.2	53	10.0.0.1	53	4329	3386035
ip(0)	10.1.1.2	0	10.0.0.2	0	4785	3719654
ip(0)	10.1.1.2	0	10.0.1.2	0	4530	3518769
udp(17)	10.1.1.2	0	10.0.7.1	0	5011	3916767
tcp(6)	10.1.1.2	20	10.3.0.1	20	1	1494
tcp(6)	10.1.1.2	20	10.168.80.1	20	1	677
tcp(6)	10.1.1.2	20	10.69.192.1	20	1	446
tcp(6)	10.1.1.2	20	10.239.240.1	20	1	1426
tcp(6)	10.1.1.2	20	10.126.160.1	20	1	889
tcp(6)	10.1.1.2	20	10.71.224.1	20	1	1046

user@router1> **show services accounting memory**

Service Accounting interface: sp-2/0/0, Local interface index: 468

Service name: (default sampling)

Memory utilization

Allocation count: 437340, Free count: 430681, Maximum allocated: 6782

Allocations per second: 3366, Frees per second: 6412

Total memory used (in bytes): 133416928, Total memory free (in bytes): 133961744

user@router1> **show services accounting packet-size-distribution**

Service Accounting interface: sp-2/0/0, Local interface index: 468

Service name: (default sampling)

Range start	Range end	Number of packets	Percentage packets
64	96	1705156	100

user@router1> **show services accounting status**

Service Accounting interface: sp-2/0/0, Local interface index: 468

Service name: (default sampling)

Interface state: Monitoring

Group index: 0

Export interval: 60 secs, Export format: cflowd v5

Protocol: IPv4, Engine type: 55, Engine ID: 5

Route record count: 13, IFL to SNMP index count: 30, AS count: 1

Time set: Yes, Configuration set: Yes

Route record set: Yes, IFL SNMP map set: Yes

```

user@router1> show services accounting usage
Service Accounting interface: sp-2/0/0, Local interface index: 468
Service name: (default sampling)
CPU utilization
  Uptime: 4790345 milliseconds, Interrupt time: 1668537848 microseconds
  Load (5 second): 71%, Load (1 minute): 63%

```

Associating Sampling Instances for Active Flow Monitoring with a Specific FPC, MPC, or DPC

The Junos OS enables you to configure sampling instances for active flow monitoring, by specifying a name for the sampling parameters and associating the instance name with a specific FPC, MPC, or DPC.

To configure active sampling instances, include the `instance` statement at the `[edit forwarding-options sampling]` hierarchy level. For more information about configuring sampling instances, see the [Junos OS Services Interfaces Library for Routing Devices](#).

To associate a configured active sampling instance with a specific FPC, MPC, or DPC, include the sampling instance name at the `[edit chassis fpc slot-number]` hierarchy level:

```

[edit chassis fpc slot-number]
sampling-instance instance-name;

```

On a TX Matrix, TX Matrix Plus router, include the `sampling-instance` statement at the `[edit chassis lcc number fpc slot-number]` hierarchy level:

```

[edit chassis lcc number fpc slot-number]
sampling-instance instance-name;

```

RELATED DOCUMENTATION

[Example: Sampling Instance Configuration | 133](#)

[sampling-instance | 1385](#)

Example: Sampling Instance Configuration

IN THIS SECTION


- [Example Network Details | 134](#)
- [Example Router Configuration | 135](#)
- [Configuration Commands Used for the Configuration Example | 138](#)
- [Verifying Your Work | 139](#)

You can configure active sampling using a sampling instance and associate that sampling instance to a particular Flexible Port Concentrator (FPC), Modular Port Concentrator (MPC), or Dense Port Concentrator (DPC). In addition, you can define multiple sampling instances associated with multiple destinations and protocol families per sampling instance destination.

Example Network Details

The following example shows the configuration of two sampling instances on an MX480 router running Junos OS Release 9.6.

Figure 18: Active Flow Monitoring—Sampling Instance Configuration Topology Diagram



In [Figure 18 on page 134](#), packets from Router 1 arrive on the monitoring router's Gigabit Ethernet ge-0/1/0 interface, the packets are sampled by the services interface sp-2/0/0 and sent to the cflowd server by the export interface ge-1/0/4. Packets from Router 3 arrive on the monitoring router's Gigabit Ethernet ge-3/1/0 interface, the packets are sampled by the services interface sp-2/1/0 and sent to the

cflowd server by the export interface ge-1/0/4. Normal traffic flow from ge-0/1/0 and ge-3/1/0 to ge-1/0/0 and on to Router 2 continues undisturbed during the sampling process. In active flow monitoring, both the input interface and exit interface can be any interface type (such as SONET/SDH, Gigabit Ethernet, and so on).

Only one sampling instance can be attached to an FPC, MPC, or DPC. Multiple families can be configured under a sampling instance. Each family can have its own collector address. You can define sampling instances and attach each instance to different FPCs, or a single sampling instance can be attached to all FPCs.

The sampling configuration for this example includes the following:

- Two sampling instances, s0 and s1, configured to collect sampling data at the [edit forwarding-options] hierarchy level. The flow-server statement includes the IP address, port, and template of the flow server. The interface statement includes the services interface, sp-2/0/0 or sp-2/1/0, for flow record processing, and the source address of the incoming router on the sampled interface.
- The binding of the two sampling instances to FPCs 0 and 3. These are configured with the sampling-instance statement at the [edit chassis fpc slot] hierarchy level.
- Sampling activated on the input interfaces ge-0/1/0 and ge-3/1/0 using the sampling statement at the [edit interfaces interface-name unit unit-number family family] hierarchy level.

In this example, the ping command is issued on Router 1 to Router 2 via the MX480 router to generate traffic. After the packets are generated, show commands are issued to verify that the sampling configuration is working as expected.

Example Router Configuration

The following output shows the configuration of an MX480 router with two sampling instances.

```
user@MX480-router> show configuration
[...Output Truncated...]
}
chassis {
    fpc 0 { # The fpc number is associated with the interface on which sampling is enabled,
ge-0/1/0 in this statement.
        sampling-instance s0;
    }
    fpc 3 { # The fpc number is associated with the interface on which sampling is enabled,
ge-3/1/0 in this statement.
        sampling-instance s1;
    }
}
```

```

interfaces {
    ge-0/1/0 { # This interface has sampling activated.
        unit 0 {
            family inet {
                sampling { # Here sampling is activated.
                    input;
                }
                address 10.0.0.1/30;
            }
        }
    }
    ge-1/0/0 { # The interface on which packets are exiting the router.
        unit 0 {
            family inet {
                address 192.0.2.1/30;
            }
        }
    }
    ge-1/0/4 { # The interface connected to the cflowd server.
        unit 0 {
            family inet {
                address 198.51.100.1/32;
            }
        }
    }
    sp-2/0/0 { # The service interface that samples the packets from Router 1.
        unit 0 {
            family inet;
        }
    }
    sp-2/1/0 { # The service interface that samples the packets from Router 3.
        unit 0 {
            family inet;
        }
    }
    ge-3/1/0 { # This interface has sampling activated.
        unit 0 {
            family inet {
                sampling { # Here sampling is activated.
                    input;
                }
                address 192.168.2.1/30;
            }
        }
    }
}

```

```

    }
  }
}
forwarding-options {
  sampling {
    instance {
      s0 {
        input {
          rate 1;
          run-length 0;
        }
        family inet {
          output {
            flow-server 198.51.100.2 { # The address of the external server.
              port 2055;
              version9 {
                template {
                  v4
                }
              }
            }
          }
          interface sp-2/0/0 {
            source-address 192.168.1.1; # Source address of the sampled packets
          }
        }
      }
    }
  }
}
s1 {
  input {
    rate 1;
    run-length 0;
  }
  family inet {
    output {
      flow-server 198.51.100.2 { # The address of the external server.
        port 2055;
        version9 {
          template {
            v4
          }
        }
      }
    }
    interface sp-2/1/0 {

```

```

        source-address 192.168.1.2; # Source address of the sampled packets
    }
}
}
}
}
}
}
}

routing-options {
    static {
        route 203.0.113.0/8 next-hop 192.0.2.2;
    }
}

services {
    flow-monitoring {
        version9 {
            template v4 {
                flow-active-timeout 30;
                flow-inactive-timeout 30;
                ipv4-template;
            }
        }
    }
}
}

```

Configuration Commands Used for the Configuration Example

The following set commands are used for the configuration of the sampling instance in this example. Replace the values in these commands with values relevant to your own network.

- set chassis fpc 0 sampling-instance s0
- set chassis fpc 3 sampling-instance s1
- set interfaces ge-0/1/0 unit 0 family inet sampling input
- set interfaces ge-0/1/0 unit 0 family inet address
- set interfaces ge-1/0/0 unit 0 family inet address
- set interfaces sp-2/0/0 unit 0 family inet
- set interfaces sp-2/1/0 unit 0 family inet

- set interfaces ge-3/1/0 unit 0 family inet sampling input
- set interfaces ge-3/1/0 unit 0 family inet address
- set forwarding-options sampling instance s0 input rate 1
- set forwarding-options sampling instance s0 input run-length 0
- set forwarding-options sampling instance s0 family inet output flow-server 198.51.100.2 port 2055
- set forwarding-options sampling instance s0 family inet output flow-server 198.51.100.2 version9 template v4;
- set forwarding-options sampling instance s0 family inet output interface sp-2/0/0 source-address 192.168.1.1
- set forwarding-options sampling instance s1 input rate 1
- set forwarding-options sampling instance s1 input run-length 0
- set forwarding-options sampling instance s1 family inet output flow-server 198.51.100.2 port 2055
- set forwarding-options sampling instance s1 family inet output flow-server 198.51.100.2 version9 template v4;
- set forwarding-options sampling instance s1 family inet output interface sp-2/1/0 source-address 192.168.1.2
- set routing-options static route 203.0.113.0/8 next-hop 192.0.2.2
- set services flow-monitoring version9 template v4 flow-active-timeout 30
- set services flow-monitoring version9 template v4 flow-inactive-timeout 30
- set services flow-monitoring version9 template v4 ipv4-template

Verifying Your Work

To verify that your configuration is working as expected, use the following commands on the router that is configured with the sampling instance:

- show services accounting aggregation template template-name *template-name*
- show services accounting flow

The following shows the output of the show commands issued on the MX480 router used in this configuration example:

```
user@MX480-router> show services accounting aggregation template template-name v4
          Src  Dst
          Port/ Port/
```

Source Address	Destination Address	ICMP Type	ICMP Code	Proto	TOS	Packet Count
10.0.0.6	203.0.113.3	100	1000	17	8	14
10.0.0.5	203.0.113.2	100	1000	17	8	15
10.0.0.3	203.0.113.3	100	1000	17	8	15
10.0.0.2	203.0.113.3	100	1000	17	8	15
10.0.0.4	203.0.113.2	100	1000	17	8	15
10.0.0.6	203.0.113.2	100	1000	17	8	15
10.0.0.4	203.0.113.3	100	1000	17	8	15
10.0.0.2	203.0.113.2	100	1000	17	8	16
10.0.0.3	203.0.113.2	100	1000	17	8	15
10.0.0.5	203.0.113.3	100	1000	17	8	15

```
user@MX480-router> show services accounting aggregation template template-name v4
```

		Src Dst					
		Port/ Port/					
Source Address	Destination Address	ICMP Type	ICMP Code	Proto	TOS	Packet Count	
10.0.0.6	203.0.113.3	100	1000	17	8	16	
10.0.0.5	203.0.113.2	100	1000	17	8	17	
10.0.0.3	203.0.113.3	100	1000	17	8	16	
10.0.0.2	203.0.113.3	100	1000	17	8	16	
10.0.0.4	203.0.113.2	100	1000	17	8	17	
10.0.0.6	203.0.113.2	100	1000	17	8	17	
10.0.0.4	203.0.113.3	100	1000	17	8	16	
10.0.0.2	203.0.113.2	100	1000	17	8	17	
10.0.0.3	203.0.113.2	100	1000	17	8	17	
10.0.0.5	203.0.113.3	100	1000	17	8	16	

```
user@MX480-router> show services accounting flow
```

Flow information

Interface name: sp-2/0/0, Local interface index: 152

Flow packets: 884, Flow bytes: **56576**

Flow packets 10-second rate: 0, Flow bytes 10-second rate: 628

Active flows: 10, Total flows: 35

Flows exported: 75, Flows packets exported: 14

Flows inactive timed out: 25, Flows active timed out: 75

```
user@MX480-router> show services accounting flow
```

Flow information

Interface name: sp-2/0/0, Local interface index: 152

Flow packets: 898, Flow bytes: **57472**

Flow packets 10-second rate: 0, Flow bytes 10-second rate: 628

```
Active flows: 10, Total flows: 35  
Flows exported: 75, Flows packets exported: 14  
Flows inactive timed out: 25, Flows active timed out: 75
```

RELATED DOCUMENTATION

[Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches | 433](#)

[Configuring Active Flow Monitoring | 42](#)

[sampling-instance | 1385](#)

Example: Sampling and Discard Accounting Configuration on M, MX and T Series Routers


IN THIS SECTION

- [Verifying Your Work | 146](#)

Discard accounting allows you to sample traffic, send it to a flow server for analysis, and discard all packets without forwarding them to their intended destination. Discard accounting is enabled with the `discard accounting group-name` statement in a firewall filter at the [edit firewall family inet filter *filter-name* term *term-name* then] hierarchy level. Then, the filter is applied to an interface with the `filter` statement at

the [edit interfaces *interface-name* unit *unit-number* family inet] hierarchy level and processed with the output statement at the [edit forwarding-options accounting *group-name*] hierarchy level.

Figure 19: Active Flow Monitoring—Sampling and Discard Accounting Topology Diagram



In [Figure 19 on page 142](#), traffic from Router 1 arrives on the monitoring router's Gigabit Ethernet **ge-2/3/0** interface. The export interface leading to the flow server is **fe-1/0/0** and there is no exit interface.

In this example, TCP traffic is sent to one accounting group and all other traffic is diverted to a second group. After being sampled and counted, the two types of traffic are acted upon by the sampling and accounting processes. These processes create flow records and send the records to the version 8 flow server for analysis. Because multiple types of traffic are sent to the same server, we recommend that you configure the **engine-id**, **engine-type**, and source-address statements manually in your accounting and sampling hierarchies. This way, you can differentiate between traffic types when they arrive at the flow server.

```
[edit]
interfaces {
  sp-2/0/0 { # This adaptive services interface creates the flow records.
    unit 0 {
```

```

        family inet {
            address 10.5.5.1/32 {
                destination 10.5.5.2;
            }
        }
    }
}

    fe-1/0/0 { # This is the interface where records are sent to the flow server.
    unit 0 {
        family inet {
            address 10.60.2.2/30;
        }
    }
}

    ge-2/3/0 { # This is the input interface where traffic enters the router.
    unit 0 {
        family inet {
            filter {
                input catch_all;
            }
            address 10.11.1.1/30;
        }
    }
}
}

forwarding-options {
    sampling { # The router samples the traffic.
        input {
            rate 100; # One out of every 100 packets is sampled.
        }
    }
    family inet {
        output { # The sampling process creates and exports flow
records.
            flow-server 10.60.2.1 { # You can configure a variety of
settings.
                port 2055;
                version 8;
                aggregation { # Aggregation is unique to flow version 8.
                    protocol-port;
                    source-destination-prefix;
                }
            }
        }
    }
}

```

```

        aggregate-export-interval 90;
        flow-inactive-timeout 60;
        flow-active-timeout 60;

        interface sp-2/0/0 { # This statement enables PIC-based
sampling.
        engine-id 5; # Engine statements are dynamic, but can be configured.
        engine-type 55;
        source-address 10.60.2.2; # You must configure this
statement.
        }
    }
}

    accounting counter1 { # This discard accounting process handles default
traffic.
        output { # This process creates and exports flow records.
            flow-inactive-timeout 65;
            flow-active-timeout 65;
            flow-server 10.60.2.1 { # You can configure a variety of settings.
                port 2055;
                version 8;
                aggregation { # Aggregation is unique to version 8.
                    protocol-port;
                    source-destination-prefix;
                }
            }
        }

        interface sp-2/0/0 { # This statement enables PIC-based discard
accounting.
        engine-id 1; # Engine statements are dynamic, but can be configured.
        engine-type 11;
        source-address 10.60.2.3; # You must configure this statement.
        }
    }
}

    accounting t2 { # The second discard accounting process handles the TCP
traffic.
        output { # This process creates and exports flow records.
            aggregate-export-interval 90;
            flow-inactive-timeout 65;
            flow-active-timeout 65;
            flow-server 10.60.2.1 { # You can configure a
variety of settings for the server.
                port 2055;
                version 8;

```

```

        aggregation { # Aggregation is unique to version 8.
            protocol-port;
            source-destination-prefix;
        }
    }

    interface sp-2/0/0 { # This statement enables PIC-based discard
accounting.
        engine-id 2; # Engine statements are dynamic, but can be configured.
        engine-type 22;
            source-address 10.60.2.4;# You must configure this statement.
        }
    }
}
firewall {
    family inet {
        filter catch_all { # Apply the firewall filter on the input interface.
            term t2 { # This places TCP traffic into one group for sampling
and
                from { # discard accounting.
                    protocol tcp;
                }
                then {
                    count c2;# The count action counts traffic as it enters
the router.
                    sample; # The sample action sends the traffic to the
sampling process.
                    discard accounting t2; # The discard accounting
discards traffic.
                }
            }
            term default { # Performs sampling and discard accounting on all
other traffic.
                then {
                    count counter; # The count action counts traffic as it
enters the router.
                    sample# The sample action sends the traffic to the
sampling process.
                    discard accounting counter1; # This activates discard
accounting.
                }
            }
        }
    }
}

```

```

    }
}

```

Verifying Your Work

To verify that your configuration is correct, use the following commands on the monitoring station that is configured for active flow monitoring:

- `show services accounting aggregation (for version 8 flows only)`
- `show services accounting errors`
- `show services accounting (flow | flow-detail)`
- `show services accounting memory`
- `show services accounting packet-size-distribution`
- `show services accounting status`
- `show services accounting usage`

The following shows the output of the `show` commands used with the configuration example:

```

user@host> show services accounting flow name t2
Service Accounting interface: sp-2/0/0, Local interface index: 468
Service name: t2
  Flow information
    Flow packets: 56130820, Flow bytes: 3592372480
    Flow packets 10-second rate: 13024, Flow bytes 10-second rate: 833573
    Active flows: 600, Total flows: 600
    Flows exported: 28848, Flows packets exported: 960
    Flows inactive timed out: 0, Flows active timed out: 35400

user@host> show services accounting
Service Name:
  (default sampling)
  counter1
  t2

user@host> show services accounting aggregation protocol-port detail name t2
Service Accounting interface: sp-2/0/0, Local interface index: 468
Service name: t2

```

Protocol: 6, Source port: 20, Destination port: 20
 Start time: 442794, End time: 6436260
 Flow count: 1, Packet count: 4294693925, Byte count: 4277471552

user@host> **show services accounting aggregation source-destination-prefix name**

t2 limit 10 order packets

Service Accounting interface: sp-2/0/0, Local interface index: 542

Service name: t2

Source Prefix	Destination Prefix	Input SNMP Index	Output SNMP Index	Flow count	Packet count	Byte count
10.1.1.2/20	10.225.0.1/0	24	26	0	13	9650
10.1.1.2/20	10.143.80.1/0	24	26	0	13	10061
10.1.1.2/20	10.59.176.1/0	24	26	0	13	10426
10.1.1.2/20	10.5.32.1/0	24	26	0	13	12225
10.1.1.2/20	10.36.16.1/0	24	26	0	13	9116
10.1.1.2/20	10.1.96.1/0	24	26	0	12	11050
10.1.1.2/20	10.14.48.1/0	24	26	0	13	10812
10.1.1.2/20	10.31.192.1/0	24	26	0	13	11473
10.1.1.2/20	10.129.144.1/0	24	26	0	13	7647
10.1.1.2/20	10.188.160.1/0	24	26	0	13	10056

user@host> **show services accounting aggregation source-destination-prefix name**

t2 extensive limit 3

Service Accounting interface: sp-2/0/0, Local interface index: 542

Service name: t2

Source address: 10.1.1.2, Source prefix length: 20

Destination address: 10.200.176.1, Destination prefix length: 0

Input SNMP interface index: 24, Output SNMP interface index: 26

Source-AS: 69, Destination-AS: 69

Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003

Flow count: 0, Packet count: 6, Byte count: 5340

Source address: 10.1.1.2, Source prefix length: 20

Destination address: 10.243.160.1, Destination prefix length: 0

Input SNMP interface index: 24, Output SNMP interface index: 26

Source-AS: 69, Destination-AS: 69

Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003

Flow count: 0, Packet count: 6, Byte count: 5490

Source address: 10.1.1.2, Source prefix length: 20

Destination address: 10.162.160.1, Destination prefix length: 0

Input SNMP interface index: 24, Output SNMP interface index: 26

Source-AS: 69, Destination-AS: 69

Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003

Flow count: 0, Packet count: 6, Byte count: 4079

Monitoring Traffic Using Passive Flow Monitoring

IN THIS CHAPTER

- [Passive Flow Monitoring Overview | 150](#)
- [Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers | 152](#)
- [Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers | 154](#)
- [Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers | 156](#)
- [Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers | 157](#)
- [Configuring Passive Flow Monitoring | 166](#)
- [Example: Passive Flow Monitoring Configuration on M, MX and T Series Routers | 168](#)
- [Configuring a Routing Table Group on an M, MX or T Series Router to Add Interface Routes into the Forwarding Instance | 187](#)
- [Using IPSec and an ES PIC on an M, MX or T Series Router to Send Encrypted Traffic to a Packet Analyzer | 188](#)
- [Applying a Firewall Filter Output Interface on an M, MX or T Series Router to Port-mirror Traffic to PICs or Flow Collection Services | 189](#)
- [Monitoring Traffic on a Router with a VRF Instance and a Monitoring Group | 190](#)
- [Specifying a Firewall Filter on an M, MX or T Series Router to Select Traffic to Monitor | 191](#)
- [Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers | 192](#)
- [Establishing a VRF Instance on an M, MX or T Series Router for Monitored Traffic | 196](#)
- [Configuring a Monitoring Group on an M, MX or T Series Router to Send Traffic to the Flow Server | 196](#)
- [Configuring Policy Options on M, MX or T Series Routers | 198](#)
- [Stripping MPLS Labels on ATM, Ethernet-Based and SONET/SDH Router Interfaces | 199](#)
- [Using an M, MX or T Series Router Flow Collector Interface to Process and Export Multiple Flow Records | 201](#)
- [Example: Configuring a Flow Collector Interface on an M, MX or T Series Router | 208](#)

Passive Flow Monitoring Overview

Using a Juniper Networks M Series, T Series, or MX Series router, a selection of PICs (including the Monitoring Services PIC, Adaptive Services [AS] PIC, Multiservices PIC, or Multiservices DPC) and other networking hardware, you can monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to do the following:


- Gather and export detailed information about IP version 4 (IPv4) traffic flows between source and destination nodes in your network.
- Sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format.
- Perform discard accounting on an incoming traffic flow.
- Encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both.
- Direct filtered traffic to different packet analyzers and present the data in its original format (port mirror).

NOTE: Monitoring Services PICs, AS PICs, and Multiservices PICs must be mounted on an Enhanced Flexible PIC Concentrator (FPC) in an M Series, T Series, or MX Series router. Multiservices DPCs installed in Juniper Networks MX Series routers support the same functionality, with the exception of the passive monitoring and flow-tap features.

The router used for passive monitoring does not route packets from the monitored interface, nor does it run any routing protocols related to those interfaces; it only receives traffic flows, collects intercepted

traffic, and exports it to cflowd servers and packet analyzers. Figure 20 on page 151 shows a typical topology for the passive flow-monitoring application.

Figure 20: Passive Monitoring Application Topology



Traffic travels normally between Router 1 and Router 2. To redirect IPv4 traffic, you insert an optical splitter on the interface between these two routers. The optical splitter copies and redirects the traffic to the monitoring station, which is an M40e, M160, M320, T Series, or MX Series router. The optical cable connects only the receive port on the monitoring station, never the transmit port. This configuration allows the monitoring station to receive traffic from the router being monitored but never to transmit it back.

If you are monitoring traffic flow, the Internet Processor II application-specific integrated circuit (ASIC) in the router forwards a copy of the traffic to the Monitoring Services, Adaptive Services, or Multiservices PIC in the monitoring station. If more than one monitoring PIC is installed, the monitoring station distributes the load of the incoming traffic across the multiple PICs. The monitoring PICs generate flow records in cflowd version 5 format, and the records are then exported to the cflowd collector.

If you are performing lawful interception of traffic between the two routers, the Internet Processor II ASIC filters the incoming traffic and forwards it to the Tunnel Services PIC. Filter-based forwarding is then applied to direct the traffic to the packet analyzers.

Optionally, the intercepted traffic or the cflowd records can be encrypted by the ES PIC or IP Security (IPsec) services and then sent to a cflowd server or packet analyzer.

RELATED DOCUMENTATION

[Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers](#) | 157

Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers

To perform passive flow monitoring, your router must meet these minimum requirements:

- Junos OS Release 22.4R1 or later for passive flow monitoring support on the MX304 router with the LMIC16-BASE line card, on the MX10004, MX10008, and MX10016 routers with the LC9600 line card and on the MX2010 and MX2020 routers with the MPC10 and MPC11 line cards.
- Junos OS Release 20.4R1 or later for passive flow monitoring support on the MX10008 router with the JNP10K-2101 line card and on the MX240/MX480/MX960/MX2008/MX2010/MX2020 routers with either the MPC7E-MRATE or MPC7E-10G line card.
- Junos OS Release 9.2 or later for passive flow monitoring support for IQ2 interfaces only on M120, M320, T320, T640, T1600 and MX Series routers
- Junos OS Release 8.5 or later for passive flow monitoring support on the MX Series MultiServices routers
- Junos OS Release 8.4 or later for passive flow monitoring support on the MultiServices 400 PIC (Type 2)
- Junos OS Release 7.6 or later to clear error and flow statistics with the `clear passive-monitoring statistics` command
- Junos OS Release 7.5 or later for support of the dynamic flow capture (DFC) Management Information Base (MIB)
- Junos OS Release 7.4 or later for dynamic flow capture on Monitoring Services III PICs installed in T Series and M320 routers, and port mirroring of IPv6 packets
- Junos OS Release 7.3 or later for passive flow monitoring on selected Ethernet-based interfaces and filter-based forwarding on output interfaces
- Junos OS Release 7.1 or later for passive flow monitoring and flow collection services on Monitoring Services II PICs installed in T Series and M320 routers
- Junos OS Release 6.4 or later for support of the next-hop IP address field in flow monitoring version 5 records

- Junos OS Release 6.2 or later for ATM2 intelligent queuing (IQ) interface passive monitoring, flow collection services, and MPLS label stripping
- Junos OS Release 6.1 or later for MPLS passive monitoring
- Junos OS Release 6.0 or later for the Monitoring Services II PIC
- Junos OS Release 5.7 or later for the automatic insertion of autonomous system (AS) numbers and SNMP index values for interfaces into flow records
- Junos OS Release 5.4 or later for the Monitoring Services PIC
- M40e, M160, M320, MX Series, or T Series router with an Internet Processor II ASIC or later
- Type 1 enhanced FPCs
- Two optical splitters
- A Tunnel Services PIC (required if you wish to send traffic to more than one analyzer)
- An input interface from the following list:
 - SONET/SDH PIC—OC3, OC12, or OC48
 - ATM2 IQ PIC—OC3 or OC12
 - 4-port Fast Ethernet PIC
 - Gigabit Ethernet PIC—4-port with small form-factor pluggable transceiver (SFP) or 10-port with SFP
 - 1-port 10-Gigabit Ethernet PIC with XENPAK
- Outgoing PICs to connect to the flow collector or packet analyzer
- Flow monitoring version 5 collector
- ES PIC and packet analyzers (optional)

RELATED DOCUMENTATION

[Active Flow Monitoring System Requirements](#) | 45

[Active Flow Monitoring PIC Specifications](#) | 49

Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers

There are several hardware and software considerations when you implement passive flow monitoring. When defining the hardware requirements of the monitoring station, keep in mind the following:

- The input interfaces on the monitoring station must be SONET/SDH interfaces (OC3, OC12, or OC48), ATM2 IQ interfaces (OC3 or OC12), 4-port Fast Ethernet interfaces, Gigabit Ethernet interfaces with SFP (4-port or 10-port), or 1-port 10-Gigabit Ethernet interfaces with XENPAK.
- To monitor the flows in both directions for a single interface, the monitoring station must have two SONET/SDH, ATM2 IQ, or Ethernet-based receive ports, one for each direction of flow. In ["Passive Flow Monitoring Application Topology" on page 156](#), the monitoring station needs one port to monitor the traffic flowing from Router 1 to Router 2, and a second port to monitor the traffic flowing from Router 2 to Router 1.
- The Monitoring Services PICs must be installed in a Type 1 enhanced FPC slot.
- Type 1 and Type 2 Tunnel Services PICs are supported.
- Use an ES PIC to encrypt the flow export.
- Symmetric hashing is not supported on the MPC10 and MPC11 line cards. You should choose a different MPC line card if you wish to support symmetrical hashing along with passive monitoring.
- You can only configure passive monitoring on a physical port and not on a logical interface or per VLAN. You cannot configure passive monitoring on an aggregated Ethernet port or on a port with Ethernet encapsulation.
- IDS servers must be directly connected to the router. You need to configure the interfaces connecting to the IDS servers as part of a link aggregation group (LAG). You need to configure static routes to route the packets onto an IDS server.

When defining a traffic monitoring strategy, keep in mind the following:

- The monitoring station collects only IPv4 packets. All other packet formats are discarded and not counted.
- You can set the amount of time a data flow can be inactive before the monitoring station terminates the flow and exports the flow data. To set the timer, include the `flow-inactive-timeout` statement at the `[edit forwarding-options monitoring group-name family inet output]` hierarchy level. The timer value can be from 15 seconds through 1800 seconds, with a default value of 60 seconds.

You can also configure the monitoring station to collect periodic flow reports for flows that last longer than the configured active timeout. To set this activity timer, include the `flow-active-timeout` statement at

the [edit forwarding-options monitoring *group-name* family inet output] hierarchy level. The timer value can be from 60 seconds through 1800 seconds, with a default value of 180 seconds.

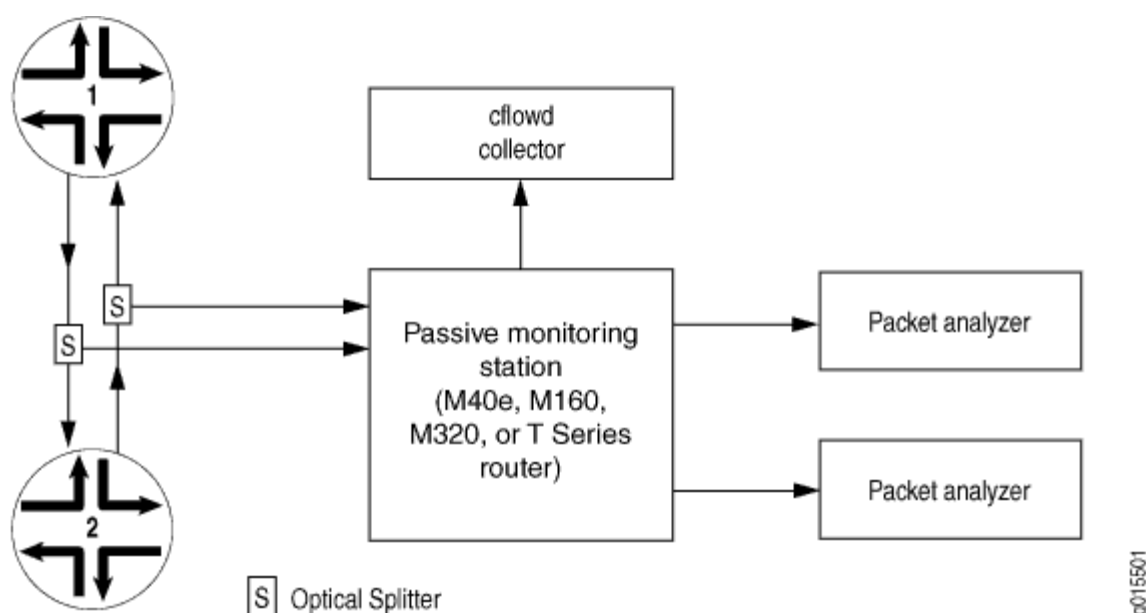
- Multiple expired flows are exported together, if possible. A UDP packet is sent when one of the following conditions is met:
 - When 30 flows are contained in the current packet, the flows are exported.
 - If there are fewer than 30 flows but the export timer expires, the flows are exported one second after the timer expires.
- TCP and UDP flows are considered differently:
 - TCP flows watch for a segment containing the **FIN** bit and a subsequent acknowledgement (**ACK**) to detect the end of a flow. Alternately, a TCP reset (**RST**) can also indicate the end of a flow. When these TCP combinations are detected, the flow expires. The **FIN+ACK** and **RST** cases cover most TCP stream closures. For all other flows, an inactive timeout is needed.
 - All non-TCP flows, such as UDP, depend on timeout mechanisms for export.
- The default MTU value for SONET/SDH interfaces is 4474 bytes; for Gigabit Ethernet and Fast Ethernet interfaces, it is 1500 bytes. If the monitoring station receives packets exceeding 4474 bytes, they are discarded; no fragmentation is performed. Note that the supported MTU size on the Gigabit Ethernet or Fast Ethernet PICs might exceed 1500 bytes, depending on the type of PIC.
- Any incoming traffic that is discarded is not forwarded to packet analyzers.
- The interfaces on the monitoring station that collect intercepted traffic must be configured with Cisco HDLC or PPP encapsulation.
- You must always use a standard interface (for example, one that follows the usual *interface-name-fpc/pic/slot* format) to send flow records to a flow server. Flow data generated by the Monitoring Services or Monitoring Services II PICs will not be delivered to the server across the **fxp0** interface.
- You can send version 5 records to multiple flow servers. You can configure up to eight servers and flow traffic is load-balanced between the servers in a round-robin fashion. If one of the servers ceases operation, flow traffic load-balances automatically between the remaining active servers. To configure, include up to eight flow-server statements at the [edit forwarding-options monitoring *group-name* output] hierarchy level.

Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers

Flow monitoring version 5 supports passive flow monitoring. Versions 8 and 9 do not support passive flow monitoring.

The M40e, M160, M320, MX Series, or T Series router that is used for passive flow monitoring does not route packets from monitored interfaces, nor does it run any routing protocols related to those interfaces; it only passes along intercepted traffic and receives traffic flows. [Figure 21 on page 156](#) shows a typical topology for the passive flow monitoring application.

Figure 21: Passive Flow Monitoring Application Topology



Traffic travels normally between Router 1 and Router 2. To redirect IPv4 traffic, you insert an optical splitter on the interface between these two routers. The optical splitter copies and redirects the traffic to the monitoring station. The optical cable connects only the receive port on the monitoring station, never the transmit port. This configuration allows the monitoring station to receive traffic only from the router being monitored but never to transmit it back.

If you are monitoring traffic flow, the Internet Processor II ASIC in the router forwards a copy of the traffic to the Monitoring Services or Monitoring Services II PIC in the monitoring station. If there is more than one Monitoring Services PIC installed, the monitoring station distributes the load of the incoming traffic across the multiple PICs. The Monitoring Services PICs generate flow records in version 5 format, and the records are exported to the flow collector.

When you are performing lawful interception of packets, the Internet Processor II ASIC filters the incoming traffic and forwards it to the Tunnel Services PIC. Filter-based forwarding is then applied to direct the traffic to the packet analyzers. Optionally, the intercepted traffic or the flow records can be encrypted by the ES PIC and then sent to their destination. With additional configuration, flow records can be processed by a flow collector and flows can be captured dynamically.

With MPLS passive monitoring, the router can process MPLS packets with label values that do not have corresponding entries in the `mpls.0` routing table. You can divert these unrecognized MPLS packets, remove the MPLS labels, and redirect the underlying IPv4 packets. This is equivalent to a default route for MPLS packets or a promiscuous label. Because this application does not use a Monitoring Services PIC, see the *Junos MPLS Applications Configuration Guide* for more information about MPLS passive monitoring.

RELATED DOCUMENTATION

[Active Flow Monitoring Overview | 54](#)

[Active Flow Monitoring Overview | 53](#)

Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers

IN THIS SECTION

- [Passive Flow Monitoring for MPLS Encapsulated Packets | 160](#)
- [Example: Enabling IPv4 Passive Flow Monitoring | 162](#)
- [Example: Enabling IPv6 Passive Flow Monitoring | 165](#)

You can monitor IPv4 traffic from another router if you have the following components installed in an M Series, MX Series, or T Series router:

- Monitoring Services, Adaptive Services, or Multiservices PICs to perform the service processing
- SONET/SDH, Fast Ethernet, or Gigabit Ethernet PICs as transit interface

On SONET/SDH interfaces, you enable passive flow monitoring by including the `passive-monitor-mode` statement at the `[edit interfaces so-fpc/pic/port unit logical-unit-number]` hierarchy level:

```
[edit interfaces so-fpc/pic/port unit logical-unit-number]
passive-monitor-mode;
```

On Asynchronous Transfer Mode (ATM), Fast Ethernet, or Gigabit Ethernet interfaces, you enable passive flow monitoring by including the `passive-monitor-mode` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
passive-monitor-mode;
```

IPv6 passive monitoring is not supported on Monitoring Services PICs. You must configure port mirroring to forward the packets from the passive monitored ports to other interfaces. Interfaces configured on the following FPCs and PIC support IPv6 passive monitoring on the T640 and T1600 Series routers:

- Enhanced Scaling FPC2
- Enhanced Scaling FPC3
- Enhanced II FPC1
- Enhanced II FPC2
- Enhanced II FPC3
- Enhanced Scaling FPC4
- Enhanced Scaling FPC4.1
- 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (supported on both WAN-PHY and LAN-PHY mode for both IPv4 and IPv6 addresses)
- Gigabit Ethernet PIC with SFP
- 10-Gigabit Ethernet PIC with XENPAK (T1600 Series router)
- SONET/SDH OC192/STM64 PIC (T1600 Series router)
- SONET/SDH OC192/STM64 PICs with XFP (T1600 Series router)
- SONET/SDH OC48c/STM16 PIC with SFP (T1600 Series router)
- SONET/SDH OC48/STM16 (Multi-Rate)

- SONET/SDH OC12/STM4 (Multi-Rate) PIC with SFP
- Type 1 SONET/SDH OC3/STM1 (Multi-Rate) PIC with SFP

To configure port mirroring, include the `port-mirroring` statement at the `[edit forwarding-options]` hierarchy level.

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine silently drops packets coming from that interface and destined to the router itself. Passive monitoring mode also stops the Routing Engine from transmitting any packet from that interface. Packets received from the monitored interface can be forwarded to monitoring interfaces. If you include the `passive-monitor-mode` statement in the configuration:

- The ATM interface is always up, and the interface does not receive or transmit incoming control packets, such as Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) cells.
- The SONET/SDH interface does not send keepalives or alarms and does not participate actively on the network.
- Gigabit and Fast Ethernet interfaces can support both per-port passive monitoring and per-VLAN passive monitoring. The destination MAC filter on the receive port of the Ethernet interfaces is disabled.
- Ethernet encapsulation options are not allowed.
- Ethernet interfaces do not support the `stacked-vlan-tagging` statement for both IPv4 and IPv6 packets in passive monitoring mode.

On monitoring services interfaces, you enable passive flow monitoring by including the `family` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level, specifying the `inet` option:

```
[edit interfaces interface-name unit logical-unit-number]
family inet;
```

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see ["Configuring Flow-Monitoring Interfaces" on page 5](#).

For conformity with the cflowd record structure, you must include the `receive-options-packets` and `receive-ttl-exceeded` statements at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number family inet]
receive-options-packets;
receive-ttl-exceeded;
```

Passive Flow Monitoring for MPLS Encapsulated Packets

On monitoring services interfaces, you can process MPLS packets that have not been assigned label values and have no corresponding entry in the `mpls.0` routing table. This allows you to assign a default route to unlabeled MPLS packets.

To configure a default label value for MPLS packets, include the `default-route` statement at the `[edit protocols mpls interface interface-name label-map]` hierarchy level:

```
[edit protocols mpls interface interface-name label-map]
default-route {
    (next-hop (address | interface-name | address/interface-name)) | (reject | discard);
    (pop | (swap <out-label>));
    class-of-service value;
    preference preference;
    type type;
}
```

For more information about static labels, see the [MPLS Applications User Guide](#).

Removing MPLS Labels from Incoming Packets

The Junos OS can forward only IPv4 packets to a Monitoring Services, Adaptive Services, or Multiservices PIC. IPv4 and IPv6 packets with MPLS labels cannot be forwarded to a monitoring PIC. By default, if packets with MPLS labels are forwarded to the monitoring PIC, they are discarded. To monitor IPv4 and IPv6 packets with MPLS labels, you must remove the MPLS labels as the packets arrive on the interface.

You can remove MPLS labels from an incoming packet by including the `pop-all-labels` statement at the `[edit interfaces interface-name (atm-options | fastether-options | gigether-options | sonet-options) mpls]` hierarchy level:

```
[edit interfaces interface-name (atm-options | fastether-options | gigether-options | sonet-
options) mpls]
pop-all-labels {
    required-depth [ numbers ];
}
```

For MX Series routers with MPCs, the `pop-all-labels` statement pops all labels by default and the `required-depth` statement is ignored.

For other configurations, you can remove up to two MPLS labels from an incoming packet. By default, the `pop-all-labels` statement takes effect for incoming packets with one or two labels. You can specify the number of MPLS labels that an incoming packet must have for the `pop-all-labels` statement to take effect by including the `required-depth` statement at the `[edit interfaces interface-name (atm-options | fastether-options | gigether-options | sonet-options) mpls pop-all-labels]` hierarchy level:

```
[edit interfaces interface-name (atm-options | fastether-options | gigether-options | sonet-
options) mpls pop-all-labels]
required-depth [ numbers ];
```

The required depth can be 1, 2, or `[1 2]`. If you include the `required-depth 1` statement, the `pop-all-labels` statement takes effect for incoming packets with one label only. If you include the `required-depth 2` statement, the `pop-all-labels` statement takes effect for incoming packets with two labels only. If you include the `required-depth [1 2]` statement, the `pop-all-labels` statement takes effect for incoming packets with one or two labels. A required depth of `[1 2]` is equivalent to the default behavior of the `pop-all-labels` statement.

When you remove MPLS labels from incoming packets, note the following:

- The `pop-all-labels` statement has no effect on IP packets with three or more MPLS labels except for MX Series routers with MPCs.
- When you enable MPLS label removal, you must configure all ports on a PIC with the same label popping mode and required depth.
- You use the `pop-all-labels` statement to enable passive monitoring applications, not active monitoring applications.
- You cannot apply MPLS filters or accounting to the MPLS labels because the labels are removed as soon as the packet arrives on the interface.

- On ATM2 interfaces, you must use a label value greater than 4095 because the lower range of MPLS labels is reserved for label-switched interface (LSI) and virtual private LAN service (VPLS) support. For more information, see the [Junos OS VPNs Library for Routing Devices](#).
- The following ATM encapsulation types are not supported on interfaces with MPLS label removal:
 - atm-ccc-cell-relay
 - atm-ccc-vc-mux
 - atm-mlppp-llc
 - atm-tcc-snap
 - atm-tcc-vc-mux
 - ether-over-atm-llc
 - ether-vpls-over-atm-llc

Example: Enabling IPv4 Passive Flow Monitoring

The following example shows a complete configuration for enabling passive flow monitoring on an Ethernet interface.

In this example, the Gigabit Ethernet interface can accept all Ethernet packets. It strips VLAN tags (if there are any) and up to two MPLS labels blindly, and passes IPv4 packets to the monitoring interface. With this configuration, it can monitor IPv4, VLAN+IPv4, VLAN+MPLS+IPv4, and VLAN+MPLS+MPLS+IPv4 labeled packets.

The Fast Ethernet interface can accept only packets with VLAN ID 100. All other packets are dropped. With this configuration, it can monitor VLAN (ID=100)+IPv4, VLAN (ID=100)+MPLS+IPv4, and VLAN (ID=100)+MPLS+MPLS+IPv4 labeled packets.

```
[edit firewall]
family inet {
    filter input-monitoring-filter {
        term def {
            then {
                count counter;
                accept;
            }
        }
    }
}
```

```

[edit interfaces]
ge-0/0/0 {
  passive-monitor-mode;
  together-options {
    mpls {
      pop-all-labels;
    }
  }
  unit 0 {
    family inet {
      filter {
        input input-monitoring-filter;
      }
    }
  }
}
fe-0/1/0 {
  passive-monitor-mode;
  vlan-tagging;
  fastether-options {
    mpls {
      pop-all-labels required-depth [ 1 2 ];
    }
  }
  unit 0 {
    vlan-id 100;
    family inet {
      filter {
        input input-monitoring-filter;
      }
    }
  }
}
mo-1/0/0 {
  unit 0 {
    family inet {
      receive-options-packets;
      receive-ttl-exceeded;
    }
  }
  unit 1 {
    family inet;
  }
}

```

```

}
[edit forwarding-options]
monitoring mon1 {
  family inet {
    output {
      export-format cflowd-version-5;
      cflowd 192.0.2.2 port 2055;
      interface mo-1/0/0.0 {
        source-address 192.0.2.1;
      }
    }
  }
}
[edit routing-instances]
monitoring-vrf {
  instance-type vrf;
  interface ge-0/0/0.0;
  interface fe-0/1/0.0;
  interface mo-1/0/0.1;
  route-distinguisher 68:1;
  vrf-import monitoring-vrf-import;
  vrf-export monitoring-vrf-export;
  routing-options {
    static {
      route 0.0.0.0/0 next-hop mo-1/0/0.1;
    }
  }
}
[edit policy-options]
policy-statement monitoring-vrf-import {
  then {
    reject;
  }
}
policy-statement monitoring-vrf-export {
  then {
    reject;
  }
}

```

Example: Enabling IPv6 Passive Flow Monitoring

The following example shows a complete configuration for enabling IPv6 passive flow monitoring on an Ethernet interface.

In this example, the Gigabit Ethernet interface can accept all Ethernet packets. It strips VLAN tags (if there are any) and up to two MPLS labels blindly, and passes IPv6 packets to the monitoring interface. With this configuration, the Gigabit Ethernet interface can monitor IPv6, VLAN+IPv6, VLAN+MPLS+IPv6, and VLAN+MPLS+MPLS+IPv6 labeled packets.

The vlan-tagged Gigabit Ethernet interface can accept only packets with VLAN ID 100. All other packets are dropped. With this configuration, it can monitor VLAN (ID=100)+IPv6, VLAN (ID=100)+MPLS+IPv6, and VLAN (ID=100)+MPLS+MPLS+IPv6 labeled packets.

```
[edit interfaces]
xe-0/1/0 {
  passive-monitor-mode;
  unit 0 {
    family inet6 {
      filter {
        input port-mirror6;
      }
      address 2001:db8::1/128;
    }
  }
}

xe-0/1/2 {
  passive-monitor-mode;
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet6 {
      filter {
        input port-mirror6;
      }
    }
  }
}

xe-0/1/1 {
  unit 0 {
    family inet6 {
      address 2001:db8::1/128;
    }
  }
}
```



```

    }

    }
[edit firewall]
family inet6 {
    filter port-mirror6 {
        term term2 {
            then {
                count count_pm;
                port-mirror;
                accept;
            }
        }
    }
}
[edit forwarding options]
port-mirroring {
    input {
        rate 1;
    }
    family inet6 {
        output {
            interface xe-0/1/1.0 {
                next-hop 2001:db8::3;
            }
            no-filter-check;
        }
    }
}

```

RELATED DOCUMENTATION

[Passive Flow Monitoring Overview](#) | 150

Configuring Passive Flow Monitoring

[Table 26 on page 167](#) shows which Juniper Networks PICs and routers support passive flow monitoring. The PICs receive passively monitored network traffic from an input interface (SONET/SDH, ATM2 IQ,

Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet), convert the received packets into flow records, and export them to a flow server for further analysis.

Table 26: Passive Flow Monitoring PIC Support

PIC Type	M40e	M160	T Series/ M320
Monitoring Services PIC	Yes	Yes	No
Monitoring Services II PIC	Yes	Yes	Yes
Monitoring Services III PIC	Yes	Yes	Yes
MultiServices 400 PIC (Type 2)	Yes	No	Yes

The key configuration hierarchy statement for passive flow monitoring is the `monitoring` statement found at the `[edit forwarding-options]` hierarchy level. At minimum, you must configure a VRF routing instance to direct the traffic to a monitoring services interface for flow processing.

However, there are several options you can use that add complexity to passive flow monitoring. For example, you can configure the router to direct traffic into a routing instance and deliver the traffic into a monitoring group. You can also use *port mirroring* and filter-based forwarding to copy and redirect traffic. Optionally, you can configure the monitoring station to encrypt flow output before it is sent to a flow server for processing, to send flow records to a flow collector, or to process on-demand monitoring requests with dynamic flow capture.

RELATED DOCUMENTATION

[Copying and Redirecting Traffic with Port Mirroring and Filter-Based Forwarding](#)

[Using an M, MX or T Series Router Flow Collector Interface to Process and Export Multiple Flow Records | 201](#)


[Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers | 154](#)

Example: Passive Flow Monitoring Configuration on M, MX and T Series Routers

IN THIS SECTION

- Verifying Your Work | 177

Figure 22: Passive Flow Monitoring—Topology Diagram



In [Figure 22 on page 168](#), traffic enters the monitoring station through interfaces **so-0/0/0** and **so-0/1/0**. After the firewall filter accepts the traffic to be monitored, the packets enter a VRF instance.

The original packets travel within the VRF instance to the Monitoring Services PIC for flow processing. The final flow packets are sent from the monitoring services interfaces out the **fe-3/0/0** interface to a flow server.

A copy of the accepted traffic is port-mirrored to the Tunnel PIC. As the copied packets enter the tunnel interface, a second firewall filter separates TCP and UDP packets and places them into two filter-based forwarding instances. The UDP instance directs the UDP packets to a packet analyzer attached to

fe-3/2/0. The TCP instance sends the TCP packets to the ES PIC for encryption and the ES PIC sends the packets to a second packet analyzer connected to **fe-3/2/1**.

Your first step is to define a firewall filter to select packets for monitoring. All filtered traffic must be accepted, and the port-mirror statement at the [edit firewall family inet filter *filter-name* term *term-name* then] hierarchy level facilitates port mirroring.

Next, configure the input SONET/SDH interfaces and apply the firewall filter that you just defined. The `passive-monitor-mode` statement disables SONET keepalives on the SONET/SDH interfaces and enables passive flow monitoring.

Configure all other interfaces that you will use with the monitoring application, including the monitoring services interfaces, the export interfaces, the tunnel interface, and the ES interface. Once the interfaces are in place, configure a VRF instance and monitoring group to direct the original packets from the input interfaces to the monitoring services interfaces for processing. The resulting flow description packets exit **fe-3/0/0** to reach the flow server.

Next, configure statements to port-mirror the monitored traffic to a tunnel interface. Design a firewall filter that selects some of this copied traffic for further analysis and some of the traffic for discarding. In this case, isolate TCP and UDP traffic and direct these two flows into separate filter-based forwarding routing instances. Remember to apply the filter to the tunnel interface to enable the separation of TCP traffic from UDP traffic. Also, import the interface routes into the forwarding instances with a routing table group.

In the filter-based forwarding instances, define static route next hops. The next hop for the TCP instance is the ES interface and the next hop for the UDP instance is the packet analyzer connected to **fe-3/2/0**. Finally, configure IPsec so that the next hop for the TCP traffic is the second packet analyzer attached to **fe-3/2/1**.

```
[edit]
interfaces {
    so-0/0/0 { # Traffic enters the router on this interface.
        description "input interface";
        encapsulation ppp;
        unit 0 {
            passive-monitor-mode; # Disables SONET keepalives.
            family inet {
                filter {
                    input input-monitoring-filter; # The firewall filter is
                    applied here.
                }
            }
        }
    }
}
```

```

        so-0/1/0 { # Traffic enters the router on this interface.
description " input interface";
encapsulation ppp;
unit 0 {
        passive-monitor-mode; # Disables SONET keepalives.
        family inet {
                filter {
                        input input-monitoring-filter; # The firewall filter
is applied here.
                }
        }
}

        es-3/1/0 { # This is where the TCP traffic enters the ES PIC.
unit 0 {
        tunnel {
                source 10.8.8.1;
                destination 10.8.8.2;
        }
        family inet {
                ipsec-sa sa-esp;
                address 192.0.2.1/32 {
                        destination 192.0.2.2;
                }
        }
}

        fe-3/0/0 { # Flow records exit here and travel to the flow server.
description " export interface to the flow server";
unit 0 {
        family inet;
        address 192.168.245.1/30;
}

        fe-3/2/0 { # This export interface for UDP traffic leads to a packet analyzer.
description " export interface to the packet analyzer";
unit 0 {
        family inet {
                address 10.9.9.1/30;
        }
}

        fe-3/2/1 { # This IPSec tunnel source exports TCP traffic to a packet

```

```

analyzer.
    unit 0 {
        family inet {
            address 10.8.8.1/30;
        }
    }
}

    mo-4/0/0 { # This marks the beginning of the monitoring services interfaces.
    unit 0 { # Unit 0 is part of the inet.0 routing table and generates flow records.
        family inet;
    }

        unit 1 { # Unit 1 receives monitored traffic and is part of the VRF
instance.
        family inet;
    }
}

    mo-4/1/0 {
    unit 0 { # Unit 0 is part of the inet.0 routing table and generates flow records.
        family inet;
    }

        unit 1 { # Unit 1 receives monitored traffic and is part of the VRF
instance.
        family inet;
    }
}

    mo-4/2/0 {
    unit 0 { # Unit 0 is part of the inet.0 routing table and generates flow records.
        family inet;
    }

        unit 1 { # Unit 1 receives monitored traffic and is part of the VRF
instance.
        family inet;
    }
}

    mo-4/3/0 {
    unit 0 { # Unit 0 is part of the inet.0 routing table and generates flow records.
        family inet;
    }

        unit 1 { # Unit 1 receives monitored traffic and is part of the VRF
instance.
        family inet;
    }
}

```

```

        vt-0/2/0 { # The tunnel services interface receives the port-mirrored
traffic.
    unit 0 {
        family inet {
            filter {
                input tunnel-interface-filter; # The filter splits
traffic into TCP and UDP
            }
        }
    }
}
forwarding-options {
    monitoring group1 { # Monitored traffic is processed by the monitoring
services
        family inet { # interfaces and flow records are sent to the flow server.
            output {
                export-format cflowd-version-5;
                flow-active-timeout 60;
                flow-inactive-timeout 30;
                flow-server 192.168.245.2 port 2055; # IP address and port
for server.
            }
            interface mo-4/0/0.1 { # Use monitoring services
interfaces for output.
                engine-id 1; # engine and interface-index statements are optional.
                engine-type 1;
                input-interface-index 44;
                output-interface-index 54;
                source-address 192.168.245.1; # This is the IP address
of fe-3/0/0.
            }
            interface mo-4/1/0.1 {
                engine-id 2; # engine and interface-index statements are optional.
                engine-type 1;
                input-interface-index 45;
                output-interface-index 55;
                source-address 192.168.245.1; # This is the IP address of fe-3/0/0.
            }
            interface mo-4/2/0.1 {
                engine-id 3; # engine and interface-index statements are optional.
                engine-type 1;
                input-interface-index 46;
                output-interface-index 56;

```

```

        source-address 192.168.245.1; # This is the IP address of fe-3/0/0.
    }
    interface mo-4/3/0.1 {
        engine-id 4; # engine and interface-index statements are optional.
        engine-type 1;
        input-interface-index 47;
        output-interface-index 57;
        source-address 192.168.245.1; # This is the IP address
of fe-3/0/0.
    }
}
}
}

    port-mirroring { # Copies the traffic and sends it to the Tunnel Services
PIC.
    family inet {
        input {
            rate 1;
            run-length 1;
        }
        output {
            interface vt-0/2/0.0;
            no-filter-check;
        }
    }
}
}

    routing-options { # This installs the interface routes into the forwarding
instances.
        interface-routes {
            rib-group inet bc-vrf;
        }
        rib-groups {
            bc-vrf {
                import-rib [inet.0 tcp-routing-table.inet.0 udp-routing-table.inet.0];
            }
        }
        forwarding-table {
            export pplb; # Applies per-packet load balancing to the forwarding table.
        }
    }
}
policy-options {
    policy-statement monitoring-vrf-import {

```



```

        then reject;
    }
    policy-statement monitoring-vrf-export {
        then reject;
    }
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}

    security { # This sets IPSec options for the ES PIC.
    ipsec {
        proposal esp-sha1-3des {
            protocol esp;
            authentication-algorithm hmac-sha1-96;
            encryption-algorithm 3des-cbc;
            lifetime-seconds 180;
        }
        policy esp-group2 {
            perfect-forward-secrecy {
                keys group2;
            }
            proposals esp-sha1-3des;
        }
        security-association sa-esp {
            mode tunnel;
            dynamic {
                ipsec-policy esp-group2;
            }
        }
    }
    ike {
        proposal ike-esp {
            authentication-method pre-shared-keys;
            dh-group group2;
            authentication-algorithm sha1;
            encryption-algorithm 3des-cbc;
            lifetime-seconds 180;
        }
        policy 10.8.8.2 {
            mode aggressive;
            proposals ike-esp;

```

```

        pre-shared-key ascii-text "$ABC123";
    }
}
}
firewall {
    family inet {
        filter input-monitoring-filter { # This filter selects traffic to send into the VRF
            term 1 { # instance and prepares the traffic for port mirroring.
                from {
                    destination-address {
                        10.7.0.0/16;
                    }
                }
                then {
                    port-mirror;
                    accept;
                }
            }
            term 2 {
                from {
                    destination-address {
                        10.6.0.0/16;
                    }
                }
                then accept;
            }
        }

        filter tunnel-interface-filter { # This filter breaks the port-
mirrored traffic into two
            term tcp { # filter-based forwarding instances: TCP packets and UDP packets.
                from {
                    protocol tcp;
                }
                then { # This counts TCP packets and sends them into a TCP instance.
                    count tcp;
                    routing-instance tcp-routing-table;
                }
            }
            term udp {
                from {
                    protocol udp;
                }
                then { # This counts UDP packets and sends them into a UDP instance.

```



```

udp-routing-table { # This is the filter-based forwarding instance for UDP traffic.
    instance-type forwarding;
    routing-options { # The next hop is the second packet analyzer.
        static {
            route 0.0.0.0/0 next-hop 10.9.1.2;
        }
    }
}
}
}

```

Verifying Your Work

To verify that your configuration is correct, use the following commands on the monitoring station that is configured for passive flow monitoring:

- `show route 0/0`
- `show passive-monitoring error`
- `show passive-monitoring flow`
- `show passive-monitoring memory`
- `show passive-monitoring status`
- `show passive-monitoring usage`

To clear statistics for the `show passive-monitoring error` and `show passive-monitoring flow` commands, issue the `clear passive-monitoring (all | interface-name)` command.

You can also view passive flow monitoring status with the Simple Network Management Protocol (SNMP). The following Management Information Base (MIB) tables are supported:

- **jnxPMonErrorTable**—Corresponds to the `show passive-monitoring error` command.
- **jnxPMonFlowTable**—Corresponds to the `show passive-monitoring flow` command.
- **jnxPMonMemoryTable**—Corresponds to the `show passive-monitoring memory` command.

The following section shows the output of the `show` commands used with the configuration example:

```

user@host> show route 0/0
<skip inet.0>

```

We are only concerned with the routing-instance route.

```
bc-vrf.inet.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
bc-vrf.inet.0:+ = Active Route, - = Last Active, * = Both
0.0.0.0/0      *[Static/5] 5d 17:34:57
                via mo-4/0/0.1
                > via mo-4/1/0.1
                via mo-4/2/0.1
                via mo-4/3/0.1
tcp-rt.inet.0: 13 destinations, 13 routes (12 active, 0 holddown, 1
hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0      *[Static/5] 19:24:39
                > via es-3/1/0.0
                : <other interface routes>
udp-rt.inet.0: 13 destinations, 13 routes (12 active, 0 holddown, 1
hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0      *[Static/5] 19:24:39
                > to 10.9.1.2 via fe-3/2/0.0
                : <other interface routes>
```

NOTE: For all `show passive-monitoring` commands, the output obtained when using a wildcard (such as `*`) or the `all` option is based on the configured interfaces listed at the `[edit forwarding-options monitoring group-name]` hierarchy level. In the output from the configuration example, you see information only for the configured interfaces `mo-4/0/0`, `mo-4/1/0`, `mo-4/2/0`, and `mo-4/3/0`.

Many of the statements you can configure in a monitoring group, such as `engine-id` and `engine-type`, are visible in the output of the `show passive-monitoring` commands.

Table 27: Output Fields for the `show passive-monitoring error` Command

Field	Explanation
Packets dropped (no memory)	Number of packets dropped because of memory.
Packets dropped (not IP)	Number of non-IP packets dropped.

Table 27: Output Fields for the show passive-monitoring error Command (*Continued*)

Field	Explanation
Packets dropped (not IPv4)	Number of packets dropped because they failed the IPv4 check.
Packets dropped (header too small)	Number of packets dropped because the packet length or IP header length was too small.
Memory allocation failures	Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space.
Memory free failures	Number of flow record memory frees.
Memory free list failures	Number of flow records received from free list that failed. Memory is nearly exhausted or too many new flows greater than 128K are being created in one second.
Memory warning	The flows have exceeded 1 million packets per second (Mpps) on a Monitoring Services PIC or 2 Mpps on a Monitoring Services II PIC. The response can be Yes or No .
Memory overload	The memory has been overloaded. The response is Yes or No .
PPS overload	In packets per second, whether the PIC is receiving more traffic than the configured threshold. The response can be Yes or No .

Table 27: Output Fields for the show passive-monitoring error Command *(Continued)*

Field	Explanation
BPS overload	In bytes per second, whether the PIC is receiving more traffic than the configured threshold. The response can be Yes or No .

```
user@host> show passive-monitoring error all
```

```
Passive monitoring interface: mo-4/0/0, Local interface index: 44
```

```
Error information
```

```
Packets dropped (no memory): 0, Packets dropped (not IP): 0
```

```
Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
```

```
Memory allocation failures: 0, Memory free failures: 0
```

```
Memory free list failures: 0
```

```
Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No
```

```
Passive monitoring interface: mo-4/1/0, Local interface index: 45
```

```
Error information
```

```
Packets dropped (no memory): 0, Packets dropped (not IP): 0
```

```
Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
```

```
Memory allocation failures: 0, Memory free failures: 0
```

```
Memory free list failures: 0
```

```
Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No
```

```
Passive monitoring interface: mo-4/2/0, Local interface index: 46
```

```
Error information
```

```
Packets dropped (no memory): 0, Packets dropped (not IP): 0
```

```
Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
```

```
Memory allocation failures: 0, Memory free failures: 0
```

```
Memory free list failures: 0
```

```
Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No
```

```
Passive monitoring interface: mo-4/3/0, Local interface index: 47
```

```
Error information
```

```
Packets dropped (no memory): 0, Packets dropped (not IP): 0
```

```
Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
```

```
Memory allocation failures: 0, Memory free failures: 0
```

```
Memory free list failures: 0
```

```
Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No
```

Table 28: Output Fields for the show passive-monitoring flow Command

Field	Explanation
Flow packets	Number of packets received by an operational PIC.
Flow bytes	Number of bytes received by an operational PIC.
Flow packets 10-second rate	Number of packets per second handled by the PIC and displayed as a 10-second average.
Flow bytes 10-second rate	Number of bytes per second handled by the PIC and displayed as a 10-second average.
Active flows	Number of currently active flows tracked by the PIC.
Total flows	Total number of flows received by an operational PIC.
Flows exported	Total number of flows exported by an operational PIC.
Flows packets exported	Total number of flow packets exported by an operational PIC.
Flows inactive timed out	Total number of flows that are exported because of inactivity.
Flows active timed out	Total number of long-lived flows that are exported because of an active timeout.

```

user@host> show passive-monitoring flow all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Flow information
Flow packets: 6533434, Flow bytes: 653343400
Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
Active flows: 0, Total flows: 1599
Flows exported: 1599, Flows packets exported: 55
Flows inactive timed out: 1599, Flows active timed out: 0

```


Passive monitoring interface: mo-4/1/0, Local interface index: 45

Flow information

Flow packets: 6537780, Flow bytes: 653778000

Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0

Active flows: 0, Total flows: 1601

Flows exported: 1601, Flows packets exported: 55

Flows inactive timed out: 1601, Flows active timed out: 0

Passive monitoring interface: mo-4/2/0, Local interface index: 46

Flow information

Flow packets: 6529259, Flow bytes: 652925900

Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0

Active flows: 0, Total flows: 1599

Flows exported: 1599, Flows packets exported: 55

Flows inactive timed out: 1599, Flows active timed out: 0

Passive monitoring interface: mo-4/3/0, Local interface index: 47

Flow information

Flow packets: 6560741, Flow bytes: 656074100

Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0

Active flows: 0, Total flows: 1598

Flows exported: 1598, Flows packets exported: 55

Flows inactive timed out: 1598, Flows active timed out: 0

Table 29: Output Fields for the show passive-monitoring memory Command

Field	Explanation
Allocation count	Number of flow records allocated.
Free count	Number of flow records freed.
Maximum allocated	Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time.
Allocations per second	Flow records allocated per second during the last statistics interval on the PIC.

Table 29: Output Fields for the show passive-monitoring memory Command (*Continued*)

Field	Explanation
Frees per second	Flow records freed per second during the last statistics interval on the PIC.
Total memory used	Total amount of memory currently used (in bytes).
Total memory free	Total amount of memory currently free (in bytes).

```
user@host> show passive-monitoring memory all
```

```
Passive monitoring interface: mo-4/0/0, Local interface index: 44
```

```
Memory utilization
```

```
Allocation count: 1600, Free count: 1599, Maximum allocated: 1600
```

```
Allocations per second: 3200, Frees per second: 1438
```

```
Total memory used (in bytes): 103579176, Total memory free (in bytes): 163914184
```

```
Passive monitoring interface: mo-4/1/0, Local interface index: 45
```

```
Memory utilization
```

```
Allocation count: 1602, Free count: 1601, Maximum allocated: 1602
```

```
Allocations per second: 3204, Frees per second: 1472
```

```
Total memory used (in bytes): 103579176, Total memory free (in bytes): 163914184
```

```
Passive monitoring interface: mo-4/2/0, Local interface index: 46
```

```
Memory utilization
```

```
Allocation count: 1600, Free count: 1599, Maximum allocated: 1600
```

```
Allocations per second: 3200, Frees per second: 1440
```

```
Total memory used (in bytes): 103579176, Total memory free (in bytes): 163914184
```

```
Passive monitoring interface: mo-4/3/0, Local interface index: 47
```

```
Memory utilization
```

```
Allocation count: 1599, Free count: 1598, Maximum allocated: 1599
```

```
Allocations per second: 3198, Frees per second: 1468
```

```
Total memory used (in bytes): 103579176, Total memory free (in bytes): 163914184
```

Table 30: Output Fields for the show passive-monitoring status Command

Field	Explanation
Interface state	Indicates whether the interface is monitoring (operating properly), disabled (administratively disabled), or not monitoring (not configured).
Group index	Integer that represents the monitoring group of which the PIC is a member. (This does not indicate the number of monitoring groups.)
Export interval	Configured export interval for flow records, in seconds.
Export format	Configured export format (only v5 is currently supported).
Protocol	Protocol the PIC is configured to monitor (only IPv4 is currently supported).
Engine type	Configured engine type that is inserted in output flow packets.
Engine ID	Configured engine ID that is inserted in output flow packets.
Route record count	Number of routes recorded.
IFL to SNMP index count	Number of logical interfaces mapped to an SNMP index.
AS count	Number of AS boundaries that the flow has crossed.
Time set	Indicates whether the time stamp is in place.
Configuration set	Indicates whether the monitoring configuration is set.
Route record set	Indicates whether routes are being recorded.

Table 30: Output Fields for the show passive-monitoring status Command (*Continued*)

Field	Explanation
IFL SNMP map set	Indicates whether logical interfaces are being mapped to an SNMP index.

```

user@host> show passive-monitoring status all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
  Interface state: Monitoring
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 1
  Route record count: 13, IFL to SNMP index count: 30, AS count: 1
  Time set: Yes, Configuration set: Yes
  Route record set: Yes, IFL SNMP map set: Yes

```

```

Passive monitoring interface: mo-4/1/0, Local interface index: 45
  Interface state: Monitoring
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 2
  Route record count: 13, IFL to SNMP index count: 30, AS count: 1
  Time set: Yes, Configuration set: Yes
  Route record set: Yes, IFL SNMP map set: Yes

```

```

Passive monitoring interface: mo-4/2/0, Local interface index: 46
  Interface state: Monitoring
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 3
  Route record count: 13, IFL to SNMP index count: 30, AS count: 1
  Time set: Yes, Configuration set: Yes
  Route record set: Yes, IFL SNMP map set: Yes

```

```

Passive monitoring interface: mo-4/3/0, Local interface index: 47
  Interface state: Monitoring
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 4
  Route record count: 13, IFL to SNMP index count: 30, AS count: 1

```

Time set: Yes, Configuration set: Yes
Route record set: Yes, IFL SNMP map set: Yes

Table 31: Output Fields for the show passive-monitoring usage Command

Field	Explanation
Uptime	Time, in milliseconds, that the PIC has been operational.
Interrupt time	Cumulative time that the PIC spent in processing packets since the last PIC reset.
Load (5 second)	CPU load on the PIC averaged over 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.
Load (1 minute)	CPU load on the PIC averaged over 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.

```

user@host> show passive-monitoring usage *
Passive monitoring interface: mo-4/0/0, Local interface index: 44
CPU utilization
  Uptime: 653155 milliseconds, Interrupt time: 40213754 microseconds
  Load (5 second): 20%, Load (1 minute): 17%

Passive monitoring interface: mo-4/1/0, Local interface index: 45
CPU utilization
  Uptime: 652292 milliseconds, Interrupt time: 40223178 microseconds
  Load (5 second): 22%, Load (1 minute): 15%

Passive monitoring interface: mo-4/2/0, Local interface index: 46
CPU utilization
  Uptime: 649491 milliseconds, Interrupt time: 40173645 microseconds
  Load (5 second): 22%, Load (1 minute): 10098862%

Passive monitoring interface: mo-4/3/0, Local interface index: 47
CPU utilization

```

```
Uptime: 657328 milliseconds, Interrupt time: 40368704 microseconds
Load (5 second): 1%, Load (1 minute): 15%
```

Configuring a Routing Table Group on an M, MX or T Series Router to Add Interface Routes into the Forwarding Instance

Next, import the interface routes into the forwarding instance. This step is necessary because the next hops specified in the forwarding instances must be installed in the forwarding instances themselves. To configure, include the `import-rib` statement at the `[edit routing-options rib-groups group-name]` hierarchy level. The export statement at the `[edit routing-options forwarding-table]` hierarchy level and the **pplb** policy enable load balancing.

```
[edit]
routing-options {
  interface-routes {
    rib-group inet bc-vrf;
  }
  rib-groups {
    bc-vrf {
      import-rib [inet.0 tcp-routing-table.inet.0 udp-routing-table.inet.0];
    }
  }
  forwarding-table {
    export pplb;
  }
}
policy-options {
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
}
```

Using IPsec and an ES PIC on an M, MX or T Series Router to Send Encrypted Traffic to a Packet Analyzer

You can send some or all of the traffic securely to the packet analyzer using IPsec (a suite of related protocols for cryptographically securing communications at the IP Packet Layer) and an Encryption Services (ES) PIC. In this case, the TCP traffic is encrypted, sent over an IPsec tunnel, and received by the packet analyzer. For more information on configuring IPsec on the ES PIC, see the *IPsec User Guide* or the *Junos System Basics Configuration Guide*.

```
[edit]
interfaces {
  es-3/1/0 {
    unit 0 {
      tunnel {
        source 10.8.8.1;
        destination 10.8.8.2;
      }
      family inet {
        ipsec-sa sa-esp;
        address 192.0.2.1/32 {
          destination 192.0.2.2;
        }
      }
    }
  }
  fe-3/2/1 {
    unit 0 {
      family inet {
        address 10.8.8.1/30;
      }
    }
  }
}
security {
  ipsec {
    proposal esp-sha1-3des {
      protocol esp;
      authentication-algorithm hmac-sha1-96;
      encryption-algorithm 3des-cbc;
      lifetime-seconds 180;
    }
  }
}
```

```

    policy esp-group2 {
        perfect-forward-secrecy {
            keys group2;
        }
        proposals esp-sha1-3des;
    }
    security-association sa-esp {
        mode tunnel;
        dynamic {
            ipsec-policy esp-group2;
        }
    }
}
ike {
    proposal ike-esp {
        authentication-method pre-shared-keys;
        dh-group group2;
        authentication-algorithm sha1;
        encryption-algorithm 3des-cbc;
        lifetime-seconds 180;
    }
    policy 10.8.8.2 {
        mode aggressive;
        proposals ike-esp;
        pre-shared-key ascii-text "$ABC123";
    }
}
}

```

Applying a Firewall Filter Output Interface on an M, MX or T Series Router to Port-mirror Traffic to PICs or Flow Collection Services

On output interfaces, you can apply a firewall filter that leads to a filter-based forwarding routing instance. This is useful if you want to port-mirror traffic to multiple Monitoring Services PICs or flow collection services interfaces. To configure, include the output statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet filter] hierarchy level.

```

[edit]
interfaces

```



```

fe-3/1/0 {
  description "export interface to flow collection services interfaces";
  unit 0 {
    family inet;
    address ip-address;
    filter {
      output output-filter-name;
    }
  }
}

```

Monitoring Traffic on a Router with a VRF Instance and a Monitoring Group

The first way you can implement passive flow monitoring is to direct traffic into a VRF routing instance and use a monitoring group to export this traffic to a flow server for analysis. Complete the following tasks:

- ["Specifying a Firewall Filter on an M, MX or T Series Router to Select Traffic to Monitor" on page 191](#)
- ["Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers" on page 192](#)
- ["Establishing a VRF Instance on an M, MX or T Series Router for Monitored Traffic" on page 196](#)
- ["Configuring a Monitoring Group on an M, MX or T Series Router to Send Traffic to the Flow Server" on page 196](#)
- ["Configuring Policy Options on M, MX or T Series Routers" on page 198](#)
- ["Stripping MPLS Labels on ATM, Ethernet-Based and SONET/SDH Router Interfaces" on page 199](#)

Specifying a Firewall Filter on an M, MX or T Series Router to Select Traffic to Monitor

When you define a firewall filter, you select the initial traffic to be monitored. To configure a firewall filter, include the filter statement at the `[edit firewall family inet]` hierarchy level. All filtered traffic to be monitored must be accepted.

```
[edit]
firewall {
  family inet {
    filter input-monitoring-filter {
      term 1 {
        from {
          destination-address {
            10.7.0.0/16;
          }
        }
        then {
          count counter1;
          accept;
        }
      }
      term 2 {
        from {
          destination-address {
            10.6.0.0/16;
          }
        }
        then {
          count counter2;
          accept;
        }
      }
    }
  }
}
```

Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers

After creating the input filter, you need to configure the interfaces where traffic will enter the router. To enable passive flow monitoring for SONET/SDH input interfaces, include the `passive-monitor-mode` statement at the `[edit interfaces so-fpc/pic/port unit unit-number]` hierarchy level. This mode disables the router from participating in the network as an active device. On SONET/SDH interfaces, passive monitor mode suppresses SONET keepalives.

For ATM2 IQ interfaces, passive monitor mode suppresses the sending and receiving of ATM Operations, Administration, and Maintenance (OAM) and Integrated Local Management Interface (ILMI) control messages. To enable passive flow monitoring for ATM2 IQ input interfaces, include the `passive-monitor-mode` statement at the `[edit interfaces at-fpc/pic/port]` hierarchy level. ATM passive monitoring supports the following interface encapsulation types: Cisco-compatible ATM Network Layer Protocol ID (NLPID) (`atm-cisco-nlpid`), ATM NLPID (`atm-nlpid`), ATM Point-to-Point Protocol (PPP) over ATM Adaptation Layer 5 (AAL5)/ logical link control (LLC) (`atm-ppp-llc`), ATM PPP over raw AAL5 (`atm-ppp-vc-mux`), ATM LLC/ subnetwork attachment point (SNAP) (`atm-snap`), and ATM virtual circuit (VC) multiplexing (`atm-vc-mux`).

Ethernet-based interfaces support both per-port passive monitoring and per-VLAN passive monitoring. For Fast Ethernet interfaces, include the `passive-monitor-mode` statement at the `[edit interfaces fe-fpc/pic/port]` hierarchy level. For Gigabit Ethernet interfaces, include the `passive-monitor-mode` statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level. On Ethernet-based interfaces, passive monitor mode disables the Routing Engine from receiving packets and prevents the routing table from transmitting packets. You can verify this by the presence of the **No-receive** and **No-transmit** interface flags in the output of the `show interfaces (fe | ge)-fpc/pic/port` command.

NOTE: The following restrictions apply to passive flow monitoring on Ethernet-based interfaces:

- No special encapsulation types are allowed, so you must configure Ethernet encapsulations only.
- When you configure the `passive-monitor-mode` statement, destination MAC address filters applied to incoming interfaces are disabled by default.
- The `flow-control` statement at the `[edit interfaces ge-fpc/pic/port gigether-options]` or `[edit interfaces fe-fpc/pic/port fastether-options]` hierarchy level does not work when passive flow monitoring is enabled.

In addition to passive monitor mode, apply the previously defined firewall filter to the interface with the filter statement at the [edit interfaces *interface-name-fpc/pic/port* unit *unit-number* family inet] hierarchy level:

```
[edit]
interfaces {
  so-0/0/0 {
    description "SONET/SDH input interface";
    encapsulation ppp;
    unit 0 {
      passive-monitor-mode;
      family inet {
        filter {
          input input-monitoring-filter;
        }
      }
    }
  }
  at-1/0/0 {
    description "ATM2 IQ input interface";
    passive-monitor-mode;
    atm-options {
      pic-type atm2;
      vpi 0 {
        maximum-vcs 255;
      }
    }
    unit 0 {
      encapsulation atm-snap;
      vci 0.100;
      family inet {
        filter {
          input input-monitoring-filter;
        }
      }
    }
  }
  ge-2/0/0 {
    description "Gigabit Ethernet input interface";
    passive-monitor-mode;
    unit 0 {
      family inet {
```

```

        filter {
            input input-monitoring-filter;
        }
    }
}
}
}
}

```

Configure the interfaces on the Monitoring Services PIC or Monitoring Services II PIC with the `family inet` statement at the `[edit interfaces mo-fpc/pic/port unit unit-number]` hierarchy level. The statement allows the interfaces to process IPv4 traffic received from the input interfaces.

When you use VRF instances, you need to configure two logical interfaces. The first (**unit 0**) is part of the `inet.0` routing table and sources the flow packets. The second (**unit 1**) is configured as part of the VRF instance so the monitoring services interface can serve as a valid next hop for packets received in the instance.

You can also capture options packets and time-to-live (TTL) exceeded information when the monitoring services interface processes flow records. To configure, include the **receive-options-packets** and `receive-ttl-exceeded` statements at the `[edit interfaces mo-fpc/pic/port unit unit-number family inet]` hierarchy level:

```

[edit]
interfaces {
    mo-4/0/0 {
        unit 0 {
            family inet {
                receive-options-packets;
                receive-ttl-exceeded;
            }
        }
        unit 1 {
            family inet;
        }
    }
    mo-4/1/0 {
        unit 0 {
            family inet;
        }
        unit 1 {
            family inet;
        }
    }
    mo-4/2/0 {

```

```

        unit 0 {
            family inet;
        }
        unit 1 {
            family inet;
        }
    }
    mo-4/3/0 {
        unit 0 {
            family inet;
        }
        unit 1 {
            family inet;
        }
    }
}

```

You must also configure the export interface where flow packets exit the monitoring station and are sent to the flow server.

On output interfaces, you can apply a firewall filter that leads to a filter-based forwarding routing instance. This is useful if you want to port-mirror traffic to multiple Monitoring Services PICs or flow collection services interfaces. To configure, include the output statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet filter] hierarchy level. For more information, see [Using Filter-Based Forwarding to Export Monitored Traffic to Multiple Destinations](#).

```

[edit]
interfaces
fe-3/0/0 {
    description "export interface to flow server";
    unit 0 {
        family inet;
        address ip-address;
        filter {
            output output-filter-name;
        }
    }
}

```

Establishing a VRF Instance on an M, MX or T Series Router for Monitored Traffic

After the firewall filter and interfaces are ready, create a VPN routing and forwarding (VRF) instance. The filtered traffic enters the VRF instance and is shared only between the input interfaces and the monitoring services output interfaces. In this case, a group of four monitoring services interfaces is used as the next hop.

```
[edit]
routing-instances {
  monitoring-vrf {
    instance-type vrf;
    interface so-0/0/0.0;
    interface so-0/1/0.0;
    interface mo-4/0/0.1;
    interface mo-4/1/0.1;
    interface mo-4/2/0.1;
    route-distinguisher 69:1;
    vrf-import monitoring-vrf-import;
    vrf-export monitoring-vrf-export;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop [mo-4/0/0.1 mo-4/1/0.1 mo-4/2/0.1];
      }
    }
  }
}
```

Configuring a Monitoring Group on an M, MX or T Series Router to Send Traffic to the Flow Server

You collect flow records by specifying output interfaces in a monitoring group. In general, the monitoring services interfaces are the output interfaces. The logical unit number on the output interfaces when used in conjunction with a VRF instance must be 1. To configure, include the output statement at the [edit forwarding-options monitoring *group-name* family inet] hierarchy level.

NOTE: Because routing instances determine the input interface, the input statement at the [edit forwarding-options monitoring *group-name* family inet] hierarchy level has been removed in Junos OS Release 6.0 and later. If you have a configuration that contains this old statement, we recommend that you update your configuration and remove the statement.

As part of the `mo-fpc/pic/port` statement at the [edit forwarding-options monitoring *group-name* family inet output interface] hierarchy level, you must specify a source address for transmission of flow information. You can use the router ID IP address, the IP address of the input interface, or any local IP address of your choice as the source address. If you provide a different source-address statement for each monitoring services output interface, you can track which interface processes a particular flow record.

All other statements at this level (**engine-id**, **engine-type**, **input-interface-index**, and **output-interface-index**) are dynamically generated, but can be configured manually. To reset outgoing interface or incoming interface indexes that were once configured manually, configure the **input-interface-index** or **outgoing-interface-index** statements with a value of 0 at the [edit forwarding-options monitoring *group-name* family inet output interface *interface-name*] hierarchy level.

To specify the flow server IP address and port number, include the `flow-server ip-address port port-number` statement at the [edit forwarding-options monitoring *group-name* family inet output] hierarchy level. You can specify up to eight flow servers in a monitoring group and the IP address for each server must be unique. Flow records are exported and load-balanced between all active flow servers.

Once you configure the VRF and monitoring group statements, traffic enters the input interfaces, passes to the monitoring services interfaces for processing, and is discarded. The resulting flow description packets exit the monitoring station through the export interface. If you want traffic to travel to destinations other than the monitoring services interfaces, or need to establish additional analysis, see the section [Copying and Redirecting Traffic with Port Mirroring and Filter-Based Forwarding](#).

NOTE: You must complete interface configuration on the Monitoring Services or Monitoring Services II PIC before an interface can be added into a monitoring group. For more information, see ["Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers"](#) on page 192.

```
[edit]
forwarding-options {
  monitoring group1 {
    family inet {
      output {
        export-format cflowd-version-5;
```



```

flow-active-timeout 60;
flow-inactive-timeout 30;
flow-server 192.168.245.1 port 2055;
flow-server 192.168.245.2 port 2055;
interface mo-4/0/0.1 {
    engine-id 1;
    engine-type 1;
    input-interface-index 44;
    output-interface-index 54;
    source-address 192.168.245.1;
}
interface mo-4/1/0.1 {
    engine-id 2;
    engine-type 1;
    input-interface-index 45;
    output-interface-index 55;
    source-address 192.168.245.1;
}
interface mo-4/2/0.1 {
    engine-id 3;
    engine-type 1;
    input-interface-index 46;
    output-interface-index 56;
    source-address 192.168.245.1;
}
}
}
}
}
}
}

```

Configuring Policy Options on M, MX or T Series Routers

When you use a group of next hops in your monitoring group, you can load-balance traffic and distribute it to the export interfaces if you configure policy options. To configure, include the `load-balance per-packet` statement at the `[edit policy-options policy-statement policy-name then]` hierarchy level. You can also reject

import and export of VRF routes by including the `reject` statement at the `[edit policy-options policy-statement policy-name then]` hierarchy level.

```
[edit]
routing-options {
  forwarding-table {
    export pplb;
  }
}
policy-options {
  policy-statement monitoring-vrf-import {
    then {
      reject;
    }
  }
  policy-statement monitoring-vrf-export {
    then {
      reject;
    }
  }
  policy-statement pplb {
    then {
      load-balance per-packet;
    }
  }
}
```

Stripping MPLS Labels on ATM, Ethernet-Based and SONET/SDH Router Interfaces

Because flow monitoring can be performed only on IPv4 packets, any packets containing MPLS labels must have the labels removed before monitoring can occur. To remove MPLS labels from packets as they enter an ATM2 IQ, Ethernet-based, or SONET/SDH interface, include the `pop-all-labels` statement at the `[edit interfaces interface-name-fpc/pic/port (atm | fastether | gigether | sonet)-options mpls]` hierarchy level. If you use static MPLS labels, we recommend you assign label values from **10000** through **99999** to avoid using the label ranges reserved by the Junos OS.

To remove a specified number of labels from selected packets with MPLS labels, include the `required-depth` statement at the `[edit interfaces interface-name-fpc/pic/port (atm | fastether | gigether | sonet)-options`

mpls pop-all-labels] hierarchy level. A **required-depth** value of **1** removes labels from all packets containing only one MPLS label, a value of **2** removes labels from all packets containing only two MPLS labels, and a value of [1 2] removes labels from all packets containing either one or two MPLS labels. The **required-depth** value of [1 2] is the default setting. When you configure the required-depth statement, you must configure the same value for all ports on the same PIC.

The labels are removed and discarded as soon as they arrive at the interface. As a result, no MPLS filters can be applied to the stripped labels, no statistics are generated for the labels, and you cannot apply an IP filter to the incoming packets. No Tunnel Services PIC is required to perform MPLS label stripping.

```
[edit]
interfaces {
  at-fpc/pic/port {
    atm-options {
      mpls {
        pop-all-labels {
          required-depth 1;
        }
      }
    }
  }
  (fe | ge)-fpc/pic/port {
    (fastether | gigether)-options {
      mpls {
        pop-all-labels {
          required-depth [1 2];
        }
      }
    }
  }
  so-fpc/pic/port {
    sonet-options {
      mpls {
        pop-all-labels {
          required-depth 2;
        }
      }
    }
  }
}
```

Using an M, MX or T Series Router Flow Collector Interface to Process and Export Multiple Flow Records

Basic passive monitoring can sometimes create a large number of flow records. However, you can manage multiple flow records with a flow collector interface. You can create a flow collector interface from a Monitoring Services II PIC. The flow collector interface combines multiple flow records received from a monitoring services interface into a compressed ASCII data file and exports the file to an FTP server.

To convert a Monitoring Services II PIC into a flow collector interface, include the `flow-collector` statement at the `[edit chassis fpc fpc-slot pic pic-slot monitoring-services application]` hierarchy level. To restore the monitoring functions of a Monitoring Services II PIC, include the `monitor` statement at the `[edit chassis fpc fpc-slot pic pic-slot monitoring-services application]` hierarchy level.

After you commit the configuration to convert the PIC between the **monitor** and **flow-collector** service types, you must take the PIC offline and then bring the PIC back online. Rebooting the router does not enable the new service type. You can use the Monitoring Services II PIC for either flow collection or monitoring, but not both types of service simultaneously.

A flow collector interface, designated by the **cp-fpc/pic/port** interface name, requires three logical interfaces for correct operation. Units 0 and 1 are used respectively as export channels 0 and 1 to send the compressed ASCII data files to an FTP server. You must include a class-of-service (CoS) configuration for these two export channels to provide adequate bandwidth for file transmission. Unit 2 is used as a flow receive channel to receive flow records from a monitoring services interface.

NOTE: Unlike conventional interfaces, IP addresses for flow collector logical interfaces set up a point-to-point connection between the Routing Engine and the flow collector. The address statement at the `[edit interfaces cp-fpc/pic/port unit unit-number family inet]` hierarchy level corresponds to the IP address of the Routing Engine. Likewise, the destination statement at the `[edit interfaces cp-fpc/pic/port unit unit-number family inet address ip-address]` hierarchy level corresponds to the IP address of the flow collector interface. As a result, you must configure the destination statement for Units 0 and 1 (export channels 0 and 1) with *local* addresses that can reach the FTP server. Similarly, configure the destination statement for Unit 2 (flow receive channel) with a *local* IP address so it can reach the monitoring services interface that sends flow records.

To activate flow collector services after the Monitoring Services II PIC is converted into a flow collector, include the `flow-collector` statement at the `[edit services]` hierarchy level. You also need to configure several additional components:

- Destination of the FTP server—Determines where the compressed ASCII data files are sent after the flow records are collected and processed. To specify the destination FTP server, include the

destinations statement at the [edit services flow-collector] hierarchy level. You can specify up to two FTP server destinations and include the password for each configured server. If two FTP servers are configured, the first server in the configuration is the primary server and the second is a backup server.

- File specifications—Preset data file formats, name formats, and transfer characteristics. Files are sent by FTP to the destination FTP server when the timer expires or when a preset number of records are received, whichever comes first. To set the data file format, include the `data-format` statement at the [edit services flow-collector **file-specification file-name**] hierarchy level. The default data format is **flow-compressed**. To set the export timer and file size thresholds, include the `transfer` statement at the [edit services flow-collector **file-specification file-name**] hierarchy level and specify values for the **timeout** and **record-level** options. The default values are 600 seconds for **timeout** and 500,000 records for **record-level**.

To set the filename format, include the `name-format` statement at the [edit services flow-collector **file-specification file-name**] hierarchy level. Common name format macros that you can use in your configuration are included in [Table 32 on page 202](#).

Table 32: Name Format Macros

Field	Expansion
<code>{am_pm}</code>	AM or PM
<code>{date}</code>	Expands to the current date, using the <code>{month}</code> , <code>{day}</code> , and <code>{year}</code> macros.
<code>{day}</code>	01 to 31
<code>{day_abbrev}</code>	Sun through Sat
<code>{day_full}</code>	Sunday through Saturday
<code>{generation_number}</code>	Expands to a unique, sequential number for each new file created.
<code>{hour_12}</code>	01 to 12
<code>{hour_24}</code>	00 to 23

Table 32: Name Format Macros (*Continued*)

Field	Expansion
<code>{ifalias}</code>	Expands to a description string for the logical interface.
<code>{minute}</code>	00 to 59
<code>{month}</code>	01 to 12
<code>{month_abbrev}</code>	Jan through Dec
<code>{month_full}</code>	January through December
<code>{num_zone}</code>	-2359 to +2359
<code>{second}</code>	00 to 60
<code>{time}</code>	Expands to the time the file is created, using the <code>{hour_24}</code> , <code>{minute}</code> , and <code>{second}</code> macros.
<code>{time_zone}</code>	Time zone code name of the locale (gmt , pst , and so on).
<code>{year}</code>	1970 , 2008 , and so on.
<code>{year_abbrev}</code>	00 to 99

- Input interface-to-flow collector interface mappings—Match an input interface with a flow collector interface and apply the preset file specifications to the input interface. To configure the default flow collector and file specifications for all input interfaces, include the **file-specification** and **collector** statements at the [edit services flow-collector interface-map] hierarchy level. To override the default settings and apply flow collector and file specifications to a specific input interface, include the **file-specification** and **collector** statements at the [edit services flow-collector **interface-map** *interface-name*] hierarchy level.
- Transfer log settings—Allow you to configure the destination FTP server where log files containing the transfer activity history for a flow collector interface are to be archived, the name for the log file,

and the amount of time the router waits before sending the log file to the FTP server. To configure, include the **archive-sites**, **filename-prefix**, and **maximum-age** statements at the [edit services flow-collector transfer-log-archive] hierarchy level. The default value for the **maximum-age** statement is 120 minutes, with a range of 1 to 360 minutes. Also, you can configure up to five FTP archive site servers to receive log files.

- **Miscellaneous settings**—Allow you to configure values for the IP address of the analyzer, an identifier for the analyzer, the maximum number of times the flow collector interface attempts to send transfer log files to the FTP server, and the amount of time the flow collector interface waits between retry attempts. To configure, include the **analyzer-address**, **analyzer-id**, **retry**, and **retry-delay** statements at the [edit services flow-collector] hierarchy level. The range for the **retry** statement is 0 through 10 retry attempts. The default for the **retry-delay** statement is 30 seconds and the range is 0 through 60 seconds.

To specify a flow collector interface as the destination for flow records coming from a Monitoring Services or Monitoring Services II PIC, include the **collector-pic** statement at the [edit forwarding-options monitoring *group-name* family inet output flow-export-destination] hierarchy level. You can select either the flow collector interface or a flow server as the destination for flow records, but you cannot select both destination types simultaneously.

There is also a Juniper Networks enterprise Management Information Base (MIB) for the flow collector interface. The Flow Collector Services MIB allows you to use SNMP to monitor the flow collector interface. The MIB provides statistics on files, records, memory, FTP, and error states of a flow collector interface. It also provides SNMP traps for unavailable destinations, unsuccessful file transfers, flow overloading, and memory overloading. For more information, see the *Junos Network Management Configuration Guide* or view the enterprise-specific Juniper Networks MIBs at <https://www.juniper.net/techpubs/software/junos/mibs.html>.

In summary, to implement the flow collector service, include statements at the [edit chassis], [edit interfaces], [edit forwarding-options], and [edit services] hierarchy levels. The excerpt on the following pages shows the flow collector service configuration hierarchy. For a full configuration example, see ["Example: Configuring a Flow Collector Interface on an M, MX or T Series Router" on page 208](#).

```
[edit]
chassis {
  fpc fpc-slot {
    pic pic-slot {
      monitoring-services {
        application flow-collector;
      }
    }
  }
}
```

```

interfaces {
  cp-fpc/pic/port {
    description "flow_collector_interface";
    unit 0 {
      family inet {
        address ip-address {
          destination ip-address;
        }
      }
    }
    unit 1 {
      family inet {
        address ip-address {
          destination ip-address;
        }
      }
    }
    unit 2 {
      family inet {
        address ip-address {
          destination ip-address;
        }
      }
    }
  }
  interface-fpc/pic/port {
    description "export_interface";
    unit 0 {
      family inet {
        address ip-address;
      }
    }
  }
  mo-fpc/pic/port {
    description "monitoring_services_interface";
    unit 0 {
      family inet;
    }
  }
  SONET/SDH, ATM2 IQ, or Ethernet-based-interface-fpc/pic/port {
    description "input_interface";
    encapsulation encapsulation-type;
    passive-monitor-mode; # Apply to the logical interface for SONET/SDH
  }
}

```



```

    }
}
forwarding-options {
    monitoring group1 {
        family inet {
            output {
                export-format cflowd-version-5;
                flow-active-timeout value;
                flow-inactive-timeout value;
                flow-export-destination collector-pic;
                interface mo-fpc/pic/port {
                    source-address ip-address;
                }
            }
        }
    }
}
services {
    flow-collector {
        analyzer-address ip-address;
        analyzer-id name;
        retry value;
        retry-delay seconds;
        destinations {
            "ftp://username@ftp-server-address-1//directory/" {
                password "encrypted-password";
            }
            "ftp://username@ftp-server-address-2//directory/" {
                password "encrypted-password";
            }
        }
        file-specification {
            file-specification-name {
            }
            data-format flow-compressed;
            transfer timeout value record-level size;
        }
    }
    interface-map {
        file-specification file-specification-name;
        collector cp-fpc/pic/port;
        interface-name {
            file-specification file-specification-name;

```

```
        collector cp-fpc/pic/port;  
    }  
}  
transfer-log-archive {  
    filename-prefix filename;  
    maximum-age timeout-value;  
    archive-sites {  
        "ftp://username@ip-address//directory/" {  
            password "encrypted-password";  
        }  
    }  
}  
}
```

Example: Configuring a Flow Collector Interface on an M, MX or T Series Router

IN THIS SECTION

- [Verifying Your Work | 217](#)

Figure 23 on page 209 shows the path traveled by monitored traffic as it passes through the router. Packets arrive at input interfaces **so-0/1/0**, **so-3/0/0**, and **so-3/1/0**. The raw packets are directed into a filter-based forwarding routing instance and processed into flow records by the monitoring services interfaces **mo-7/1/0**, **mo-7/2/0**, and **mo-7/3/0**. The flow records are compressed into files at the flow collector interfaces **cp-6/0/0** and **cp-7/0/0** and sent to the FTP server for analysis. Finally, a mandatory class-of-service (CoS) configuration is applied to export channels 0 and 1 on the flow collector interfaces to manage the outgoing processed files.

Router 1

```
[edit]
chassis {
  fpc 6 {
    pic 0 {
      monitoring-services {
        application flow-collector; # This converts a Monitoring Services II PIC
      } # into a flow collector interface.
    }
  }
  fpc 7 {
    pic 0 {
      monitoring-services {
        application flow-collector; # This converts a Monitoring Services II PIC
      } # into a flow collector interface.
    }
  }
}
interfaces {
  cp-6/0/0 {
    unit 0 { # Logical interface .0 on a flow collector interface is export
      family inet { # channel 0 and sends records to the FTP server.
        filter {
          output cp-ftp; # Apply the CoS filter here.
        }
        address 10.0.0.1/32 {
          destination 10.0.0.2;
        }
      }
    }
    unit 1 { # Logical interface .1 on a flow collector interface is
      export
      family inet { # channel 1 and sends records to the FTP server.
```

```

        filter {
            output cp-ftp; # Apply the CoS filter here.
        }
        address 10.1.1.1/32 {
            destination 10.1.1.2;
        }
    }
}

    unit 2 { # Logical interface .2 on a flow collector interface is the
flow
    family inet { # receive channel that communicates with the Routing Engine.
        address 10.2.2.1/32 { # Do not apply a CoS filter on logical interface .2.
            destination 10.2.2.2;
        }
    }
}

    cp-7/0/0 {
        unit 0 { # Logical interface .0 on a flow collector interface is
export
    family inet { # channel 0 and sends records to the FTP server.
        filter {
            output cp-ftp; # Apply the CoS filter here.
        }
        address 10.3.3.1/32 {
            destination 10.3.3.2;
        }
    }
}

        unit 1 { # Logical interface .1 on a flow collector interface is
export
    family inet { # channel 1 and sends records to the FTP server.
        filter {
            output cp-ftp; # Apply the CoS filter here.
        }
        address 10.4.4.1/32 {
            destination 10.4.4.2;
        }
    }
}

        unit 2 { # Logical interface .2 on a flow collector interface is the
flow
    family inet { # receive channel that communicates with the Routing Engine.

```

```

        address 10.5.5.1/32 { # Do not apply a CoS filter on logical interface .2.
            destination 10.5.5.2;
        }
    }
}

    fe-1/3/0 { # This is the exit interface leading to the first FTP server.
unit 0 {
    family inet {
        address 192.168.56.90/30;
    }
}

    ge-1/0/0 { # This is the exit interface leading to the second FTP server.
unit 0 {
    family inet {
        address 192.168.252.2/24;
    }
}

    mo-7/1/0 { # This is the first interface that creates flow records.
unit 0 {
    family inet;
}

    mo-7/2/0 { # This is the second interface that creates flow records.
unit 0 {
    family inet;
}

    mo-7/3/0 { # This is the third interface that creates flow records.
unit 0 {
    family inet;
}

    so-0/1/0 { # This is the first input interface that receives traffic to be
monitored.
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode; # This allows the interface to be passively
monitored.
        family inet {
            filter {

```

```

input catch; # The filter-based forwarding filter is
applied here.
    }
    }
}

so-3/0/0 { # This is the second interface that receives traffic to be
monitored.
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode; # This allows the interface to be passively
monitored.
        family inet {
            filter {
                input catch; # The filter-based forwarding filter is
applied here.
            }
        }
    }
}

so-3/1/0 { # This is the third interface that receives traffic to be
monitored.
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode; # This allows the interface to be passively
monitored.
        family inet {
            filter {
                input catch; # The filter-based forwarding filter is
applied here.
            }
        }
    }
}

forwarding-options {
    monitoring group1 { # Always define your monitoring group here.
        family inet {
            output {
                export-format cflowd-version-5;
                flow-active-timeout 60;
                flow-inactive-timeout 15;
                flow-export-destination collector-pic; # Sends records to

```


the flow collector.

```

        interface mo-7/1/0.0 {
            source-address 192.168.252.2;
        }
        interface mo-7/2/0.0 {
            source-address 192.168.252.2;
        }
        interface mo-7/3/0.0 {
            source-address 192.168.252.2;
        }
    }
}

routing-options {
    interface-routes {
        rib-group inet common;
    }
    rib-groups {
        common {
            import-rib [ inet.0 fbf_instance.inet.0 ];
        }
    }
    forwarding-table {
        export pplb;
    }
}

policy-options {
    policy-statement pplb {
        then {
            load-balance per-packet;
        }
    }
}

class-of-service { # A class-of-service configuration for the flow collector interface
    interfaces { # is mandatory when implementing flow collector services.
        cp-6/0/0 {
            scheduler-map cp-map;
        }
        cp-7/0/0 {
            scheduler-map cp-map;
        }
    }
}

```

```

scheduler-maps {
    cp-map {
        forwarding-class best-effort scheduler Q0;
        forwarding-class expedited-forwarding scheduler Q1;
        forwarding-class network-control scheduler Q3;
    }
}
schedulers {
    Q0 {
        transmit-rate remainder;
        buffer-size percent 90;
    }
    Q1 {
        transmit-rate percent 5;
        buffer-size percent 5;
        priority strict-high;
    }
    Q3 {
        transmit-rate percent 5;
        buffer-size percent 5;
    }
}
}
firewall {
    family inet {
        filter cp-ftp { # This filter provides CoS for flow collector interface traffic.
            term t1 {
                then forwarding-class expedited-forwarding;
            }
        }
    }
    filter catch { # This firewall filter sends incoming traffic into the
        interface-specific; # filter-based forwarding routing instance.
        term def {
            then {
                count counter;
                routing-instance fbf_instance;
            }
        }
    }
}
routing-instances {
    fbf_instance { # This instance sends traffic to the monitoring services

```

```

interface.
    instance-type forwarding;
    routing-options {
        static {
            route 0.0.0.0/0 next-hop mo-7/1/0.0;
        }
    }
}
}
services {
    flow-collector { # Define properties for flow collector interfaces here.
        analyzer-address 10.10.10.1; # This is the IP address of the analyzer.
        analyzer-id server1; # This helps to identify the analyzer.
        retry 3; # Maximum number of attempts by the PIC to send a file
transfer log.
        retry-delay 30; # The time interval between attempts to send a file
transfer log.
        destinations { # This defines the FTP servers that receive flow
collector output.
            "ftp://user@192.168.56.89//tmp/collect1/" { # The primary FTP
server.
                password "$ABC123"; # SECRET-DATA
            }
            "ftp://user@192.168.252.1//tmp/collect2/" { # The second FTP
server.
                password "$ABC123"; # SECRET-DATA
            }
        }
        file-specification { # Define sets of flow collector characteristics
here.
            def-spec {
            }
            data-format flow-compressed; # The default compressed output
format.
        }
        f1 {
            name-format "cFlowd-py69Ni69-0-%D_%T-%I_%N.bcp.bi.gz";
            data-format flow-compressed; # The default compressed output
format.
            transfer timeout 1800 record-level 1000000; # Here are configured
values.
        }
    }
}

```

```

        interface-map { # Allows you to map interfaces to flow collector interfaces.
            file-specification def-spec; # Flows generated for default traffic are
sent to the

            collector cp-7/0/0; # default flow collector interface cp-7/0/0.
            so-0/1/0.0 {# Flows generated for the so-0/1/0 interface are sent
                collector cp-6/0/0; # to cp-6/0/0, and the file-specification
used is "default".
            }

            so-3/0/0.0 { # Flows generated for the so-3/0/0 interface are sent
                file-specification f1; # to cp-6/0/0, and the file-specification
used is "f1."

                collector cp-6/0/0;
            }

            so-3/1/0.0; # Because no settings are defined, flows generated for this
        }

        transfer-log-archive { # Sends flow collector interface log files to an FTP
server.
            filename-prefix so_3_0_0_log;
            maximum-age 15;
            archive-sites {
                "ftp://user@192.168.56.89//tmp/transfers/" {
                    password "$ABC123";
                }
            }
        }
    }
}

```

Verifying Your Work

To verify that your flow collector configuration is working, use the following commands on the monitoring station that is configured for flow collection:

- `clear services flow-collector statistics`
- `request services flow-collector change-destination (primary | secondary)`
- `request services flow-collector test-file-transfer`
- `show services flow-collector file interface (detail | extensive | terse)`
- `show services flow-collector (detail | extensive)`
- `show services flow-collector input interface (detail | extensive | terse)`

The following section shows the output of the show commands used with the configuration example:

```

user@router1> show services flow-collector input interface cp-6/0/0 detail
Interface                Packets      Bytes
mo-7/1/0.0              6170       8941592

user@router1> show services flow-collector interface all detail
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
  Packets    Bytes    Flows Uncompressed  Compressed  FTP bytes FTP files
              Bytes      Bytes
      6736  9757936   195993   21855798   3194148         0         0
Flow collector interface: cp-7/0/0
Interface state: Collecting flows
  Packets    Bytes    Flows Uncompressed  Compressed  FTP bytes FTP files
              Bytes      Bytes
         0         0         0         0         0         0         0

user@router1> show services flow-collector input interface cp-6/0/0 extensive
Interface                Packets      Bytes
mo-7/1/0.0              6260       9074096

user@router1> show services flow-collector interface cp-6/0/0 extensive
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Memory:
  Used: 19593212, Free: 479528656
Input:
  Packets: 6658, per second: 0, peak per second: 0
  Bytes: 9647752, per second: 12655, peak per second: 14311
  Flow records processed: 193782, per second: 252, peak per second: 287
Allocation:
  Blocks allocated: 174, per second: 0, peak per second: 0
  Blocks freed: 0, per second: 0, peak per second: 0
  Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
  Files created: 1, per second: 0, peak per second: 0
  Files exported: 0, per second: 0, peak per second: 0
  Files destroyed: 0, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 21075152, per second: 52032, peak per second: 156172
  Compressed bytes: 3079713, per second: 7618, peak per second: 22999

```

```

Packet drops:
  No memory: 0, Not IP: 0
  Not IPv4: 0, Too small: 0
  Fragments: 0, ICMP: 0
  TCP: 0, Unknown: 0
  Not JUNOS flow: 0
File Transfer:
  FTP bytes: 0, per second: 0, peak per second: 0
  FTP files: 0, per second: 0, peak per second: 0
  FTP failure: 0
Export channel: 0
  Current server: Secondary
  Primary server state: OK, Secondary server state: OK
Export channel: 1
  Current server: Secondary
  Primary server state: OK, Secondary server state: OK

user@router1> show services flow-collector file interface cp-6/0/0 terse
File name                                     Flows State
cFlowd-py69Ni69-0-20031112_014301-so_3_0_0_0.bcp.bi.gz  185643 Active

user@router1> show services flow-collector file interface cp-6/0/0 detail
Filename: cFlowd-py69Ni69-0-20031112_014301-so_3_0_0_0.bcp.bi.gz
Throughput:
  Flow records: 187067, Uncompressed bytes: 21121960, Compressed bytes: 2965643
Status:
  State: Active, Transfer attempts: 0

user@router1> show services flow-collector file interface cp-6/0/0 extensive
Filename: cFlowd-py69Ni69-0-20031112_014301-so_3_0_0_0.bcp.bi.gz
Throughput:
  Flow records: 188365, per second: 238, peak per second: 287
  Uncompressed bytes: 21267756, per second: 27007, peak per second: 32526
  Compressed bytes: 2965643, per second: 0, peak per second: 22999
Status:
  Compressed blocks: 156, Block count: 156
  State: Active, Transfer attempts: 0

```

To clear statistics for a flow collector interface, issue the `clear services flow-collector statistics interface (all | interface-name)` command.

Another useful flow collector option allows you to change the FTP server from primary to secondary and test for FTP transfers. To force the flow collector interface to use a primary or secondary FTP

server, include the **primary** or **secondary** option when you issue the `request services flow-collector change-destination interface cp-fpc/pic/port` command.

If you configure only one primary server and issue this command with the **primary** option, you receive the error message “Destination change not needed.” If the secondary server is not configured and you issue this command with the **secondary** option, you receive the error message “Destination not configured.” Otherwise, when both servers are configured properly, successful output appears as follows.

```
user@router1> request services flow-collector change-destination interface      cp-6/0/0
primary
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful

user@router1> request services flow-collector change-destination interface cp-6/0/0 secondary
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

Other options for the `request services flow-collector change-destination interface cp-fpc/pic/port` command are **immediately** (which forces an instant switchover), **gracefully** (the default behavior that allows a gradual switchover), **clear-files** (which purges existing data files), and **clear-logs** (which purges existing log files).

To verify that transfer log files are being scheduled for delivery to the FTP servers, issue the `request services flow-collector test-file-transfer filename interface cp-fpc/pic/port` command. Include the desired export channel (zero or one) and target FTP server (primary or secondary) with this command.

```
user@router1> request services flow-collector test-file-transfer test_file      interface
cp-6/0/0 channel-one primary
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Response: Test file transfer successfully scheduled
```

Another way you can check for the success of your file transfers is by analyzing the transfer log. A transfer log sends detailed information about files that are collected and processed by the flow collector interface. [Table 33 on page 221](#) explains the various fields available in the transfer log.

Table 33: Flow Collector Interface Transfer Log Fields

Field	Explanation
fn	Filename
sz	File size
nr	Number of records
ts	Timestamp with the format of year (4 digits), month (2 digits), day (2 digits), hours (2 digits), minutes (2 digits), and seconds (2 digits).
sf	Success flag—The values are 1 for success and 0 for failure.
ul	Server URL
rc	FTP result code
er	FTP error text
tt	Transfer time

This is an example of a successful transfer log:

```
fn="cFlowd-py69Ni69-0-20040227_230438-at_4_0_0_4_3.bcp.bi.gz":sz=552569
:nr=20000:ts="20040227230855":sf=1:ul="ftp://10.63.152.1/tmp/server1/:"rc=250:
er="":tt=3280
```

This is an example of a transfer log when an FTP session fails:

```
fn="cFlowd-py69Ni69-0-20040227_230515-at_4_0_0_2_8.bcp.bi.gz":sz=560436
:nr=20000:ts="20040227230855":sf=1:ul="ftp://10.63.152.1/tmp/server1/:"rc=250
:er="":tt=3290
```


As the flow collector interface receives and processes flow records, the PIC services logging process (fsad) handles the following tasks:

- When the flow collector interface transfers a file to the FTP server, a temporary log file is created in the **/var/log/flowc** directory. The temporary log file has this filenames convention:

<hostname>_<filename_prefix>_YYYYMMDD_hhmmss.tmp

hostname is the hostname of the transfer server, **filename_prefix** is the same value defined with the filename-prefix statement at the **[edit services flow-collector transfer-log-archive]** hierarchy level, **YYYYMMDD** is the year, month, and date, and **hhmmss** is the timestamp indicating hours, minutes, and seconds.

- After the log file has been stored in the router for the length of time specified by the **maximum-age** statement at the **[edit services flow-collector transfer-log-archive]** hierarchy level (the default is 120 minutes), the temporary log file is converted to an actual log file and the temporary file is deleted. The new log file retains the same naming conventions, except the extension is ***.log**.
- When the final log file is created and compressed, the PIC services logging process (fsad) tries to send the log file from the **/var/log/flowc** directory to an FTP server. You can specify up to five FTP servers to receive the log files by including the **archive-sites** statement at the **[edit services flow-collector transfer-log-archive]** hierarchy level. The logging process attempts to send the log file to one server at a time, in order of their appearance in the configuration. Upon the first successful transfer, the log file is deleted and the logging process stops sending log files to the remaining FTP servers in the list.
- If the log file transfer is not successful, the log file is moved to the **/var/log/flowc/failed** directory. Every 30 minutes, the logging process tries to resend the log files. After the log files are transferred successfully, they are deleted from the **/var/log/flowc/failed** directory.

NOTE: If the memory for a flow collector interface is full, the interface might drop incoming packets.

After the flow collector interface successfully delivers the processed information file to the FTP server, you can analyze the file. The file contains detailed information about the flows collected and processed by the flow collector interface. [Table 34 on page 223](#) explains the various fields available in the flow collector interface file.

Table 34: Flow Collector Interface File Fields in Order of Appearance

Field	Explanation
linkDir	Link directory—A randomly generated number used to identify the record
analyzer-address	Analyzer address
analyzer-ID	Analyzer identifier
ifAlias	Interface identifier
source-address	Source address
destination-address	Destination address
packets	Number of packets
bytes	Number of bytes
start-time	Start time
end-time	End time
source-port	Source port
destination-port	Destination port
tcp_flag	TCP flag
protocol	IP protocol number
src_AS_number	Source AS number

Table 34: Flow Collector Interface File Fields in Order of Appearance *(Continued)*

Field	Explanation
dst_AS_number	Destination AS number

This is an example of output from a flow collector interface file:

```
11799241612374557782|10.10.10.1|server1|at_4_0_0_4|192.168.10.100|10.0.0.1|8|
3136|1077926402|1077926402|8224|12336|27|6|0|0
```

Processing and Exporting Multiple Records Using Flow Collection

IN THIS CHAPTER

- [Flow Collection Overview | 225](#)
- [Configuring Flow Collection | 226](#)
- [Example: Configuring Flow Collection | 231](#)
- [Sending cflowd Records to Flow Collector Interfaces | 239](#)
- [Configuring Flow Collection Mode and Interfaces on Router Services PICs on M and T Series Routers | 239](#)

Flow Collection Overview

You can process and export multiple cflowd records with a flow collector interface. You create a flow collector interface on a Monitoring Services II or Multiservices 400 PIC. The flow collector interface combines multiple cflowd records into a compressed ASCII data file and exports the file to an FTP server. To convert a services PIC into a flow collector interface, include the `flow-collector` statement at the `[edit chassis fpc fpc-slot pic pic-slot monitoring-services application]` hierarchy level.

You can use the services PIC for either flow collection or monitoring, but not for both types of service simultaneously. When converting the PIC between service types, you must configure the `flow-collector` statement, take the PIC offline, and then bring the PIC back online. Restarting the router does not enable the new service type.

A flow collector interface, designated by the **cp-fpc/pic/port** interface name, requires three logical interfaces for correct operation. Units 0 and 1 are used to send the compressed ASCII data files to an FTP server, while Unit 2 is used to receive cflowd records from a monitoring services interface.

NOTE: Unlike conventional interfaces, the address statement at the `[edit interfaces cp-fpc/pic/port unit unit-number family inet]` hierarchy level corresponds to the IP address of the Routing Engine. Likewise, the destination statement at the `[edit interfaces cp-fpc/pic/port unit unit-number family inet`

address *ip-address*] hierarchy level corresponds to the IP address of the flow collector interface. As a result, you must configure the destination statement for Unit 0 and 1 with */oca/*addresses that can reach the FTP server. Similarly, configure the destination statement for Unit 2 with a */oca/*IP address so it can reach the monitoring services interface that sends cflowd records.

To activate flow collector services after the services PIC is converted into a flow collector, include the `flow-collector` statement at the `[edit services]` hierarchy level.

After you activate the flow collector, you need to configure the following components:

- Destination of the FTP server
- File specifications
- Input interface-to-flow collector interface mappings
- Transfer log settings

RELATED DOCUMENTATION

[Configuring Flow Collection | 226](#)

[Sending cflowd Records to Flow Collector Interfaces | 239](#)

[Configuring Flow Collection Mode and Interfaces on Router Services PICs on M and T Series Routers | 239](#)

Configuring Flow Collection

IN THIS SECTION

- [Configuring Destination FTP Servers for Flow Records | 227](#)
- [Configuring a Packet Analyzer | 227](#)
- [Configuring File Formats | 228](#)
- [Configuring Interface Mappings | 229](#)
- [Configuring Transfer Logs | 229](#)
- [Configuring Retry Attempts | 230](#)

Configuring Destination FTP Servers for Flow Records

Flow collection destinations are where the compressed ASCII data files are sent after the cflowd records are collected and processed. To specify the destination FTP server, include the `destinations` statement at the `[edit services flow-collector]` hierarchy level. You can specify up to two FTP server destinations and include the password for each configured server. If two FTP servers are configured, the first server in the configuration is the primary server and the second is a backup server.

To configure a destination for flow collection files, include the `destinations` statement at the `[edit services flow-collector]` hierarchy level:

```
[edit services flow-collector]
destinations {
  ftp:url {
    password "password";
  }
}
```

To specify the destination FTP server, include the `ftp:url` statement. The value `url` is the FTP server address for the primary flow collection destination and can include macros.

When you include macros in the `ftp:url` statement, a directory can be created only for a single level. For example, the path `ftp://10.2.2.2/%m/%Y` expands to `ftp://10.2.2.2/01/2005`, and the software attempts to create the directory `01/2005` on the destination FTP server. If the `01/` directory already exists on the destination FTP server, the software creates the **/2005/ directory** one level down. If the `01/` directory does not exist on the destination FTP server, the software cannot create the **/2005/ directory**, and the FTP server destination fails. For more information about macros, see ["ftp" on page 1138](#).

To specify the FTP server password, include the `password "password"` statement. The password must be enclosed in quotation marks. You can specify up to two destination FTP servers. The first destination specified is considered the primary destination.

Configuring a Packet Analyzer

You can specify values for the IP address and identifier of a packet analyzer to which the flow collector interface sends traffic for analysis. The values you specify here override any default values configured elsewhere.

To configure an IP address and identifier for the packet analyzer, include the `analyzer-address` and `analyzer-id` statements at the `[edit services flow-collector]` hierarchy level:

```
[edit services flow-collector]
analyzer-address address;
analyzer-id name;
```

Configuring File Formats

You configure data file formats, name formats, and transfer characteristics for the flow collection files. File records are sent to the destination FTP server when the timer expires or when a preset number of records are received, whichever comes first.

To configure the flow collection file format, include the `file-specification` statement at the `[edit services flow-collector]` hierarchy level:

```
[edit services flow-collector]
file-specification {
  variant variant-number {
    data-format format;
    name-format format;
    transfer {
      record-level number;
      timeout seconds;
    }
  }
}
```

To set the data file format, include the `data-format` statement. To set the file name format, include the `name-format` statement. To set the export timer and file size thresholds, include the `transfer` statement and specify values for the `timeout` and `record-level` options.

For example, you can specify the name format as follows:

```
[edit services flow-collector file-specification variant variant-number]
name-format "cFlowd-py69Ni69-0-%D_%T-%I_%N.bcp.bi.gz";
```

In this example, `cFlowd-py69Ni69-0` is the static portion used verbatim, `%D` is the date in `YYYYMMDD` format, `%T` is the time in `HHMMSS` format, `%I` is the value of `ifAlias`, `%N` is the generation number, and `bcp.bi.gz` is a user-configured string. A number of macros are supported for expressing the date and time

information in different ways; for a complete list, see the summary section for ["name-format" on page 1251](#).

Configuring Interface Mappings

You can match an input interface with a flow collector interface and apply the preset file specifications to the input interface.

To configure an interface mapping, include the `interface-map` statement at the `[edit services flow-collector]` hierarchy level:

```
[edit services flow-collector]
interface-map {
  collector interface-name;
  file-specification variant-number;
  interface-name {
    collector interface-name;
    file-specification variant-number;
  }
}
```

To configure the default flow collector and file specifications for all input interfaces, include the `file-specification` and `collector` statements at the `[edit services flow-collector interface-map]` hierarchy level. To override the default settings and apply flow collector and file specifications to a specific input interface, include the `file-specification` and `collector` statements at the `[edit services flow-collector interface-map interface-name]` hierarchy level.

Configuring Transfer Logs

You can configure the filename, export interval, maximum size, and destination FTP server for log files containing the transfer activity history for a flow collector interface.

To configure a transfer log, include the `transfer-log-archive` statement at the `[edit services flow-collector]` hierarchy level:

```
[edit services flow-collector]
transfer-log-archive {
  archive-sites {
    ftp:url {
      password "password";
      username username;
```



```

    }
  }
  filename-prefix prefix;
  maximum-age minutes;
}

```

To configure the destination for archiving files, include the `archive-sites` statement. Specify the filename as follows:

```

[edit services flow-collector transfer-log]
filename "cFlowd-py69Ni69-0-%D_%T";

```

where `cFlowd-py69Ni69-0` is the static portion used verbatim, `%D` is the date in YYYYMMDD format, and `%T` is the time in HHMMSS format.

You can optionally include the following statements:

- `filename-prefix`—Sets a standard prefix for all the logged files.
- `maximum-age`—Specifies the duration a file remains on the server. The range is 1 through 360 minutes.

Configuring Retry Attempts

You can specify values for situations in which the flow collector interface needs more than one attempt to transfer log files to the FTP server:

- Maximum number of retry attempts
- Amount of time the flow collector interface waits between successive retries

To configure retry settings, include the `retry` and `retry-delay` statements at the `[edit services flow-collector]` hierarchy level:

```

retry number;
retry-delay seconds;

```

The `retry` value can be from 0 through 10. The `retry-delay` value can be from 0 through 60 seconds.

RELATED DOCUMENTATION


[Flow Collection Overview](#) | 225

[Sending cflowd Records to Flow Collector Interfaces](#) | 239

Example: Configuring Flow Collection

Figure 24 on page 231 shows the path traveled by monitored traffic as it passes through the router. Packets arrive at input interfaces so-0/1/0, so-3/0/0, and so-3/1/0. The raw packets are directed into a filter-based forwarding routing instance and processed into cflowd records by the monitoring services interfaces mo-7/1/0, mo-7/2/0, and mo-7/3/0. The cflowd records are compressed into files at the flow collector interfaces cp-6/0/0 and cp-7/0/0 and sent to the FTP server for analysis. Finally, a mandatory class-of-service (CoS) configuration is applied to export channels 0 and 1 on the flow collector interfaces to manage the outgoing processed files.

Figure 24: Flow Collector Interface Topology Diagram



```
[edit]
chassis {
  fpc 6 {
    pic 0 {
```

```

        monitoring-services {
            application flow-collector; # This converts a Monitoring Services II or
                                      # Multiservices 400 PIC into a flow collector interface.
        }
    }
}
fpc 7 {
    pic 0 {
        monitoring-services {
            application flow-collector; # This converts a Monitoring Services II or
                                      # Multiservices 400 PIC into a flow collector interface.
        }
    }
}
}
interfaces {
    cp-6/0/0 {
        unit 0 { # Logical interface .0 on a flow collector interface is export
            family inet { # channel 0 and sends records to the FTP server.
                filter {
                    output cp-ftp; # Apply the CoS filter here.
                }
                address 10.0.0.1/32 {
                    destination 10.0.0.2;
                }
            }
        }
        unit 1 { # Logical interface .1 on a flow collector interface is export
            family inet { # channel 1 and sends records to the FTP server.
                filter {
                    output cp-ftp; # Apply the CoS filter here.
                }
                address 10.1.1.1/32 {
                    destination 10.1.1.2;
                }
            }
        }
        unit 2 { # Logical interface .2 on a flow collector interface is the flow
            family inet { # receive channel that communicates with the Routing Engine.
                address 10.2.2.1/32 { # Do not apply a CoS filter on logical interface .2.
                    destination 10.2.2.2;
                }
            }
        }
    }
}

```

```

    }
}
cp-7/0/0 {
    unit 0 {# Logical interface .0 on a flow collector interface is export
        family inet {# channel 0 and sends records to the FTP server.
            filter {
                output cp-ftp;# Apply the CoS filter here.
            }
            address 10.3.3.1/32 {
                destination 10.3.3.2;
            }
        }
    }
}
unit 1 {# Logical interface .1 on a flow collector interface is export
    family inet {# channel 1 and sends records to the FTP server.
        filter {
            output cp-ftp;# Apply the CoS filter here.
        }
        address 10.4.4.1/32 {
            destination 10.4.4.2;
        }
    }
}
unit 2 {# Logical interface .2 on a flow collector interface is the flow
    family inet {# receive channel that communicates with the Routing Engine.
        address 10.5.5.1/32 {# Do not apply a CoS filter on logical interface .2.
            destination 10.5.5.2;
        }
    }
}
}
fe-1/3/0 { # This is the exit interface leading to the first FTP server.
    unit 0 {
        family inet {
            address 192.168.56.90/30;
        }
    }
}
ge-1/0/0 { # This is the exit interface leading to the second FTP server.
    unit 0 {
        family inet {
            address 192.168.252.2/24;
        }
    }
}

```

```

    }
}
mo-7/1/0 { # This is the first interface that creates cflowd records.
    unit 0 {
        family inet;
    }
}
mo-7/2/0 { # This is the second interface that creates cflowd records.
    unit 0 {
        family inet;
    }
}
mo-7/3/0 { # This is the third interface that creates cflowd records.
    unit 0 {
        family inet;
    }
}
so-0/1/0 { # This is the first input interface that receives traffic to be monitored.
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode; # This allows the interface to be passively monitored.
        family inet {
            filter {
                input catch; # The filter-based forwarding filter is applied here.
            }
        }
    }
}
so-3/0/0 { # This is the second interface that receives traffic to be monitored.
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode; # This allows the interface to be passively monitored.
        family inet {
            filter {
                input catch; # The filter-based forwarding filter is applied here.
            }
        }
    }
}
so-3/1/0 { # This is the third interface that receives traffic to be monitored.
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode; # This allows the interface to be passively monitored.
    }
}

```

```

    family inet {
        filter {
            input catch; # The filter-based forwarding filter is applied here.
        }
    }
}

forwarding-options {
    monitoring group1 {# Always define your monitoring group here.
        family inet {
            output {
                export-format cflowd-version-5;
                flow-active-timeout 60;
                flow-inactive-timeout 15;
                flow-export-destination collector-pic; # Sends records to the flow collector.
                interface mo-7/1/0.0 {
                    source-address 192.168.252.2;
                }
                interface mo-7/2/0.0 {
                    source-address 192.168.252.2;
                }
                interface mo-7/3/0.0 {
                    source-address 192.168.252.2;
                }
            }
        }
    }
}

firewall {
    family inet {
        filter cp-ftp { # This filter provides CoS for flow collector interface traffic.
            term t1 {
                then forwarding-class expedited-forwarding;
            }
        }
    }

    filter catch { # This firewall filter sends incoming traffic into the
        interface-specific;# filter-based forwarding routing instance.
        term def {
            then {
                count counter;
                routing-instance fbf_instance;
            }
        }
    }
}

```

```

    }
  }
  routing-options {
    interface-routes {
      rib-group inet common;
    }
    rib-groups {
      common {
        import-rib [inet.0 fbf_instance.inet.0];
      }
    }
    forwarding-table {
      export pplb;
    }
  }
  policy-options {
    policy-statement pplb {
      then {
        load-balance per-packet;
      }
    }
  }
  routing-instances {
    fbf_instance { # This instance sends traffic to the monitoring services interface.
      instance-type forwarding;
      routing-options {
        static {
          route 0.0.0.0/0 next-hop mo-7/1/0.0;
        }
      }
    }
  }
  class-of-service { # A class-of-service configuration for the flow collector interface
    interfaces { # is required for flow collector services.
      cp-6/0/0 {
        scheduler-map cp-map;
      }
      cp-7/0/0 {
        scheduler-map cp-map;
      }
    }
  }
  scheduler-maps {

```

```

cp-map {
    forwarding-class best-effort scheduler Q0;
    forwarding-class expedited-forwarding scheduler Q1;
    forwarding-class network-control scheduler Q3;
}
}
schedulers {
    Q0 {
        transmit-rate remainder;
        buffer-size percent 90;
    }
    Q1 {
        transmit-rate percent 5;
        buffer-size percent 5;
        priority strict-high;
    }
    Q3 {
        transmit-rate percent 5;
        buffer-size percent 5;
    }
}
}
services {
    flow-collector { # Define properties for flow collector interfaces here.
        analyzer-address 10.10.10.1; # This is the IP address of the analyzer.
        analyzer-id server1; # This helps to identify the analyzer.
        retry 3; # Maximum number of attempts by the PIC to send a file transfer log.
        retry-delay 30; # The time interval between attempts to send a file transfer log.
        destinations { # This defines the FTP servers that receive flow collector output.
            "ftp://user@192.168.56.89//tmp/collect1/" { # The primary FTP server.
                password "$ABC123"; # SECRET-DATA
            }
            "ftp://user@192.168.252.1//tmp/collect2/" { # The secondary FTP server.
                password "$ABC123"; # SECRET-DATA
            }
        }
    }
    file-specification { # Define sets of flow collector characteristics here.
        def-spec {
            name-format "default-allInt-0-%D_%T-%I_%N.bcp.bi.gz";
            data-format flow-compressed; # The default compressed output format.
        } # When no overrides are specified, a collector uses default transfer values.
        f1 {
            name-format "cFlowd-py69Ni69-0-%D_%T-%I_%N.bcp.bi.gz";
            data-format flow-compressed; # The default compressed output format.
        }
    }
}

```



```

        transfer timeout 1800 record-level 1000000; # Here are configured values.
    }
}
interface-map { # Allows you to map interfaces to flow collector interfaces.
    file-specification def-spec; # Flows generated for default traffic are sent to the
    collector cp-7/0/0; # default flow collector interface "cp-7/0/0".
    so-0/1/0.0 { # Flows generated for the so-0/1/0 interface are sent
        collector cp-6/0/0; # to cp-6/0/0, and the file-specification used is
    } # "default."
    so-3/0/0.0 { # Flows generated for the so-3/0/0 interface are sent
        file-specification f1; # to cp-6/0/0, and the file-specification used is "f1."
        collector cp-6/0/0;
    }
    so-3/1/0.0; # Because no settings are defined, flows generated for this
} # interface use interface cp-7/0/0 and the default file specification.
transfer-log-archive { # Sends flow collector interface log files to an FTP server.
    filename-prefix so_3_0_0_log;
    maximum-age 15;
    archive-sites {
        "ftp://user@192.168.56.89//tmp/transfers/" {
            password "$ABC123";
        }
    }
}
}
}
}

```

RELATED DOCUMENTATION

[Flow Collection Overview | 225](#)

[Configuring Flow Collection | 226](#)

[Sending cflowd Records to Flow Collector Interfaces | 239](#)

[Configuring Flow Collection Mode and Interfaces on Router Services PICs on M and T Series Routers | 239](#)

Sending cflowd Records to Flow Collector Interfaces

To specify a flow collector interface as the destination for cflowd records coming from a services PIC, include the `collector-pic` statement at the `[edit forwarding-options monitoring group-name family inet output flow-export-destination]` hierarchy level:

```
[edit forwarding-options monitoring group-name family inet output flow-export-destination]
collector-pic;
```

You can select either the flow collector interface or a cflowd server as the destination for cflowd records, but not both at the same time.

RELATED DOCUMENTATION

[Flow Collection Overview | 225](#)

[Configuring Flow Collection | 226](#)

[Configuring Flow Collection Mode and Interfaces on Router Services PICs on M and T Series Routers | 239](#)

[Example: Configuring Flow Collection | 231](#)

Configuring Flow Collection Mode and Interfaces on Router Services PICs on M and T Series Routers

You can select the services PIC to run in either flow collection mode or monitoring mode, but not both.

To set the services PIC to run in flow collection mode, include the `flow-collector` statement at the `[edit chassis fpc slot-number pic pic-number monitoring-services application]` hierarchy level:

```
[edit chassis fpc slot-number pic pic-number monitoring-services application]
flow-collector;
```

To specify flow collection interfaces, you configure the `cp` interface at the `[edit interfaces]` hierarchy level:

```
[edit interfaces]
cp-fpc/pic/port {
```

```
...  
}
```

RELATED DOCUMENTATION

[Flow Collection Overview | 225](#)

[Configuring Flow Collection | 226](#)

[Sending cflowd Records to Flow Collector Interfaces | 239](#)

[Example: Configuring Flow Collection | 231](#)

Logging Flow Monitoring Records with Version 9 and IPFIX Templates for NAT Events

IN THIS CHAPTER

- Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241
- Configure Active Flow Monitoring Logs for NAT44/NAT64 | 254
- Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256
- Exporting Syslog Messages to an External Host Without Flow Monitoring Formats Using an MX Series Router or NFX250 | 258
- Exporting Version 9 Flow Data Records to a Log Collector Overview Using an MX Series Router or NFX250 | 259
- Understanding Exporting IPFIX Flow Data Records to a Log Collector Using an MX Series Router or NFX250 | 260
- Mapping Between Field Values for Version 9 Flow Templates and Logs Exported From an MX-Series Router or NFX250 | 262
- Mapping Between Field Values for IPFIX Flow Templates and Logs Exported From an MX Series Router or NFX250 | 265
- Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272
- Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275

Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250

Starting with Junos OS Release 14.2R2 and 15.1R1, you can configure MX Series routers with MS-MPCs and MS-MICs to log network address translation (NAT) events using the Junos Traffic Vision (previously known as Jflow) version 9 or IPFIX (version 10) template format. You can also configure MX Series routers with MX-SPC3 services cards with this capability starting from Junos OS Release 19.3R2.

NAT event logger generates messages in flow monitoring format for various NAT events, such as the creation of a NAT entry, deletion of a NAT entry, and for invalid NAT processing (such as NAT address pools or address values being exhausted for allocation). These events also support NAT64 translations (translation of IPv6 addresses to IPv4 addresses), binding information base (BIB) events, and more detailed error generation. The generated records or logs for NAT events in flow template format are sent from the MS-MIC or MS-MPC or MX-SPC3 to the specified host or external device that functions as the NetFlow collector. This method of generating flow monitoring records for NAT events enables cohesive and streamlined analysis of NAT traffic and troubleshooting of NAT-related problems. You can enable the capability to send flow monitoring records for NAT operations to an external collector and the capability to use the system logging protocol (syslog) to generate session logging for different services at the same time.

The flow records and the templates are encapsulated in an UDP or IP packet and sent to the collector. However, TCP-based logging of monitoring records for NAT events is not supported. Carrier-grade NAT (CGN) devices are required to log events creation and deletion of translations and information about the resources it manages. Flow monitoring logs can be optionally configured in your network topology in addition to the system logging (syslog) capability, which causes logs to be saved from the PIC to either the in the **/var/log** directory of the Routing Engine (local) or to an external server (remote). Generally, flow collectors are the part of a vast network infrastructure containing several third-party devices, which perform various correlations and mappings with logs of other databases. Therefore, collection of NAT-related flow monitoring records as logs or template records is useful on the hosts or devices that function as collectors in an overall and comprehensive perspective. You can enable logging of flow monitoring records for NAT events at the service-set level to enable version 9 or IPFIX flow records to be generated as logs when NAT is configured on the router.

The NetFlow collector receives flow records in version 9 or IPFIX format from one or more exporters. It processes the received export packets by parsing and saving the flow record details. Flow records can be optionally aggregated before being stored on the hard disk. The NetFlow collector is also referred to as the collector. The exporter monitors packets entering an observation point and creates flows from these packets. The information from these flows is exported in the form of flow records to the NetFlow Collector. An observation point is a location in the network where IP packets can be overseen and monitored; for example, one or a set of interfaces on a network device such as a router. Every observation point is associated with an observation domain, which is a cluster of observation points, and constitutes the largest aggregatable set of flow information at the network device with NetFlow services enabled.

A FlowSet is a generic term for a collection of Flow Records that have a similar pattern or format. In an export packet, one or more FlowSets follow the packet header. A Template FlowSet comprises one or more template records that have been grouped together in an export packet. An Options Template FlowSet contains one or more Options Template records that are combined together in an export packet. A Data FlowSet is one or more records, of the same type, that are grouped together in an export packet. Each record is either a flow data record or an options data record that has been previously specified by a Template Record or an Options Template Record. One of the essential elements in the NetFlow format is the Template FlowSet. Templates vastly enhance the flexibility of the Flow Record

format because they allow the collector to process Flow Records without necessarily knowing the interpretation of all the data in the Flow Record.

You can configure the capability to transmit records or log messages in version 9 and IPFIX traffic flow formats generated for NAT events to an external, off-box high-speed NetFlow collector for easy and effective monitoring and diagnosis of the logs. By default, this functionality is disabled. With a high number of NAT events, this mechanism of exporting logs to an external log collector might cause scaling considerations such as loss of a few flow records. To enable the mechanism to record logging messages in flow monitoring format for NAT events, you can now include the `jflow-log` statement at the `[edit services]` hierarchy level. You can configure a collector, which is an external host to which the flow monitoring formatted logs are sent, or a group of collectors. A group of collectors is useful in scenarios in which you want to combine a set of collector devices and define common settings for logging NAT events for all the collectors in the cluster or group.

To configure a collector and its parameters, such as the source IP address from which the records are sent and the destination address of the collector, include the `collector collector-name` statement and its substatements at the `[edit services jflow-log]` hierarchy level. To specify a collector group or a cluster, include the `collector-group collector-group-name` statement and its substatements at the `[edit services jflow-log]` hierarchy level.

You need to configure a template profile and associate it with the collector. The profile defines the characteristics of the flow monitoring record template, such as the version of flow monitoring (version 9 or IPFIX), the refresh rate, in either packets or seconds, and the type of service or application (NAT in this case) for which flow records must be sent to the collector. To specify a template profile, include the `template-profile template-profile-name` statement at the `[edit services jflow-log]` hierarchy level. To specify the maximum number of messages to be collected per second for NAT error events, include the `message-rate-limit messages-per-second` statement at the `[edit interfaces ms-interface-name service-options jflow-log]` hierarchy level.

Use of version 9 and IPFIX allows you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration. You must define a template profile properties for a NAT service and associate the defined template profile with a service set to enable the flow monitoring log functionality for NAT events. To define the template profile characteristics for recording flow monitoring logs for NAT events, include the `template-profile template-profile-name` statement at the `[edit services jflow-log]` hierarchy level. To associate the template profile for recording flow monitoring logs for NAT events with a service-set level, which applies for all the services in the system, include the `template-profile template-profile-name` statement at the `[edit services service-set service-set-name]` hierarchy level.

To view statistical information on the logs generated in flow monitoring format for the interfaces and service sets configured on the system, use the `show services service-sets statistics jflow-log` command.

The following system log messages for various NAT events are logged using the system logging (syslog) capability:

- JSERVICES_SESSION_OPEN
- JSERVICES_SESSION_CLOSE
- JSERVICES_NAT_OUTOF_ADDRESSES
- JSERVICES_NAT_OUTOF_PORTS
- JSERVICES_NAT_RULE_MATCH
- JSERVICES_NAT_POOL_RELEASE
- JSERVICES_NAT_PORT_BLOCK_ALLOC
- JSERVICES_NAT_PORT_BLOCK_RELEASE
- JSERVICES_NAT_PORT_BLOCK_ACTIVE

The following NAT events are logged using the flow monitoring log capability using version 9 and IPFIX flow templates:

- NAT44 session create
- NAT44 session delete
- NAT addresses exhausted
- NAT64 session create
- NAT64 session delete
- NAT44 BIB create
- NAT44 BIB delete
- NAT64 BIB create
- NAT64 BIB delete
- NAT ports exhausted
- NAT quota exceeded
- NAT Address binding create
- NAT Address binding delete
- NAT port block allocation
- NAT port block release

- NAT port block active

Table 35 on page 245 describes the flow template format for NAT44 session creation and deletion events. The Information Element (IE) names and their IANA IDs are as defined in the IP Flow Information Export (IPFIX) Entities specification by the Internet Assigned Numbering Authority (IANA).

Table 35: Flow Template Format for NAT44 Session Creation and Deletion

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
sourceIPv4Address	32	8
postNATSourceIPv4Address	32	225
protocolIdentifier	8	4
sourceTransportPort	16	7
postNAPTsourceTransportPort	16	227
destinationIPv4Address	32	12
postNATDestinationIPv4Address	32	226
destinationTransportPort	16	11
postNAPTdestinationTransportPort	16	228
natOriginatingAddressRealm	8	229
natEvent	8	230
flowDurationMilliseconds	32	161

Table 35: Flow Template Format for NAT44 Session Creation and Deletion (Continued)

Information Element (IE)	Size (bits)	IANA ID
initiatorPackets	64	298
responderPackets	64	299
flowDirection	8	61

[Table 36 on page 246](#) describes the flow template format for NAT64 session creation and deletion events.

Table 36: Flow Template Format for NAT64 Session Creation and Deletion

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
sourceIPv6Address	128	27
postNATSourceIPv6Address	32	225
protocolIdentifier	8	4
sourceTransportPort	16	7
postNAPTsourceTransportPort	16	227
destinationIPv6Address	128	28
postNATDestinationIPv6Address	32	226
destinationTransportPort	16	11
postNAPTdestinationTransportPort	16	228

Table 36: Flow Template Format for NAT64 Session Creation and Deletion (Continued)

Information Element (IE)	Size (bits)	IANA ID
natOriginatingAddressRealm	8	229
natEvent	8	230
flowDurationMilliseconds	32	161
initiatorPackets	64	298
responderPackets	64	299
flowDirection	8	61

[Table 37 on page 247](#) describes the flow template format for NAT44 binding information base (BIB) creation and deletion events.

Table 37: Flow Template Format for NAT44 BIB Creation and Deletion

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
sourceIPv4Address	32	8
postNATSourceIPv4Address	32	225
protocolIdentifier	8	4
sourceTransportPort	16	7
postNAPTsourceTransportPort	16	227
natEvent	8	230

[Table 38 on page 248](#) describes the flow template format for NAT64 binding information base (BIB) creation and deletion events.

Table 38: Flow Template Format for NAT64 BIB Creation and Deletion

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
sourceIPv6Address	128	27
postNATSourceIPv6Address	32	225
protocolIdentifier	8	4
sourceTransportPort	16	7
postNAPTsourceTransportPort	16	227
natEvent	8	230

[Table 39 on page 248](#) describes the flow template format for addresses exhaustion events.

Table 39: Flow Template Format for Address Exhausted Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
natEvent	8	230
natPoolName	512	284

[Table 40 on page 249](#) describes the flow template format for ports exhaustion events.

Table 40: Flow Template Format for Ports Exhausted Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
natEvent	8	230
postNATSourceIPv4Address	32	225
protocolIdentifier	8	4

[Table 41 on page 249](#) describes the flow template format for NAT44 quota exceeded events.

Table 41: Flow Template Format for NAT44 Quota Exceeded Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
natEvent	8	230
sourceIPv4Address	32	8

[Table 42 on page 249](#) describes the flow template format for NAT64 quota exceeded events.

Table 42: Flow Template Format for NAT64 Quota Exceeded Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
natEvent	8	230
sourceIPv6Address	128	27

[Table 43 on page 250](#) describes the flow template format for NAT44 address binding creation and deletion events.

Table 43: Flow Template Format for NAT44 Address Binding Creation and Deletion Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
natEvent	8	230
sourceIPv4Address	32	8
postNATSourceIPv4Address	32	225

[Table 44 on page 250](#) describes the flow template format for NAT64 address binding creation and deletion events.

Table 44: Flow Template Format for NAT64 Address Binding Creation and Deletion Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
natEvent	8	230
sourceIPv6Address	128	27
postNATSourceIPv4Address	32	225

[Table 45 on page 251](#) describes the flow template format for NAT44 port block allocation and deallocation events.

Table 45: Flow Template Format for NAT44 Port Block Allocation and Deallocation Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
sourceIPv4Address	32	8
postNATSourceIPv4Address	32	225
portRangeStart	16	361
portRangeEnd	16	362
portRangeStepSize	16	363
portRangeNumPorts	16	364
observationTimeMilliseconds (time when PBA allocated) NOTE: This IE is not included in flow templates when using the MX-SPC3 services card.	64	323
natEvent	8	230

[Table 46 on page 251](#) describes the flow template format for NAT64 port block allocation and deallocation events.

Table 46: Flow Template Format for NAT64 Port Block Allocation and Deallocation Events

Information Element (IE)	Size (bits)	IANA ID
observationTimeMilliseconds	64	323
sourceIPv6Address	128	27

Table 46: Flow Template Format for NAT64 Port Block Allocation and Deallocation Events *(Continued)*

Information Element (IE)	Size (bits)	IANA ID
postNATSourceIPv4Address	32	225
portRangeStart	16	361
portRangeEnd	16	362
portRangeStepSize	16	363
portRangeNumPorts	16	364
observationTimeMilliseconds (time when port block allocation (PBA) is configured) NOTE: This IE is not included in flow templates when using the MX-SPC3 services card.	64	323
natEvent	8	230

In all of the aforementioned templates, the natEvent field maps to one of the values listed in [Table 47 on page 252](#), depending on the type of event.

Table 47: Association Between natEvent Values and Names

natEvent Value	natEvent Name
1	NAT44 Session create
2	NAT44 Session delete
3	NAT Addresses exhausted

Table 47: Association Between natEvent Values and Names (Continued)

natEvent Value	natEvent Name
4	NAT64 Session create
5	NAT64 Session delete
6	NAT44 BIB create
7	NAT44 BIB delete
8	NAT64 BIB create
9	NAT64 BIB delete
10	NAT ports exhausted
11	NAT Quota exceeded
12	NAT Address binding create
13	NAT Address binding delete
14	NAT port block allocation
15	NAT port block release
16	NAT port block active

Release History Table

Release	Description
19.3R2	You can also configure MX Series routers with MX-SPC3 services cards with this capability starting from Junos OS Release 19.3R2.

RELATED DOCUMENTATION

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

Configure Active Flow Monitoring Logs for NAT44/NAT64

IN THIS SECTION

- [Overview | 254](#)
- [Requirements | 254](#)
- [Configuration | 254](#)

Overview

Active Flow Monitoring logs are generated for NAT44 /NAT64 sessions to create or delete events on MX-SPC3 devices.

Requirements

This example uses the following hardware and software components:

- MX480 and MX960 with MX-SPC3
- Junos OS Release 21.2R1

Configuration

IN THIS SECTION

- [Results | 256](#)

To configure Active Flow Monitoring logging on MX-SPC3 devices, perform these tasks:

1. Configure the collectors on an interface.

```
[edit]
user@host# set services jflow-log collector c1 destination-address 10.30.1.2
user@host# set services jflow-log collector c1 destination-port 1055
user@host# set services jflow-log collector c1 source-ip 10.30.1.1
```

2. Configure the collector groups.

```
[edit]
user@host# set services jflow-log collector-group cg1 collector c1
```

3. Configure the template profiles and associate the template profile with the collector or collector group.

```
[edit]
user@host# set services jflow-log template-profile t1 collector-group cg1
```

4. Associate the template profile with the template type.

```
[edit]
user@host# set services jflow-log template-profile t1 template-type nat
```

5. Associate the template profile with the version.

```
[edit]
user@host# set services jflow-log template-profile t1 version ipfix
```

6. Assign the refresh-rate values.

```
[edit]
user@host# set services jflow-log template-profile t1 refresh-rate packets 100
user@host# set services jflow-log template-profile t1 refresh-rate seconds 60
```

7. Associate the template profile with the service set.

```
[edit]
user@host# set services service-set ss1 jflow-log template-profile t1
```

Results

From the configuration mode, confirm your configuration by entering the `show services jflow-log` command in configuration mode. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show services jflow-log
collector c1 {
  destination-address 10.30.1.2;
  destination-port 1055;
  source-ip 10.30.1.1;
}
collector-group cg1 {
  collector c1;
}
template-profile t1 {
  collector-group cg1;
  template-type nat;
  version ipfix;
  refresh-rate {
    packets 100;
    seconds 60;
  }
}
```

Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250

Keep the following points in mind when you configure the capability to generate logs or records in flow monitoring format for NAT events:

- Enabling syslog and Jflow capabilities at the same time might result in scaling impacts because both these mechanisms use a separate infrastructure to transfer records out to the collector.

- High number of NAT events can cause scalability considerations because of the flow monitoring framework too requiring system processes.
- The flow monitoring log infrastructure uses data CPUs to send the logs to the external flow server, which might cause a slight impact on performance.
- An explicit, separate maximum limit on the number of flow monitoring messages that are generated for NAT error events is implemented. You can control the maximum number of NAT error events for which logs in flow monitoring format must be recorded by including the `message-rate-limit messages-per-second` option at the `[edit interfaces interface-name services-options jflow-log]` hierarchy level. These records for NAT error events are generated when addresses for allocation from the NAT pool are not available, when ports for allocation to a subscriber are not available, or when the allocated quota is exceeded for NAT events (more than the configured number of ports is requested). Also, you can configure the `message-rate-limit` option that previously existed at the `[edit interfaces interface-name services-options syslog]` hierarchy level to specify the maximum number of system log messages per second that can be sent from the PIC to either the Routing Engine (local) or to an external server (remote).
- NAT error events such as “Out of Ports”, “Out of Addresses” and “Quota Exceeded” are rate limited. Default rate limit is 10,000 events per second. This setting is also configurable at PIC level.
- The template for NAT event logging is in accordance with IETF as *IPFIX Information Elements for logging NAT Events—draft-ietf-behave-ipfix-nat-logging-02*.
- Only UDP-based logging is supported, which is an unreliable protocol.
- This functionality is supported on MX Series routers with Junos OS Extension-Provider packages installed and configured on the device, and on MS-MPCs, MS-PICs, and MX-SPC3s. It is not supported on MS-DPCs with MX Series routers.
- Transmission of logs occurs in clear-text format similar to other log messages that the services PICs do not encrypt. It is assumed that the transport of logs and the positioning of the log collector are within a secured realm. Because the messages do not contain sensitive details such as username or passwords, the messages do not cause any security or reliability risks.
- Template IDs 0 through 255 are reserved for template sets and the maximum number of templates supported for logging events in flow monitoring format is 255. When you modify a template-profile configuration (changes to the collector or version, or a deactivation and activation of the service set associated with the template), the specific template is deregistered and reregistered. However, the flow monitoring infrastructure requires 10 minutes by default as the delay period for the template IDs to be freed up. As a result, if you modify the template-profile settings many times within the 10-minute period, the maximum limit on the template IDs of 255 is exceeded and further templates are not registered.
In such a case, when templates are not being registered, you must wait until the delay period for deleting a deregistered template of 10 minutes before you perform any more configuration changes

to have templates registered with the flow monitoring application. To examine whether a template has been registered, you can use the `show services service-sets statistics jflow-log` command. If the Sent field displays a non-zero value for template records, it denotes that templates are successfully registered.

- In a scenario in which the capability to log NAT events in flow monitoring format is enabled at the service-set level, and if the PIC boots up, the flow monitoring log templates are registered with the flow monitoring application. During the registration process, a first set of 12 template records are sent to the collector. However, all of the template records might not reach the collector from the PIC on the router or might not be transmitted out of the router because the interface might not be up from the perspective of the Packet Forwarding Engine. After the refresh time of a template expires, next set of template records are sent out to the collector. For example, if the template refresh time is 60 seconds, only after 60 seconds from the time of booting of the PIC, template records are properly sent to the collector.
- If no problems occur in the transmission of flow monitoring log messages to the collector from the PIC, the Sent field is incremented to indicate NAT events being logged for every event. Also, the tcpdump utility at the destination IP address of the collector denotes the reception of UDP packets. If NAT processing occurs and the value in the Dropped section of the output of the `show services service-sets statistics jflow-log service-set service-set-name` command is incremented or not incremented, you must examine the debugging statistics and counters to determine if any problems exist in the network for transmission of the flow monitoring log messages.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

Exporting Syslog Messages to an External Host Without Flow Monitoring Formats Using an MX Series Router or NFX250

Until Junos OS Release 14.2R1, the only mechanism you can use to generate logs for NAT sessions was by enabling system logging for service sets and transferring syslog messages to either the internal local host on the Routing Engine or to an external host server. When a syslog is enabled with the class or

component being NAT logs and session logs configured, NAT events are recorded. A sample of one such syslog output is as follows:

```
{service_set_3}[jservices-nat]: JSERVICES_NAT_RULE_MATCH: proto 17(UDP) app: any, xe-3/1/1.0#012
192.0.2.2/18575 -> 23.0.0.2/63,Match NAT rule-set (null) rule nat-basic_1
term t1
{service_set_3}MSVCS_LOG_SESSION_OPEN: App:none, xe-3/1/1.0#012 24.0.0.2:18575
[198.51.100.17:1048] -> 23.0.0.2:63 (UDP)
{service_set_3}MSVCS_LOG_SESSION_CLOSE: App:none, xe-3/1/1.0#012 24.0.0.2:18575
[198.51.100.17:1048] -> 23.0.0.2:63 (UDP)
```

From the preceding syslog output, it denotes that NAT create log (NAT translation) and delete log (NAT release) are generated during session events as a part of session-logs configuration. Another important log that is NAT pool exhaustion (not illustrated in the preceding example) is generated as a part of NAT-logs configuration. Such an event message might be caused by Address pooling paired (APP), endpoint-independent mapping (EIM), or address and port exhaustion.

Exporting Version 9 Flow Data Records to a Log Collector Overview Using an MX Series Router or NFX250

A flow record template defines a collection of fields with corresponding descriptions of the format and syntax for the elements or attributes that are contained in it. Network elements (such as routers and switches), which are called exports, accumulate the flow data and export the information to collectors, which are hosts or external devices that can save a large volume of such system log messages for events or system operations. The collected data provides granular, finer-level metering and statistical data for highly flexible and detailed resource usage accounting. Templates that are sent to the collector contain the structural information about the exported flow record fields; therefore, if the collector cannot interpret the formats of the new fields, it can still process the flow record.

The version 9 flow template has a predefined format. An export packet consists of a packet header followed by one or more FlowSet fields. The FlowSet fields can be any of the possible three types—Template, Data, or Options Template. The template flowset describes the fields that are in the data flowsets (or flow records). Each data flowset contains the values or statistics of one or more flows with the same template ID. An interleaved NetFlow version 9 export packet contains the packet header, Template FlowSet, and Data FlowSet fields. A Template FlowSet field signifies each event such as the creation of a NAT entry or the release of a NAT entry allocated, and the Data FlowSet field denotes the NAT sessions for which the Template FlowSet (or the event type) is associated. For example, if a NAT address entry creation, exhaustion of addresses in a NAT pool, and a NAT entry deletion or release occur, an interleaved version 9 export packet contains the packet header, one Template FlowSet field for

NAT address creation, two Data FlowSet fields for the two sessions for which address creation is performed, another TemplateSet field for NAT address deletion, two Data FlowSet fields for the two sessions for which address deletion event occurs, and the other TemplateSet field for NAT pool consumption having exceeded the configured number of pools.

The following are the possible combinations that can occur in an export packet:

- An export packet that consists of interleaved template and data FlowSets—A collector device should not assume that the template IDs defined in such a packet have any specific relationship to the data FlowSets within the same packet. The collector must always cache any received templates, and examine the template cache to determine the appropriate template ID to interpret a data record.
- An export packet consisting entirely of data FlowSets—After the appropriate template IDs have been defined and transmitted to the collector device, most of the export packets consist solely of data FlowSets.
- An export packet consisting entirely of template FlowSets—Although this case is the exception, it is possible to receive packets containing only template records. Ordinarily, templates are appended to data FlowSets. However, in some instances only templates are sent. When a router first boots up or reboots, it attempts to synchronize with the collector device as quickly as possible. The router can send template FlowSets at an accelerated rate so that the collector device has sufficient information to parse any subsequent data FlowSets. Also, template records have a limited lifetime, and they must be periodically refreshed. If the refresh interval for a template occurs and no appropriate data FlowSet that needs to be sent to the collector device is present, an export packet consisting only of template FlowSets is sent.

Understanding Exporting IPFIX Flow Data Records to a Log Collector Using an MX Series Router or NFX250

The IPFIX protocol enables you to access IP flow information on MX-Series Routers or an NFX250. The IPFIX collection process receives the flow information traversing through multiple network elements within the data network in a consistent, identical manner of representation and communication of traffic flows from the network elements to the collection point. An IPFIX device hosts at least one exporting process, which transmits flow records to collecting processes. A collector is a device that performs the collecting processes and an exporter is a device that performs the transfer to data to a collector. An IPFIX message consists of a message header followed by one or more Sets. The Sets can be any of the possible three types: Data Set, Template Set, or Options Template Set. Flow monitoring version 10 (IPFIX) message formats are very similar to version 9 message patterns.

The message header contains the following fields:

- **Version**—Version of the flow record format exported in this message. The value of this field is 0x000a.
- **Length**—Total length of the IPFIX message, measured in octets, including the header and Sets fields.
- **Export Time**—Time, in seconds, since midnight Coordinated Universal Time (UTC) of January 1, 1970, at which the IPFIX message header leaves the exporter. **Sequence Number**—Incremental sequence counter with a value of 2^{32} (2 raised to the power of 32) of all IPFIX data records sent from the current Observation Domain by the exporting process. Template and Options Template records do not increase the Sequence Number attribute.
- **Observation Domain ID**—A 32-bit identifier of the Observation Domain that is locally unique to the exporter.

One of the essential elements in the IPFIX record format is the Template FlowSet record. Templates vastly enhance the flexibility of the Flow Record format because they allow the collector to process Flow Records without necessarily knowing the interpretation of all the data in the Flow Record. A Template Record contains any combination of Internet Assigned Numbers Authority (IANA)-assigned and/or enterprise-specific information element identifiers.

The format of the Template Record signifies a template record header and one or more Field Specifier attributes. The Template FlowSet record contains the following fields:

- **Enterprise bit**—This is the first bit of the Field Specifier. If this bit is zero, the Information Element Identifier identifies an IETF-specified Information Element, and the four-octet Enterprise Number field must not be present. If this bit is one, the Information Element identifier identifies an enterprise-specific Information Element, and the Enterprise Number field must be present.
- **Information Element identifier**—An Information Element is a protocol and encoding-independent description of an attribute that can appear in an IPFIX Record. It is a numeric value that represents the type of Information Element.
- **Field Length**—Length of the corresponding encoded Information Element, in octets. The value 65535 is reserved for variable-length Information Elements.
- **Enterprise Number**—IANA enterprise number of the authority defining the Information Element identifier in this Template Record.

The Data Records are sent in Data Sets. The Data Record field consists only of a Set Header and one or more Field Values. The Template ID to which the Field Values belong is encoded in the Set Header field "Set ID" ("Set ID" = "Template ID"). Interpretation of the Data Record format can be done only if the Template Record corresponding to the Template ID is available at the collecting procedure. Field Values do not necessarily have a length of 16 bits and are encoded according to their data type specified.

Mapping Between Field Values for Version 9 Flow Templates and Logs Exported From an MX-Series Router or NFX250

The following table describes different field IDs or values for flow monitoring logs generated for NAT events in version 9 flow record formats and the events that correspond to the field values:

Field ID	Name	Size (Bytes)	Description
8	ipv4 src address	4	IPv4 source address
225	natInsideGlobalAddress	4	It reports a modified value caused by a NAT middlebox (forwarding class and loss priority) represents function after the packet passed the Observation Point.
12	ipv4 destination address	4	IPv4 destination address
226	natOutsideGlobalAddress	4	It reports a modified value caused by a NAT middlebox function after the packet passed the Observation Point.
7	transport source-port	2	TCP/UDP source port
227	postNAPTSourceTransportPort	2	It reports a modified value caused by a Network Address Port Translation (NAPT) middlebox function after the packet passed the Observation Point.
11	transport destination-port	2	TCP/UDP destination port
228	postNAPTDestinationTransportPort	2	It reports a modified value caused by a Network Address Port Translation (NAPT) middlebox function after the packet passed the Observation Point.
234	ingressVRFID	4	Unique identifier of the VRF name where the packets of this flow are being received. This identifier is unique per Metering Process.

(Continued)

Field ID	Name	Size (Bytes)	Description
235	egressVRFID	4	Unique identifier of the VRF name where the packets of this flow are being sent. This identifier is unique per Metering Process.
4	Ip protocol	1	IP protocol byte
229	natOriginatingAddressRealm	1	Indicates whether the session was created because traffic originated in the private or public address realm. postNATSourceIPv4Address, postNATDestinationIPv4Address, postNAPTSourceTransportPort, and postNAPTDestinationTransportPort are qualified with the address realm in perspective. The allowed values are: Private: 1 Public: 2
230	natEvent	1	Indicates a NAT event. The allowed values are: 1 - Create event. 2 - Delete event. 3 - Pool exhausted. A Create event is generated when a NAT translation is created, whether dynamically or statically. A Delete event is generated when a NAT translation is deleted.
1	inBytes	N	Incoming counter with length N x 8 bits for the number of bytes associated with an IP Flow. By default N is 4
2	inPkts	N	Incoming counter with length N x 8 bits for the number of packets associated with an IP Flow. By default N is 4

(Continued)

Field ID	Name	Size (Bytes)	Description
323	observationTimeMilliseconds	8	Specifies the absolute time in milliseconds of an observation that represents a time value in units of milliseconds based on coordinated universal time (UTC). The choice of an epoch, for example, 00:00 UTC, January 1, 1970, is left to corresponding encoding specifications for this type. Leap seconds are excluded. Note that transformation of values might be required between different encodings if different epoch values are used.
27	sourceIPv6Address	16	IPv6 source address
284	natPoolName	64	NAT resource pool name
361	portRangeStart	2	The port number identifying the start of a range of ports. A value of zero indicates that the range start is not specified, ie the range is defined in some other way.
362	portRangeEnd	2	The port number identifying the end of a range of ports. A value of zero indicates that the range end is not specified, and the range is defined in some other way.
363	portRangeStepSize	2	The step size in a port range. The default step size is 1, which indicates contiguous ports. A value of zero indicates that the step size is not specified, and the range is defined in some other way.
364	portRangeNumPorts	2	The number of ports in a port range. A value of zero indicates that the number of ports is not specified, and the range is defined in some other way.

Consider a sample scenario of a NAT address creation event. Based on the fields in the preceding table, for translations that are not available (such as natOutsideGlobalAddress) is set to 0. Ingress and Egress VRF of the flow can be made available. Also, natEvent is equal to 1 (create). The inBytes field is assumed

to be 0 or number of bytes of the incoming packet and the inPkts field is either 0 or 1 because it is the first packet into the system when translation happens. The observationTimeMilliseconds field denotes the time when this address translation creation is recorded.

For a NAT address deletion event, for translations that are not available (such as natOutsideGlobalAddress) is set to 0. Ingress and Egress VRF of the flow can be made available. Also, natEvent is equal to 2 (create). The inBytes field denotes the number of bytes for this flow in both the forward or upward, the value of the inPkts field denotes the number of packets for this flow in both the upward and backward directions. observationTimeMilliseconds is the time when this deletion of translation is recorded.

When the NAT pool is exhausted and no further addresses are remaining for allocation, for translations that are not available (such as natOutsideGlobalAddress) is set to 0. Ingress and Egress VRF of the flow can be made available. Also, the natEvent field is set to 3 (Pool exhausted). All resource failures are combined as a single event. The inBytes field is assumed to be 0 or number of bytes of the incoming packet and the inPkts field is either 0 or 1 because it is the first packet into the system when translation happens. The value of the observationTimeMilliseconds field is the time when this failed translation is recorded.

Mapping Between Field Values for IPFIX Flow Templates and Logs Exported From an MX Series Router or NFX250

An IETF draft defining IPFIX Information Elements for logging various NAT events is available in IETF as *IPFIX Information Elements for logging NAT Events—draft-ietf-behave-ipfix-nat-logging-02*. The flow monitoring template format for flow monitoring logs generated for NAT events comply with the templates defined in this draft for logging NAT44/NAT64 session create/delete, binding information base (BIB) create/delete, address exhaust, pool exhaustion, quota exceeded, address binding create/delete, port block allocation and de-allocation events. Also, this draft has an extension for NAT64. Support is implemented for logging events for both NAT44 and NAT64. Apart from those templates defined in this draft, no new user-defined templates are created for logging any NAT events.

The following table lists the extensions to the NAT events. The data record contains the corresponding natEvent value to identify the event that is being logged.

Event Name	Values
NAT44 Session create	1
NAT44 Session delete	2

(Continued)

Event Name	Values
NAT Addresses exhausted	3
NAT64 Session create	4
NAT64 Session delete	5
NAT44 BIB create	6
NAT44 BIB delete	7
NAT64 BIB create	8
NAT64 BIB delete	9
NAT ports exhausted	10
Quota exceeded	11
Address binding create	12
Address binding delete	13
Port block allocation	14
Port block deallocation	15

The following table describes the field IDs or values and the corresponding names for IPv6 addresses for IPFIX flows:

Field ID	Name	Size (Bytes)	Description
27	sourceIPv6Address	16	IPv6 source address
28	destinationIPv6Address	16	IPv6 destination address
281	postNATSourceIPv6Address	16	Translated source IPv6 address
282	postNATDestinationPv6Address	16	Translated destination IPv6 address

The following table describes the field names and whether they are required or not for NAT64 session creation and deletion events:

Field Name	Size (Bits)	Whether the Field Is Mandatory
timeStamp	64	Yes
vlanID/ingressVRFID	32	No
sourceIPv4Address	128	Yes
postNATSourceIPv4Address	32	Yes
protocolIdentifier	8	Yes
sourceTransportPort	16	Yes
postNAPTsourceTransportPort	16	Yes
destinationIPv4Address	128	No
postNATDestinationIPv4Address	32	No
destinationTransportPort	16	No

(Continued)

Field Name	Size (Bits)	Whether the Field Is Mandatory
postNAPTdestinationTransportPort	16	No
natOriginatingAddressRealm	8	No
initiatorOctets	64	No
responderOctets	64	No
flowEndReason	8	No
natEvent	8	Yes

A NAT44 session creation template record can contain the following fields. The natEvent field contains a value of 1, which indicates a NAT44 session creation event. An example of such a template is as follows:

Field Name	Size (Bits)	Value
timeStamp	64	09:20:10:789
sourceIPv4Address	32	192.168.16.1
postNATSourceIPv4Address	32	192.0.2.100
protocolIdentifier	8	TC
sourceTransportPort	16	14800
postNAPTsourceTransportPort	16	1024
destinationIPv4Address	32	198.51.100.104

(Continued)

Field Name	Size (Bits)	Value
postNATDestinationIPv4Address	32	198.51.100.104
destinationTransportPort	16	80
postNAPTdestinationTransportPort	16	80
natOriginatingAddressRealm	8	0
initiatorOctets	64	No
responderOctets	64	No
flowEndReason	8	No
natEvent	8	1

A NAT44 session deletion template record can contain the following fields. The natEvent field contains a value of 2, which indicates a NAT44 session deletion event. An example of such a template is as follows:

Field Name	Size (Bits)	Value
timeStamp	64	09:20:10:789
sourceIPv4Address	32	192.168.16.1
postNATSourceIPv4Address	32	192.0.2.100
protocolIdentifier	8	TC
sourceTransportPort	16	14800

(Continued)

Field Name	Size (Bits)	Value
postNAPTsourceTransportPort	16	1024
destinationIPv4Address	32	198.51.100.104
postNATDestinationIPv4Address	32	198.51.100.104
destinationTransportPort	16	80
postNAPTdestinationTransportPort	16	80
natOriginatingAddressRealm	8	0
natEvent	8	2

To support all session termination reasons on NAT, existing `flowEndReason` information element is extended. A new CLI command `session-end-reason` is introduced to configure `flowEndReason` to be a part of J-Flow IPFIX template.

If the CLI is not configured or configured as default, the `flowEndReason` exports the default set information to fill in the data records. If the CLI is configured as custom, the `flowEndReason` exports the custom set information to fill in the data records.

The table lists the set of session termination values that can be exported:

Table 48: Session Termination Values

Session Close Reason	Session Close Reason string	Scenarios/Remark	Custom Set values	Default Set values
NAT_SESSION_CREATION	idle Timeout	When any session gets timeout	0x01	0x01

Table 48: Session Termination Values *(Continued)*

Session Close Reason	Session Close Reason string	Scenarios/Remark	Custom Set values	Default Set values
NAT_SESSION_CLOSE_TCP_CLIENT_RST	TCP CLIENT RST	Receives a TCP packet from Client with RST FLAG set	0x13	0xFF
NAT_SESSION_CLOSE_TCP_SERVER_RST	TCP SERVER RST	Receives a TCP packet from Server with RST FLAG set	0x23	0xFF
NAT_SESSION_CLOSE_TCP_FIN	TCP FIN	Receives FIN Packet	0x03	0x03
NAT_SESSION_CLOSE_ICMP_ERR	ICMP Error	Receiving ICMP Error packet in Fast path. icmp related error messages mentioned below	0x10	0xFF
NAT_SESSION_CLOSE_NSRRP	HA	<p>Create a NAT session on active router. Now, Switch to backup Router Manually or by bringing down the pic on active router.</p> <p>Wait for the switchover and send traffic. Ensure the session is synchronized.</p> <p>Now close the session.</p>	0x20	0xFF

Table 48: Session Termination Values *(Continued)*

Session Close Reason	Session Close Reason string	Scenarios/Remark	Custom Set values	Default Set values
NAT_SESSION_CLOSE_POLICY_DELETE	policy delete	When you delete Policy rematch configuration with active session.	0x50	0xFF
NAT_SESSION_CLOSE_POLICY_UPDATE	policy update	When you Update Policy rematch configuration with active session.	0x60	0xFF
NAT_SESSION_CLOSE_JSF_PLUGIN	application failure or action	It is a very rare scenario and would be difficult to simulate. Please don't have test case for this.	0x70	0xFF
NAT_SESSION_CLOSE_IFP_ZONECHANGED_SSCAN	session interface zone changed	when redundancy switchover happens in ams interface	0x80	0xFF
NAT_SESSION_CLOSE_CLI	CLI	Force clear the session	0x04	0x04

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats

You can configure MX Series routers with MS-MPCs, MS-MICs, and MX-SPC3s to log network address translation (NAT) events using the Junos Traffic Vision (previously known as Jflow) version 9 or IPFIX (version 10) template format. NAT event logger generates logs or template records in flow monitoring format and transmits them to the specified external collector or server for various NAT events, such as NAT44 and NAT64 session creation and deletion, and NAT44 and NAT64 binding information base events.

NOTE: This functionality is supported on MX Series routers with Junos OS Extension-Provider packages installed and configured on the device, and on MS-MPCs, MS-PICs, and MX-SPC3s. It is not supported on MS-DPCs with MX Series routers.

You can configure the mechanism to record logging messages in flow monitoring format for NAT events. You need to define collectors, and template profiles that contain the properties for flow monitoring logs. You can create a template profile for a particular NAT service on an MX Series router with MS-MPCs, MS-MICs, or MX-SPC3s, or for a service set, which applies for all of the NAT services. You can define a template profile to generate flow monitoring logs in a specific flow template format and associate the specified template profile with a service set.

To enable the flow monitoring log capability for NAT events and configure the transmission of logs to collectors at a service level:

1. Define the flow monitoring log service to be applied on an interface to control the maximum number of flow monitoring logs generated for NAT error events.

```
[edit]
user@host# set interfaces ms-fpc/pic/port services-options jflow-log message-rate-limit
messages-per-second
```

For example:

```
[edit]
user@host# set interfaces ms-5/0/0 services-options jflow-log message-rate-limit 50
```

2. Configure the collectors and collector groups.

```
[edit]
user@host# set services jflow-log collector collector-name destination-address address
destination-port port-number source-ip address
user@host# set services jflow-log collector-group collector-group-name collector [ collector-
name1 collector-name2]
```

For example:

```
[edit]
user@host# set services jflow-log collector c1 destination-address 203.0.113.3 destination-
```

```
port 1 source-ip 192.0.2.1
user@host# set services jflow-log collector-group cg1 collector c1
```

3. Configure the template profiles and associate the template profile with the collector or collector group.

```
[edit]
user@host# set services jflow-log template-profile template-profile-name collector collector-name version (ipfix | v9) template-type nat refresh-rate packets packets seconds seconds
user@host# set services jflow-log template-profile template-profile-name collector-group collector-group-name version (ipfix | v9) template-type nat refresh-rate packets packets seconds seconds
```

For example:

```
[edit]
user@host# set services jflow-log template-profile t1 collector c1 version ipfix template-type nat refresh-rate packets 20 seconds 20
user@host# set services jflow-log template-profile t1 collector-group cg1
user@host# set services jflow-log template-profile t2 collector c2 version v9 template-type nat refresh-rate packets 20 seconds 20
```

4. Associate the template profile with the service set.

```
[edit]
user@host# set services service-set service-set-name jflow-log template-profile template-profile-name
```

For example:

```
[edit]
user@host# set services service-set sset_0 jflow-log template-profile t1
```

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting

IN THIS SECTION

- Requirements | 275
- Generation of Log Messages Using Flow Templates for NAT Operations on MS-MPCs, MS-MICs, and MX-SPC3s | 276
- Configuration | 276
- Verification | 280

You can configure MX Series routers with MS-MPCs, MS-MICs, and MX-SPC3s to log network address translation (NAT) events using the Junos Traffic Vision (previously known as Jflow) version 9 or IPFIX (version 10) template format. This method of generating flow monitoring records for NAT events, such as NAT44 and NAT64 session creation and deletion, and NAT44 and NAT64 binding information base events, enables cohesive and streamlined analysis of NAT traffic and troubleshooting of NAT-related problems.

NOTE: This functionality is supported on MX Series routers with Junos OS Extension-Provider packages installed and configured on the device, and on MS-MPCs, MS-PICs, and MX-SPC3s. It is not supported on MS-DPCs with MX Series routers.

This example describes how to configure flow monitoring log generation in flow monitoring format for NAT events at the service-set level on MS-MIC, MS-MPC, and MX-SPC3, and contains the following sections:

NOTE: This configuration example is for an Interface-Style service set.

Requirements

This example uses the following hardware and software components:

- One MX Series router with an MS-MPC, MS-MIC, or MX-SPC3
- Junos OS Release 14.2R2 or later for MX Series routers

Generation of Log Messages Using Flow Templates for NAT Operations on MS-MPCs, MS-MICs, and MX-SPC3s

You can configure the mechanism to record logging messages in flow monitoring format for NAT events. You can create a template profile for a particular NAT service on an MX Series router with MS-MPCs, MS-MICs, or MX-SPC3s, or for a service set, which applies for all of the NAT services. You must define a template profile to generate flow monitoring logs in a specific flow template format and attach the template profile with a service set. You must configure a collector or a group of collectors, which are hosts that receive the log messages for NAT events from the service PIC or the exporter. You need to associate a template profile with the collector. The profile defines the characteristics of the flow monitoring record template, such as the version of flow monitoring (version 9 or IPFIX), the refresh rate, in either packets or seconds, and the type of service or application (NAT in this case) for which flow records must be sent to the collector.

Assume a sample deployment in which two collectors, c1 and c2, are defined. These collectors are clustered into two groups. The collector group, cg1, contains c1 and c2, and the collector group, cg2, contains c2. Two template profiles named t1 and t2 are defined. The profiles, t1 and t2, are associated with collectors, c1 and c2, respectively.

These profiles describe the properties or attributes for transmission of logs, such as the flow template format to be used, the rate at which the logs must be refreshed, and the service or event, such as NAT, for which logs must be sent to the specified collector.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 277](#)
- [Procedure | 277](#)
- [Results | 279](#)

To enable the flow monitoring log capability for NAT events and configure the transmission of logs to collectors, perform these tasks:

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configuring Service Set Properties

```
set services service-set sset_0 interface-service service-interface ms-5/0/0.0
```

Applying Flow Monitoring Log Service on an Interface

```
set interfaces ms-5/0/0 services-options jflow-log message-rate-limit 50000
```

Enabling and Configuring Flow Monitoring Logs for a Service Set

```
set services jflow-log collector c1 destination-address 192.0.2.3 destination-port 1 source-ip 198.51.100.1
set services jflow-log collector c2 destination-address 203.0.113.5 destination-port 3 source-ip 198.51.100.2
set services jflow-log collector-group cg1 collector [ c1 c2 ]
set services jflow-log template-profile t1 collector c1 version ipfix template-type nat refresh-rate packets 20 seconds 20
set services jflow-log template-profile t2 collector c2 version v9 template-type nat refresh-rate packets 20 seconds 20
set services jflow-log template-profile t1 collector-group cg1
```

Associating the Template Profile with a Service Set

```
set services service-set sset_0 jflow-log template-profile t1
```

Procedure

Step-by-Step Procedure

To configure the generation and transmission of flow monitoring template logs for NAT events:

1. Create a service set properties.

```
[edit]
user@host# set services service-set sset_0 interface-service service-interface ms-5/0/0.0
```

2. Define the flow monitoring log service to be applied on an interface.

```
[edit]
user@host# set interfaces ms-5/0/0 services-options jflow-log message-rate-limit 50000
```

3. Configure the collectors and collector groups.

```
[edit]
user@host# set services jflow-log collector c1 destination-address 192.0.2.3 destination-port
1 source-ip 198.51.100.1
user@host# set services jflow-log collector c2 destination-address 203.0.113.5 destination-
port 3 source-ip 198.51.100.2
user@host# set services jflow-log collector-group cg1 collector [ c1 c2 ]
user@host# set services jflow-log collector-group cg2 collector c2
```

4. Configure the template profiles and associate the template profile with the collector.

```
[edit]
user@host# set services jflow-log template-profile t1 collector c1 version ipfix template-
type nat refresh-rate packets 20 seconds 20
user@host# set services jflow-log template-profile t2 collector c2 version v9 template-type
nat refresh-rate packets 20 seconds 20
```

5. Associate the template profile with the service set.

```
[edit]
user @ host# set services service-set sset_0 jflow-log template-profile t1
```

Results

From the configuration mode, confirm your configuration by entering the `show services`, `show services jflow-log`, and `show services service-set sset_0 jflow-log` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@host# show services
service-set sset_0 {
    interface-service {
        service-interface ms-5/0/0;
    }
}
[edit interfaces]
ms-5/0/0 {
    services-options {
        jflow-log {
            message-rate-limit 50000;
        }
    }
}

user@host# show services jflow-log
collector c1 {
    destination-address 192.0.2.3;
    destination-port 1;
    source-ip 198.51.100.1;
}
collector c2 {
    destination-address 203.0.113.5;
    destination-port 3;
    source-ip 198.51.100.2;
}
collector-group cg1 {
    collector [ c2 c1 ];
}
collector-group cg2 {
    collector c2;
}
template-profile t2 {
    collector c2;
    template-type nat;
    refresh-rate packets 20 seconds 20;
}

```

```

    version v9;
}
template-profile t1 {
    collector c1;
    template-type nat;
    refresh-rate packets 20 seconds 20;
    version ipfix;
}

[edit]
user@host# show services service-set sset_0 jflow-log
template-profile t2;

```

Verification

IN THIS SECTION

- [Verifying That the Flow Monitoring Logs Are Generated and Sent to Collectors | 280](#)

To confirm that the configuration is working properly, perform the following:

Verifying That the Flow Monitoring Logs Are Generated and Sent to Collectors

Purpose

Verify that the flow monitoring log messages in the defined template format, such as IPFIX or version 9, are generated and transmitted to the configured collectors for the different NAT operations.

Action

From operational mode, use the `show services service-sets statistics jflow-log` command:

```

user@host> show services service-sets statistics jflow-log
Interface: ms-5/0/0
Rate limit: 1000
Template records:
Sent: 36

```

```

Dropped: 0
Data records:
  Sent: 2
  Dropped: 0

Service-set: sset_0
  Unresolvable collectors: 0
  Template records:
    Sent: 36
    Dropped: 0
  Data records:
    Sent: 2
    Dropped: 0

```

From operational mode, use the `show services service-sets statistics jflow-log detail` command:

```
user@host> show services service-sets statistics jflow-log detail
```

```

Interface: ms-5/0/0
Rate limit: 1000
Template records:
  Sent: 48
  Dropped: 0
Data records:
  Sent: 4
  Dropped: 0

Service-set: sset_0
  Unresolvable collectors: 0
  Template records:
    Sent: 48
    Dropped: 0
  Data records:
    Sent: 4
    Dropped: 0
  NAT44 Session logs:
    Template records:
      Sent: 4
      Dropped: 0 (socket send error: 0, no memory: 0)

```

```

Data records:
  Sent: 4
  Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 Session logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 BIB logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 BIB logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Address Exhausted logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Port Exhausted logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Quota Exceeded logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:

```

```

    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 Quota Exceeded logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Address Bind logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 Address Bind logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 PBA logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 PBA logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)

```

Meaning

The output shows that the log messages in flow monitoring format associated with the specified service set and interface are generated for the different NAT events.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

2

PART

Flow Capture Services

[Dynamically Capturing Packet Flows Using Junos Capture Vision](#) | 286

[Detecting Threats and Intercepting Flows Using Junos Packet Vision](#) | 302

[Using Flow-Tap to Monitor Packet Flow](#) | 321

Dynamically Capturing Packet Flows Using Junos Capture Vision

IN THIS CHAPTER

- [Understanding Junos Capture Vision | 286](#)
- [Configuring Junos Capture Vision | 289](#)
- [Example: Configuring Junos Capture Vision on M and T Series Routers | 297](#)
- [Monitoring a Capture Group Using SNMP or Show Services Commands | 301](#)

Understanding Junos Capture Vision

IN THIS SECTION

- [Junos Capture Vision Architecture | 286](#)
- [Liberal Sequence Windowing | 288](#)
- [Intercepting IPv6 Flows | 288](#)

Junos Capture Vision (known as dynamic flow capture in Junos OS Releases earlier than 13.2) enables you to capture packet flows on the basis of dynamic filtering criteria. Specifically, you can use this feature to forward passively monitored packet flows that match a particular filter list to one or more destinations using an on-demand control protocol.

Junos Capture Vision Architecture

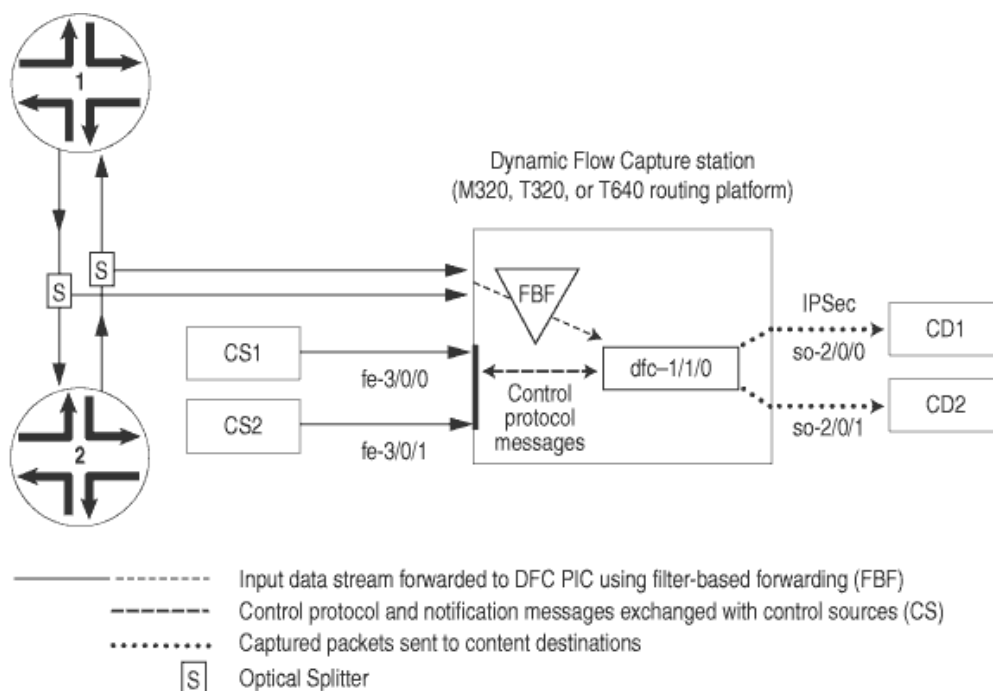
The architecture consists of one or more *control sources* that send requests to a Juniper Networks router to monitor incoming data, and then forward any packets that match specific filter criteria to a set of one or more *content destinations*. The architectural components are defined as follows:

- Control source—A client that monitors electronic data or voice transfer over the network. The control source sends filter requests to the Juniper Networks router using the Dynamic Task Control Protocol (DTCP), specified in draft-cavuto-dtcp-03.txt at <http://www.ietf.org/internet-drafts>. The control source is identified by a unique identifier and an optional list of IP addresses.
- Monitoring platform—A T Series or M320 router containing one or more Dynamic Flow Capture (DFC) PICs, which support dynamic flow capture processing. The monitoring platform processes the requests from the control sources, creates the filters, monitors incoming data flows, and sends the matched packets to the appropriate content destinations.
- Content destination—Recipient of the matched packets from the monitoring platform. Typically the matched packets are sent using an IP Security (IPsec) tunnel from the monitoring platform to another router connected to the content destination. The content destination and the control source can be physically located on the same host. For more information on IPsec tunnels, see *Understanding Junos VPN Site Secure*.

NOTE: The Junos Capture Vision PIC (either a Monitoring Services III PIC or Multiservices 400 PIC) forwards the entire packet content to the content destination, rather than to a content record as is done with cflowd or flow aggregation version 9 templates.

Figure 25 on page 288 shows a sample topology. The number of control sources and content destinations is arbitrary.

Figure 25: Junos Capture Vision Topology



g017075

Liberal Sequence Windowing

Each DTCP packet (add, delete, list, and refresh packets) contains a 64-bit sequence number to identify the order of the packets. Because the network is connectionless, the DTCP packets can arrive out of order to the router running the Junos Capture Vision application.

The *liberal sequence window* feature implements a negative window for the sequence numbers received in the DTCP packets. It enables the Junos Capture Vision application to accept not only DTCP packets with sequence numbers greater than those previously received, but also DTCP packets with lesser sequence numbers, up to a certain limit. This limit is the negative window size; the positive and negative window sizes are +256 and -256 respectively, relative to the current maximum sequence number received. No configuration is required to activate this feature; the window sizes are hard-coded and nonconfigurable.

Intercepting IPv6 Flows

Starting with Junos OS Release 11.4, Junos Capture Vision also supports intercepting IPv6 flows in M320, T320, T640, and T1600 routers with a Multiservices 400 or Multiservices 500 PIC. Junos

Capture Vision can intercept passively monitored IPv6 traffic only. All support for IPv4 interception remains the same. The interception of IPv6 traffic happens in the same way the filters capture IPv4 flows. With the introduction of IPv6 interception, both IPv4 and IPv6 filters can coexist. The mediation device, however, cannot be located in an IPv6 network.

Junos Capture Vision does not support interception of VPLS and MPLS traffic. The application cannot intercept Address Resolution Protocol (ARP) or other Layer 2 exception packets. The interception filter can be configured to timeout based on factors like total time (seconds), idle time (seconds), total packets or total data transmitted (bytes).

RELATED DOCUMENTATION

[Configuring Junos Capture Vision | 289](#)

[Example: Configuring Junos Capture Vision on M and T Series Routers | 297](#)

Configuring Junos Capture Vision

IN THIS SECTION

- [Configuring the Capture Group | 289](#)
- [Configuring the Content Destination | 291](#)
- [Configuring the Control Source | 292](#)
- [Configuring the DFC PIC Interface | 293](#)
- [Configuring the Firewall Filter | 294](#)
- [Configuring System Logging | 295](#)
- [Configuring Tracing Options for Junos Capture Vision Events | 295](#)
- [Configuring Thresholds | 296](#)
- [Limiting the Number of Duplicates of a Packet | 296](#)

Configuring the Capture Group

A capture group defines a profile of Junos Capture Vision configuration information. The static configuration includes information about control sources, content destinations, and notification

destinations. Dynamic configuration is added through interaction with control sources using a control protocol.

To configure a capture group, include the `capture-group` statement at the `[edit services dynamic-flow-capture]` hierarchy level:

```
capture-group client-name {
  content-destination identifier {
    address address;
    hard-limit bandwidth;
    hard-limit-target bandwidth;
    soft-limit bandwidth;
    soft-limit-clear bandwidth;
    ttl hops;
  }
  control-source identifier {
    allowed-destinations [ destinations ];
    minimum-priority value;
    no-syslog;
    notification-targets address port port-number;
    service-port port-number;
    shared-key value;
    source-addresses [ addresses ];
  }
  duplicates-dropped-periodicity seconds;
  input-packet-rate-threshold rate;
  interfaces interface-name;
  max-duplicates number;
  pic-memory-threshold percentage percentage;
}
```

To specify the capture-group, assign it a unique *client-name* that associates the information with the requesting control sources.

Configuring the Content Destination

You must specify a destination for the packets that match DFC PIC filter criteria. To configure the content destination, include the `content-destination` statement at the `[edit services dynamic-flow-capture capture-group client-name]` hierarchy level:

```
content-destination identifier {
  address address;
  hard-limit bandwidth;
  hard-limit-target bandwidth;
  soft-limit bandwidth;
  soft-limit-clear bandwidth;
  ttl hops;
}
```

Assign the content-destination a unique *identifier*. You must also specify its IP address and you can optionally include additional settings:

- **address**—The DFC PIC interface appends an IP header with this destination address on the matched packet (with its own IP header and contents intact) and sends it out to the content destination.
- **ttl**—The time-to-live (TTL) value for the IP-IP header. By default, the TTL value is 255. Its range is 0 through 255.
- **Congestion thresholds**—You can specify per-content destination bandwidth limits that control the amount of traffic produced by the DFC PIC during periods of congestion. The thresholds are arranged in two pairs: `hard-limit` and `hard-limit-target`, and `soft-limit` and `soft-limit-clear`. You can optionally include one or both of these paired settings. All four settings are 10-second average bandwidth values in bits per second. Typically `soft-limit-clear < soft-limit < hard-limit-target < hard-limit`. When the content bandwidth exceeds the `soft-limit` setting:
 1. A congestion notification message is sent to each control source of the criteria that point to this content destination
 2. If the control source is configured for syslog, a system log message is generated.
 3. A latch is set, indicating that the control sources have been notified. No additional notification messages are sent until the latch is cleared, when the bandwidth falls below the `soft-limit-clear` value.

When the bandwidth exceeds the `hard-limit` value:

1. Junos Capture Vision begins deleting criteria until the bandwidth falls below the `hard-limit-target` value.

2. For each criterion deleted, a CongestionDelete notification is sent to the control source for that criterion.
3. If the control source is configured for syslog, a log message is generated.

The application evaluates criteria for deletion using the following data:

- Priority—Lower priority criteria are purged first, after adjusting for control source minimum priority.
- Bandwidth—Higher bandwidth criteria are purged first.
- Timestamp—The more recent criteria are purged first.

Configuring the Control Source

You configure information about the control source, including allowed source addresses and destinations and authentication key values. To configure the control source information, include the `control-source` statement at the `[edit services dynamic-flow-capture capture-group client-name]` hierarchy level:

```
control-source identifier {
  allowed-destinations [ destination-identifiers ];
  minimum-priority value;
  no-syslog;
  notification-targets address port port-number;
  service-port port-number;
  shared-key value;
  source-addresses [ addresses ];
}
```

Assign the `control-source` statement a unique *identifier*. You can also include values for the following statements:

- `allowed-destinations`—One or more content destination identifiers to which this control source can request that matched data be sent in its control protocol requests. If you do not specify any content destinations, all available destinations are allowed.
- `minimum-priority`—Value assigned to the control source that is added to the priority of the criteria in the DTCP ADD request to determine the total priority for the criteria. The lower the value, the higher the priority. By default, `minimum-priority` has a value of 0 and the allowed range is 0 through 254.
- `notification-targets`—One or more destinations to which the DFC PIC interface can log information about control protocol-related events and other events such as PIC bootup messages. You configure

each notification-target entry with an IP address value and a User Datagram Protocol (UDP) port number.

- **service-port**—UDP port number to which the control protocol requests are directed. Control protocol requests that are not directed to this port are discarded by DFC PIC interfaces.
- **shared-key**—20-byte authentication key value shared between the control source and the DFC PIC monitoring platform.
- **source-addresses**—One or more allowed IP addresses from which the control source can send control protocol requests to the DFC PIC monitoring platform. These are /32 addresses.

Configuring the DFC PIC Interface

You specify the interface that interacts with the control sources configured in the same capture group. A Monitoring Services III PIC can belong to only one capture group, and you can configure only one PIC for each group.

To configure a DFC PIC interface, include the `interfaces` statement at the `[edit services dynamic-flow-capture capture-group client-name]` hierarchy level:

```
interfaces interface-name;
```

You specify DFC interfaces using the `dfc-` identifier at the `[edit interfaces]` hierarchy level. You must specify three logical units on each DFC PIC interface, numbered 0, 1, and 2. You cannot configure any other logical interfaces.

- **unit 0** processes control protocol requests and responses.
- **unit 1** receives monitored data.
- **unit 2** transmits the matched packets to the destination address.

The following example shows the configuration necessary to set up a DFC PIC interface and intercept both IPv4 and IPv6 traffic:

```
[edit interfaces dfc-0/0/0]
unit 0 {
  family inet {
    filter {
      output high; #Firewall filter to route control packets
      # through 'network-control' forwarding class. Control packets
      # are loss sensitive.
    }
  }
}
```



```

        address 10.1.0.0/32 { # DFC PIC address
            destination 10.36.100.1; # DFC PIC address used by
            # the control source to correspond with the
            # monitoring platform
        }
    }
    unit 1 { # receive data packets on this logical interface
        family inet; # receive IPv4 traffic for interception
        family inet6; # receive IPv6 traffic for interception
    }
    unit 2 { # send out copies of matched packets on this logical interface
        family inet;
    }
}

```

In addition, you must configure Junos Capture Vision to run on the DFC PIC in the correct chassis location. The following example shows this configuration at the `[edit chassis]` hierarchy level:

```

fpc 0 {
    pic 0 {
        monitoring-services application dynamic-flow-capture;
    }
}

```

Configuring the Firewall Filter

You can specify the firewall filter to route control packets through the network control forwarding class. The control packets are loss sensitive. To configure the firewall filter, include the following statements at the `[edit]` hierarchy level:

```

firewall {
    family inet {
        filter high {
            term all {
                then forwarding-class network-control;
            }
        }
    }
}

```

Configuring System Logging

By default, control protocol activity is logged as a separate system log facility, `dfc`. To modify the filename or level at which control protocol activity is recorded, include the following statements at the `[edit syslog]` hierarchy level:

```
file dfc.log {
    dfc any;
}
```

To cancel logging, include the `no-syslog` statement at the `[edit services dynamic-flow-capture capture-group client-name control-source identifier]` hierarchy level:

```
no-syslog;
```

NOTE: Junos Capture Vision (`dfc-`) interface supports up to 10,000 filter criteria. When more than 10,000 filters are added to the interface, the filters are accepted, but system log messages are generated indicating that the filter is full.

Configuring Tracing Options for Junos Capture Vision Events

You can enable tracing options for Junos Capture Vision events by including the `traceoptions` statement at the `[edit services dynamic-flow-capture]` hierarchy level.

When you include the `traceoptions` configuration, you can also specify the trace file name, maximum number of trace files, the maximum size of trace files, and whether the trace file can be read by all users or not.

To enable tracing options for Junos Capture Vision events, include the following configuration at the `[edit services dynamic-flow-capture]` hierarchy level:

```
traceoptions{
    file filename <files number> <size size> <world-readable | non-world-readable>;
}
```

To disable tracing for Junos Capture Vision events, delete the `traceoptions` configuration from the `[edit services dynamic-flow-capture]` hierarchy level.

NOTE: In Junos OS releases earlier than 9.2R1, tracing of Junos Capture Vision was enabled by default, and the logs were saved to the `/var/log/dfcd` directory.

Configuring Thresholds

You can optionally specify threshold values for the following situations in which warning messages be recorded in the system log:

- Input packet rate to the DFC PIC interfaces
- Memory usage on the DFC PIC interfaces

To configure threshold values, include the `input-packet-rate-threshold` or `pic-memory-threshold` statements at the `[edit services dynamic-flow-capture capture-group client-name]` hierarchy level:

```
input-packet-rate-threshold rate;
pic-memory-threshold percentage percentage;
```

If these statements are not configured, no threshold messages are logged. The threshold settings are configured for the capture group as a whole.

The range of configurable values for the `input-packet-rate-threshold` statement is 0 through 1 Mpps. The PIC calibrates the value accordingly; the Monitoring Services III PIC caps the threshold value at 300 Kpps and the Multiservices 400 PIC uses the full configured value. The range of values for the `pic-memory-threshold` statement is 0 to 100 percent.

Limiting the Number of Duplicates of a Packet

You can optionally specify the maximum number of duplicate packets the DFC PIC is allowed to generate from a single input packet. This limitation is intended to reduce the load on the PIC when packets are sent to multiple destinations. When the maximum number is reached, the duplicates are sent to the destinations with the highest criteria class priority. Within classes of equal priority, criteria having earlier timestamps are selected first.

To configure this limitation, include the `max-duplicates` statement at the `[edit services dynamic-flow-capture capture-group client-name]` hierarchy level:

```
max-duplicates number;
```

You can also apply the limitation on a global basis for the DFC PIC by including the `g-max-duplicates` statement at the `[edit services dynamic-flow-capture]` hierarchy level:

```
g-max-duplicates number;
```

By default, the maximum number of duplicates is set to 3. The range of allowed values is 1 through 64. A setting for `max-duplicates` for an individual capture-group overrides the global setting.

In addition, you can specify the frequency with which the application sends notifications to the affected control sources that duplicates are being dropped because the threshold has been reached. You configure this setting at the same levels as the maximum duplicates settings, by including the `duplicates-dropped-periodicity` statement at the `[edit services dynamic-flow-capture capture-group client-name]` hierarchy level or the `g-duplicates-dropped-periodicity` statement at the `[edit services dynamic-flow-capture]` hierarchy level:

```
duplicates-dropped-periodicity seconds;
g-duplicates-dropped-periodicity seconds;
```

As with the `g-max-duplicates` statement, the `g-duplicates-dropped-periodicity` statement applies the setting globally for the application and is overridden by a setting applied at the capture-group level. By default, the frequency for sending notifications is 30 seconds.

RELATED DOCUMENTATION

[Understanding Junos Capture Vision | 286](#)

[Example: Configuring Junos Capture Vision on M and T Series Routers | 297](#)

Example: Configuring Junos Capture Vision on M and T Series Routers

The following example includes all parts of a complete Junos Capture Vision configuration.

Configure the Junos Capture Vision PIC interface:

```
[edit interfaces dfc-0/0/0]
unit 0 {
  family inet {
    filter {
      output high; #Firewall filter to route control packets
```

```

        # through 'network-control' forwarding class. Control packets
        # are loss sensitive.
    }
    address 10.1.0.0/32 { # DFC PIC address
        destination 10.36.100.1; # DFC PIC address used by
        # the control source to correspond with the
        # monitoring platform
    }
}
unit 1 { # receive data packets on this logical interface
    family inet;
    family inet6;
}
unit 2 { # send out copies of matched packets on this logical interface
    family inet;
}

```

Configure the capture group:

```

services dynamic-flow-capture {
    capture-group g1 {
        interfaces dfc-0/0/0;
        input-packet-rate-threshold 90k;
        pic-memory-threshold percentage 80;
        control-source cs1 {
            source-addresses 10.36.41.1;
            service-port 2400;
            notification-targets {
                10.36.41.1 port 2100;
            }
            shared-key "$ABC123";
            allowed-destinations cd1;
        }
        content-destination cd1 {
            address 10.36.70.2;
            ttl 244;
        }
    }
}

```

Configure filter-based forwarding (FBF) to the Junos Capture Vision PIC interface, logical unit 1.

For more information about configuring passive monitoring interfaces, see ["Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers"](#) on page 157.

```

interfaces so-1/2/0 {
  encapsulation ppp;
  unit 0 {
    passive-monitor-mode;
    family inet {
      filter {
        input catch;
      }
    }
  }
}

```

Configure the firewall filter:

```

firewall {
  filter catch {
    interface-specific;
    term def {
      then {
        count counter;
        routing-instance fbf_inst;
      }
    }
  }
  family inet {
    filter high {
      term all {
        then forwarding-class network-control;
      }
    }
  }
}

```

Configure a forwarding routing instance. The next hop points specifically to the logical interface corresponding to unit 1, because only this particular logical unit is expected to relay monitored data to the Junos Capture Vision PIC.

```
routing-instances fbf_inst {
  instance-type forwarding;
  routing-options {
    static {
      route 0.0.0.0/0 next-hop dfc-0/0/0.1;
    }
  }
}
```

Configure routing table groups:

```
[edit]
routing-options {
  interface-routes {
    rib-group inet common;
  }
  rib-groups {
    common {
      import-rib [ inet.0 fbf_inst.inet.0 ];
    }
  }
  forwarding-table {
    export pplb;
  }
}
```

Configure interfaces to the control source and content destination:

```
interfaces fe-4/1/2 {
  description "to cs1 from dfc";
  unit 0 {
    family inet {
      address 10.36.41.2/30;
    }
  }
}
interfaces ge-7/0/0 {
```

```
description "to cd1 from dfc";
unit 0 {
    family inet {
        address 10.36.70.1/30;
    }
}
}
```

RELATED DOCUMENTATION

[Understanding Junos Capture Vision | 286](#)

[Configuring Junos Capture Vision | 289](#)

Monitoring a Capture Group Using SNMP or Show Services Commands

In Junos OS Release 7.5 and later, the Dynamic Flow Capture MIB provides a way to monitor dynamic flow capture information by using Simple Network Management Protocol (SNMP). The MIB provides the same information that you can view with the `show services dynamic-flow-capture content-destination`, `show services dynamic-flow-capture control-source`, and `show services dynamic-flow-capture statistics` commands. For more information, see the *Junos Network Management Configuration Guide*.

Detecting Threats and Intercepting Flows Using Junos Packet Vision

IN THIS CHAPTER

- [Understanding Junos Packet Vision | 302](#)
- [Configuring Junos Packet Vision on MX, M and T Series Routers | 303](#)
- [Examples: Configuring Junos Packet Vision on M, T, and MX Series Routers | 306](#)
- [Sending Packets to a Mediation Device on MX, M and T Series Routers | 309](#)
- [Example: Configuring IPv6 Support for FlowTapLite on an M120 Router With Enhanced III FPCs | 310](#)

Understanding Junos Packet Vision

Junos Capture Vision (previously known as dynamic flow capture) enables you to capture packet flows on the basis of dynamic filtering criteria, using Dynamic Tasking Control Protocol (DTCP) requests. Junos Packet Vision is a Junos OS application that performs lawful intercept of packet flows, using Dynamic Tasking Control Protocol (DTCP). The application extends the use of DTCP to intercept IPv4 and IPv6 packets in an active monitoring router and send a copy of packets that match filter criteria to one or more content destinations. Junos Packet Vision was previously known as flow-tap application.

Junos Packet Vision data can be used in the following applications:

- Flexible trend analysis for detection of new security threats
- Lawful intercept

Junos Packet Vision is supported on M Series and T Series routers, except M160 and TX Matrix routers. Junos Packet Vision filters are applied on all IPv4 traffic and do not add any perceptible delay in the forwarding path. Junos Packet Vision filters can also be applied on IPv6 traffic. For security, filters installed by one client are not visible to others and the CLI configuration does not reveal the identity of the monitored target. A lighter version of the application is supported on MX Series routers only.

RELATED DOCUMENTATION

[Configuring Junos Packet Vision on MX, M and T Series Routers | 303](#)

[Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs | 329](#)

[Examples: Configuring Junos Packet Vision on M, T, and MX Series Routers | 306](#)

Configuring Junos Packet Vision on MX, M and T Series Routers

IN THIS SECTION

- [Configuring the Junos Packet Vision Interface | 303](#)
- [Strengthening Junos Packet Vision Security | 304](#)
- [Restrictions on Junos Packet Vision Services | 305](#)

This topic explains Junos Packet Vision (previously known as Flow-Tap) configuration.

Configuring the Junos Packet Vision Interface

To configure an adaptive services interface for flow-tap service, include the `interface` statement at the `[edit services flow-tap]` hierarchy level:

```
interface sp-fpc/pic/port.unit-number;
```

You can assign any Adaptive Services or Multiservices PIC in the active monitoring router for Junos Packet Vision, and use any logical unit on the PIC.

You can specify the type of traffic for which you want to apply the Junos Packet Vision service by including the `family inet | inet6` statement. If the `family` statement is not included, the Junos Packet Vision service is, by default, applied to the IPv4 traffic. To apply Junos Packet Vision service to IPv6 traffic, you must include the `family inet6` statement in the configuration. To enable the Junos Packet Vision service for IPv4 and IPv6 traffic, you must explicitly configure the `family` statement for both `inet` and `inet6` families.

NOTE: You cannot configure Junos Capture Vision (previously known as dynamic flow capture) and Junos Packet Vision services on the same router simultaneously.

You must also configure the logical interface at the `[edit interfaces]` hierarchy level:

```
interface sp-fpc/pic/port {
    unit logical-unit-number {
        family inet;
        family inet6;
    }
}
```

NOTE: If you do not include the `family inet6` statement in the configuration, IPv6 flows are not intercepted. Note that the Flow-Tap solution did not support IPv6.

Strengthening Junos Packet Vision Security

You can add an extra level of security to Dynamic Tasking Control Protocol (DTCP) transactions between the mediation device and the router by enabling DTCP sessions on top of the SSH layer. To configure SSH settings, include the `flow-tap-dtcp` statement at the `[edit system services]` hierarchy level:

```
flow-tap-dtcp {
    ssh {
        connection-limit value;
        rate-limit value;
    }
}
```

To configure client permissions for viewing and modifying Junos Packet Vision configurations and for receiving tapped traffic, include the `permissions` statement at the `[edit system login class class-name]` hierarchy level:

```
permissions [permissions];
```

The permissions needed to use Junos Packet Vision features are as follows:

- flow-tap—Can view Junos Packet Vision configuration
- flow-tap-control—Can modify Junos Packet Vision configuration
- flow-tap-operation—Can tap flows

You can also specify user permissions on a RADIUS server, for example:

```
Bob Auth-Type := Local, User-Password = = "abc123"
Juniper-User-Permissions = "flow-tap-operation"
```

Starting in Junos OS Release 16.2, MX Series routers can process mediation device DTCP ADD requests that contain up to 15 source-destination port pairs. Multiple source-destination port pairs must be separated by commas. For example:

```
ADD DTCP/0.7
Csource-ID: ftap
Cdest-ID: cd2
Source-Port: 2000,8001,4000,5000,6000,6001,6002
Dest-Port: 2000,9001,4000,5000,6000,9000
```

For details on [edit system] and RADIUS configuration, see the [User Access and Authentication Administration Guide](#).

Restrictions on Junos Packet Vision Services

The following restrictions apply to Junos Packet Vision services:

- You cannot configure Junos Capture Vision and Junos Packet Vision features on the same router simultaneously.
- On routers that support LMNR-based FPCs, you cannot configure the Junos Packet Vision for IPv6 along with port mirroring or sampling of IPv6 traffic. This restriction applies even if the router does not have any LMNR-based FPC installed in it. However, there is no restriction on configuring Junos Packet Vision on routers that are configured for port mirroring or sampling of IPv4 traffic.
- Junos Packet Vision does not support interception of MPLS and virtual private LAN service (VPLS).
- Junos Packet Vision cannot intercept Address Resolution Protocol (ARP) and other Layer 2 exceptions.
- IPv4 and IPv6 intercept filters can coexist on a system, subject to a combined maximum of 100 filters.

- When Junos Capture Vision process or the Adaptive Services or Multiservices PIC configured for Junos Packet Vision restarts, all filters are deleted and the mediation devices are disconnected.
- Only the first fragment of an IPv4 fragmented packet stream is sent to the content destination.
- Port mirroring might not work in conjunction with Junos Packet Vision.
- Running the Junos Packet Vision over an IPsec tunnel on the same router can cause packet loops and is not supported.
- M10i routers do not support the standard Junos Packet Vision, but do support FlowTapLite (see ["Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs" on page 329](#)). Junos Packet Vision and FlowTapLite cannot be configured simultaneously on the same chassis.
- PIC-based flow-tap is not supported on M7i and M10i routers equipped with an Enhanced Compact Forwarding Engine Board (CFEB-E).
- You cannot configure Junos Packet Vision on channelized interfaces.

Release History Table

Release	Description
16.2	Starting in Junos OS Release 16.2, MX Series routers can process mediation device DTCP ADD requests that contain up to 15 source-destination port pairs. Multiple source-destination port pairs must be separated by commas.

RELATED DOCUMENTATION

| [Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs](#) | 329

Examples: Configuring Junos Packet Vision on M, T, and MX Series Routers

The following example shows all parts of a complete Junos Packet Vision configuration with IPv4 and IPv6 flow intercepts:

NOTE: The following example applies only to M Series and T Series routers, except M160 and TX Matrix routers. For MX Series routers, because the flow-tap application resides in the Packet Forwarding Engine rather than a service PIC or Dense Port Concentrator (DPC), the Packet

Forwarding Engine must send the packet to a tunnel logical (vt-) interface to encapsulate the intercepted packet. In such a scenario, you need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use.

```

services {
    flow-tap {
        interface sp-1/2/0.100;
    }
}
interfaces {
    sp-1/2/0 {
        unit 100 {
            family inet;
            family inet6;
        }
    }
}
system {
    services {
        flow-tap-dtcp {
            ssh {
                connection-limit 5;
                rate-limit 5;
            }
        }
    }
    login {
        class ft-class {
            permissions flow-tap-operation;
        }
        user ft-user1 {
            class ft-class;
            authentication {
                encrypted-password "xxxx";
            }
        }
    }
}

```

The following example shows a FlowTapLite configuration that intercepts IPv4 and IPv6 flows:

```
system {
  login {
    class flowtap {
      permissions flow-tap-operation;
    }
    user ftap {
      uid 2000;
      class flowtap;
      authentication {
        encrypted-password "$ABC123"; ## SECRET-DATA
      }
    }
  }
  services {
    flow-tap-dtcp {
      ssh;
    }
  }
}

chassis {
  fpc 0 {
    pic 0 {
      tunnel-services {
        bandwidth 10g;
      }
    }
  }
}

interfaces {
  vt-0/0/0 {
    unit 0 {
      family inet;
      family inet6;
    }
  }
}

services {
  flow-tap {
    tunnel-interface vt-0/0/0.0;
```

```
}
}
```

RELATED DOCUMENTATION

[Understanding Junos Packet Vision | 302](#)

[Configuring Junos Packet Vision on MX, M and T Series Routers | 303](#)

[Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs | 329](#)

Sending Packets to a Mediation Device on MX, M and T Series Routers

Dynamic flow capture enables you to capture passively monitored packet flows on the basis of dynamic filtering criteria, using Dynamic Tasking Control Protocol (DTCP) requests. The flow-tap application extends the use of DTCP to intercept IPv4 packets in an active flow monitoring station and send a copy of packets that match filter criteria to one or more content destinations. Flow-tap data can be used for lawful intercept purposes and provides flexible trend analysis for detection of new security threats. The flow-tap application is supported on M Series and T Series routers, except M160 routers and TX Matrix platforms.

NOTE: . For information about DTCP, see Internet draft draft-cavuto-dtcp-01.txt at <http://www.ietf.org/internet-drafts>.

For detailed information about the flow-tap application, see the following sections:

- ["Understanding Flow-Tap Architecture" on page 321](#)
- ["Configuring a Flow-Tap Interface on MX, M and T Series Routers " on page 325](#)
- ["Configuring Flow-Tap Security Properties on MX, M and T Series Routers" on page 326](#)
- ["Flow-Tap Application Restrictions " on page 327](#)
- ["Example: Flow-Tap Configuration on T and M Series Routers" on page 327](#)

Example: Configuring IPv6 Support for FlowTapLite on an M120 Router With Enhanced III FPCs

IN THIS SECTION

- Requirements | 311
- Overview and Topology | 312
- Configuration | 313
- Verification | 317

This example describes how to configure IPv6 support for FlowTapLite on an M120 router with Enhanced III FPCs. The configuration of FlowTapLite is similar on an M320 router and an MX Series router with Enhanced III FPCs. However, because the MX Series routers do not support Tunnel Services PICs, you configure a DPC and the corresponding Packet Forwarding Engine to use tunneling services at the `[edit chassis]` hierarchy level.

With Junos OS Release 10.1, the FlowTapLite service supports lawful interception of IPv6 packets; previously only interception of IPv4 packets was supported. The intercepted packets are sent to a content destination, while the flow of original packets to the actual destination is unaffected.

A mediation device installs dynamic filters on the router (or server) by sending DTCP requests. These filters include the quintuple information (source address, destination address, source port, destination port, and protocol) about the intercepted flows and the details (IP addresses and port information) of the content destination.

Below is an example of such a filter:

```
ADD DTCP/0.8
Csource-ID: ftap
Cdest-ID: cd1
Source-Address: 2001:db8:abcd:ef12:3456:78ab:abc8:1235/112
Dest-Address: 2001:db8:affe::1:1
Source-Port: 1234
Dest-Port: 2345
Protocol: *
Priority: 2
X-JTap-Input-Interface: ge-2/0/1
X-JTap-Cdest-Dest-Address: 192.0.2.5
```

```

X-JTap-Cdest-Dest-Port: 2300
X-JTap-Cdest-Source-Address: 198.51.100.9
X-JTap-Cdest-Source-Port: 65535
X-JTap-Cdest-TTL: 255
X-JTap-IP-Version: ipv6
Flags: STATIC

```

Following are descriptions of the parameters in the dynamic filter:

- **Csource-ID**—The username configured in the router at the [edit system login user] hierarchy level.
- **Cdest-ID**—The content destination identifier.
- **Source-Address, Dest-Address Source-Port, Dest-Port, Protocol**—Parameters that determine which packet flows need to be intercepted.
- **X-JTap-Input-Interface**—The interface through which the actual flows are coming into the router. Depending on the type of filters installed, the value in this field can include the following: X-JTap-Output-Interface to install output interface filters; X-JTap-VRF-NAME to install VRF filters; and to install global filters, no parameters are specified.
- **X-JTap-Cdest-Dest**—All parameters that start with this string specify different parameters associated with the content destination.
- **X-JTap-IP-Version**—Differentiates between IPv6 and IPv4 filters.

From the Packet Forwarding Engine console, you can verify that the filters are installed and working correctly.

Requirements

This example uses the following hardware and software components:

- Junos OS Release 10.1 or later
- M120 router with a tunnel (vt) interface

Before you configure IPv6 FlowTapLite on your router, be sure you have:

- A tunnel PIC that is up
- A connection from the router to the mediation device and the content destination
- Traffic flow to and from the router

Overview and Topology

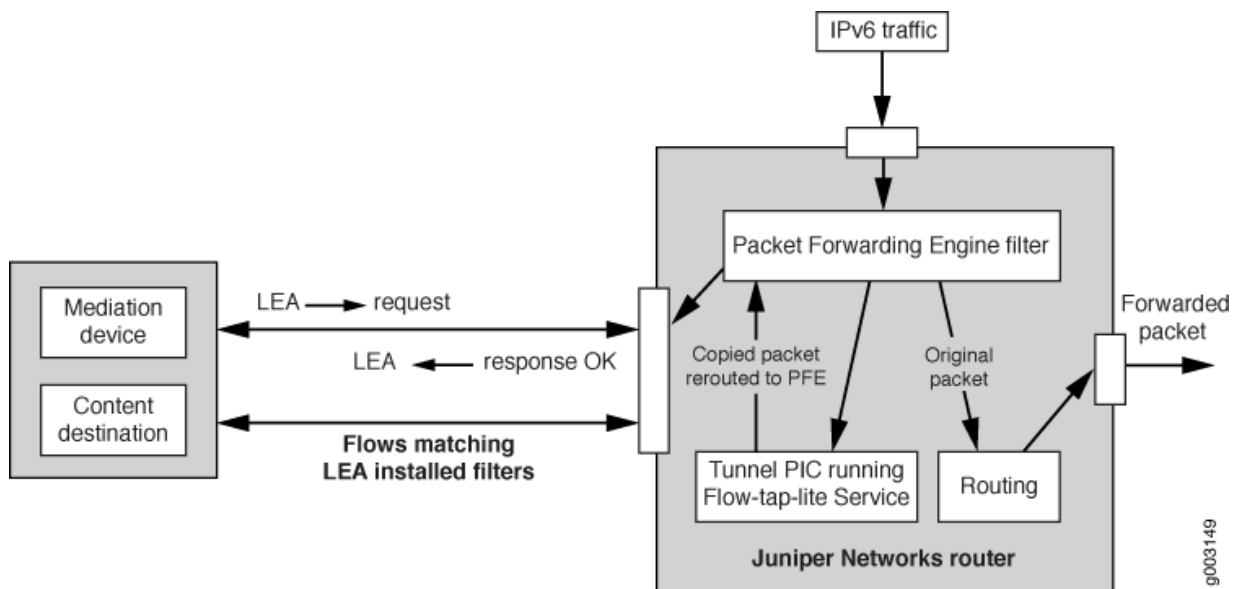
IN THIS SECTION

- Topology | 312

Figure 26 on page 312 shows the FlowTapLite configuration for one M120 router to lawfully intercept packets.

Topology

Figure 26: FlowTapLite Topology



In this example, the IPv6 packets enter the Packet Forwarding Engine and, depending on the filters installed, a new flow is created for the intercepted packets while the original packets are forwarded normally. The new flow is rerouted through the tunnel PIC back to the Packet Forwarding Engine for a route lookup, and then on to the content destination.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 313](#)
- [Configuring User Credentials | 313](#)
- [Configuring the Tunnel Interface for FlowTapLite | 314](#)
- [Configuring the Logical Tunnel Interface | 315](#)
- [Configuring FlowTapLite | 315](#)
- [Results | 316](#)

CLI Quick Configuration

To quickly configure IPv6 FlowTapLite, copy the following commands and paste them into the CLI:

```
set system login class flowtap permissions flow-tap-operation
set system login user ftap uid 2000
set system login user ftap class flowtap
set system login user ftap authentication encrypted-password "$ABC123"
set system services flow-tap-dtcp ssh
set interfaces vt-4/0/0 unit 0 family inet
set interfaces vt-4/0/0 unit 0 family inet6
set services flow-tap tunnel-interface vt-4/0/0.0
```

Configuring User Credentials

Step-by-Step Procedure

The username and password configured here are used by the mediation device when connecting and sending out DTCP requests.

1. Define a login class called flowtap:

```
[edit system]
user@router# set login class flowtap permissions flow-tap-operation
```

2. For the mediation device, configure a user called `ftap` with a unique identifier (UID):

```
[edit system]
user@router# set login user ftap uid 2000
```

3. Apply the `flowtap` class to the `ftap` user:

```
[edit system]
user@router# set login user ftap class flowtap
```

4. Configure the encrypted password used by the mediation device:

```
[edit system]
user@router# set login user ftap authentication encrypted-password $ABC123
```

5. Commit the configuration:

```
[edit system]
user@router# commit
```

Configuring the Tunnel Interface for FlowTapLite

Step-by-Step Procedure

You can add an extra level of security to DTCP transactions between the mediation device and the router by enabling DTCP sessions on top of the SSH layer.

1. Configure SSH from the `[edit system]` hierarchy level:

```
[edit system]
user@router# set services flow-tap-dtcp ssh
```

2. Commit the configuration:

```
[edit system]
user@router# commit
```

Configuring the Logical Tunnel Interface

Step-by-Step Procedure

1. Configure the logical interface and assign it to the dynamic flow control process (dfcd) at the [edit interfaces] hierarchy level:

```
[edit interfaces]
user@router# set vt-4/0/0 unit 0 family inet
```

2. Include the mandatory inet6 statement:

```
[edit interfaces]
user@router# set vt-4/0/0 unit 0 family inet6
```

3. Commit the configuration:

```
[edit interfaces]
user@router# commit
```

Configuring FlowTapLite

Step-by-Step Procedure

1. Include the flow-tap statement and the tunnel interface at the [edit services] hierarchy level:

```
[edit services]
user@router# set flow-tap tunnel-interface vt-4/0/0.0
```

2. Commit the configuration:

```
[edit services]
user@router# commit
```

Results

Check the results of the configuration:

```
[edit]
user@router# show
system {
  [...Output Truncated...]
  login {
    class flowtap {
      permissions flow-tap-operation;
    }
    user ftap {
      uid 2000;
      class flowtap;
      authentication {
        encrypted-password "$ABC123"; ## SECRET-DATA
      }
    }
  }
  services {
    telnet;
    flow-tap-dtcp {
      ssh;
    }
  }
}
interfaces {
  vt-4/0/0 {
    unit 0 {
      family inet;
      family inet6;
    }
  }
}
[...Output Truncated...]
services {
  flow-tap {
    tunnel-interface vt-4/0/0.0;
  }
}
```

Verification

IN THIS SECTION

- [Verifying That the Router Received the Filter Request | 317](#)
- [Checking That Filters Are Installed and Working on the Router | 317](#)
- [Sending a List Request | 319](#)

To confirm that the configuration is working properly, perform the following tasks:

Verifying That the Router Received the Filter Request

Purpose

After the mediation device sends the filters to the router, the mediation device must receive a message from the router confirming that the router has received the filter request.

Action

Check that the mediation device has received a message similar to the one below:

```
DTCP/0.8 200 OK
SEQ: 1
CRITERIA-ID: 1
TIMESTAMP: 2009-09-29 06:12:05.725
AUTHENTICATION-INFO: 55f9dc3debd3c7356951410f165f2a9cc5606063
```

Meaning

The message above is an example of a successfully received filter request.

Checking That Filters Are Installed and Working on the Router

Purpose

Action

Use the `show filter` and the `show filter index` commands to check that filters are installed:

```
user@router# show filter
```

Program Filters:

Index	Dir	Cnt	Text	Bss	Name
1	104	0	20	20	__default_bpdu_filter__
17000	52	0	4	4	__default_arp_policer__
57007	104	144	16	16	__flowtap_inet__
65280	52	0	4	4	__auto_policer_template__
65281	104	0	16	16	__auto_policer_template_1__
65282	156	0	32	32	__auto_policer_template_2__
65283	208	0	48	48	__auto_policer_template_3__
65284	260	0	64	64	__auto_policer_template_4__
65285	312	0	80	80	__auto_policer_template_5__
65286	364	0	96	96	__auto_policer_template_6__
65287	416	0	112	112	__auto_policer_template_7__
65288	468	0	128	128	__auto_policer_template_8__
37748736	156	144	80	80	__ftaplite_filter__ifl__70__out__ipv6__
37748737	156	144	80	80	__ftaplite_filter__vrf__4__in__ipv6__
37748738	156	144	80	80	__ftaplite_filter__ifl__71__in__ipv6__
37748739	156	144	80	80	__ftaplite_filter__vrf__0__in__ipv6__

```
user@router# show filter index 37748738 counters
```

Filter Counters/Policers:

Index	Packets	Bytes	Name
37748738	8851815	601923420	__ftaplite_term_ftap_3__counter

Meaning

The last four filters in the output for the `show filter` command above are the filters installed on the Packet Forwarding Engine. The `show filter index` command shows a non-zero packet count, indicating that the packets are hitting the filter.

Sending a List Request

Purpose

To verify that the correct filters are installed in the Packet Forwarding Engine.

Action

Use client software to send a list request to the Packet Forwarding Engine. In your list request, you can include the following three parameters individually or together: CSource-Id, CDest-ID, and Criteria-ID. With all requests, you must include the CSource-Id. Below is an example of a list request using the CSource-Id:

```
LIST DTCP/0.8
Csource-ID: ftap1
Flags: Both
```

Below is an example of a response:

```
DTCP/0.8 200 OK
SEQ: 51
TIMESTAMP: 2009-10-04 07:56:43.003
CRITERIA-ID: 1
CSOURCE-ID: ftap1
CDEST-ID: cd1
CSOURCE-ADDRESS: 10.209.152.15
FLAGS: Static
AVERAGE-BANDWIDTH: 0
MATCHING-PACKETS: 0
MATCHING-BYTES: 0
NUM-REFRESH: 0
LAST-REFRESH: 2009-10-04 07:54:30.870
X-JTAP-INPUT-INTERFACE: ge-2/1/1.0,ge-2/1/1.1,ge-2/1/1.2
SOURCE-ADDRESS: 203.0.113.1
DEST-ADDRESS: 192.168.0.1/32
SOURCE-PORT: 1000
DEST-PORT: 2000
PROTOCOL: 17
X-JTAP-CDEST-DEST-ADDRESS: 192.168.99.81
X-JTAP-CDEST-DEST-PORT: 8001
X-JTAP-CDEST-SOURCE-ADDRESS: 192.168.208.9
X-JTAP-CDEST-SOURCE-PORT: 34675
```

```
X-JTAP-CDEST-TTL: 64
CRITERIA-NUM: 1
CRITERIA-COUNT: 1
AUTHENTICATION-INFO: 0f49ff600a3d8d7d312c5031f74cc17540bc9200
```

You can also delete the request. Below is an example of a delete request:

```
DELETE DTCP/0.8
Csource-ID: ftap
Cdest-ID: cd1
Flags: STATIC
```

RELATED DOCUMENTATION

[Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs | 329](#)

[flow-tap | 1130](#)

Tunnel Interface Configuration on MX Series Routers Overview

Using Flow-Tap to Monitor Packet Flow

IN THIS CHAPTER

- Understanding Flow-Tap Architecture | 321
- Configuring a Flow-Tap Interface on MX, M and T Series Routers | 325
- Configuring Flow-Tap Security Properties on MX, M and T Series Routers | 326
- Flow-Tap Application Restrictions | 327
- Example: Flow-Tap Configuration on T and M Series Routers | 327
- Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs | 329

Understanding Flow-Tap Architecture

The flow-tap architecture consists of one or more *mediation devices* that send requests to a Juniper Networks router to monitor incoming data. Any packets that match specific filter criteria are forwarded to a set of one or more *content destinations*.

- Mediation device—A client that monitors electronic data or voice transfer over the network. The mediation device sends filter requests to the Juniper Networks router using the DTCP. The clients are not identified for security reasons, but have permissions defined by a set of special login classes.
- Monitoring platform—A Juniper Networks M Series or T Series router containing one or more Adaptive Services (AS) PICs, which are configured to support the flow-tap application. The monitoring platform processes the requests from the mediation devices, applies the dynamic filters, monitors incoming data flows, and sends the matched packets to the appropriate content destinations.
- Content destination—Recipient of the matched packets from the monitoring platform. Typically the matched packets are sent using an IP Security (IPSec) tunnel from the monitoring platform to another router connected to the content destination. The content destination and the mediation device can be physically located on the same host.
- Dynamic filters—The Packet Forwarding Engine automatically generates a *firewall filter* that is applied to all IPv4 routing instances. Each term in the filter includes a **flow-tap** action that is similar to the existing **sample** or **port-mirroring** actions. As long as one of the filter terms matches an incoming

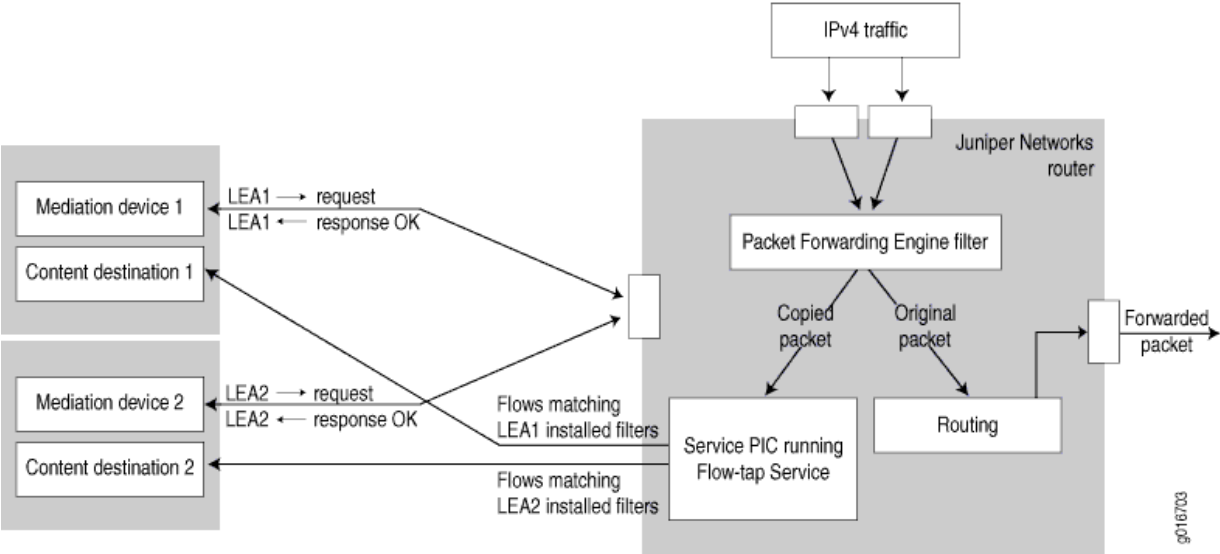
packet, the router copies the packet and forwards it to the AS PIC that is configured for flow-tap service. The AS PIC runs the packet through the client filters and sends a copy to each matching content destination. For security, filters installed by one client are not visible to others and the CLI configuration does not reveal the identity of the monitored target.

Following is a sample filter configuration; note that it is dynamically generated by the router (no user configuration is required):

```
filter combined_LEA_filter {  
  term LEA1_filter {  
    from {  
      source-address 192.0.2.;  
      destination-address 198.51.100.6;  
    }  
    then {  
      flow-tap;  
    }  
  }  
  term LEA2_filter {  
    from {  
      source-address 10.1.1.1;  
      source-port 23;  
    }  
    then {  
      flow-tap;  
    }  
  }  
}
```

[Figure 27 on page 324](#) shows a sample topology that uses two mediation devices and two content destinations.

Figure 27: Flow-Tap Topology Diagram



RELATED DOCUMENTATION

[Configuring a Flow-Tap Interface on MX, M and T Series Routers | 325](#)

[Configuring Flow-Tap Security Properties on MX, M and T Series Routers | 326](#)

[Flow-Tap Application Restrictions | 327](#)

[Example: Flow-Tap Configuration on T and M Series Routers | 327](#)

Configuring a Flow-Tap Interface on MX, M and T Series Routers

To configure an AS PIC interface for the flow-tap service, include the interface statement at the **[edit services flow-tap]** hierarchy level:

```
interface sp-fpc/pic/port.unit-number;
```

You can assign any AS PIC in the active monitoring station for flow-tap service, and use any logical unit on the PIC.

NOTE: You cannot configure dynamic flow capture and flow-tap features on the same router simultaneously.

You must also configure the logical interface at the **[edit interfaces]** hierarchy level:

```
interface sp-fpc/pic/port {
  unit logical-unit-number {
    family inet;
  }
}
```

RELATED DOCUMENTATION

[Understanding Flow-Tap Architecture | 321](#)

[Configuring Flow-Tap Security Properties on MX, M and T Series Routers | 326](#)

[Flow-Tap Application Restrictions | 327](#)

[Example: Flow-Tap Configuration on T and M Series Routers | 327](#)

Configuring Flow-Tap Security Properties on MX, M and T Series Routers

You can add an extra level of security to DTCP transactions between the mediation device and the router by enabling DTCP sessions on top of the SSH layer. To configure, include the `flow-tap-dtcp` statement at the `[edit system services]` hierarchy level:

```
flow-tap-dtcp {
  ssh {
    connection-limit value;
    rate-limit value;
  }
}
```

To configure client permissions for viewing and modifying flow-tap configurations and for receiving tapped traffic, include the `permissions` statement at the `[edit system login class class-name]` hierarchy level:

```
permissions [ permissions ];
```

The permissions needed to use flow-tap features are as follows:

- **flow-tap**—Can view flow-tap configuration.
- **flow-tap-control**—Can modify flow-tap configuration.
- **flow-tap-operation**—Can tap flows.

You can also specify user permissions on a RADIUS server, for example:

```
Bob Auth-Type := Local, User-Password = "abc123"
Juniper-User-Permissions = "flow-tap-operation"
```

For details on `[edit system]` and RADIUS configuration, see the *Junos System Basics Configuration Guide*.

RELATED DOCUMENTATION

[Understanding Flow-Tap Architecture | 321](#)

[Configuring a Flow-Tap Interface on MX, M and T Series Routers | 325](#)

[Flow-Tap Application Restrictions | 327](#)

[Example: Flow-Tap Configuration on T and M Series Routers | 327](#)

Flow-Tap Application Restrictions

The following restrictions apply to flow-tap services:

- You cannot configure dynamic flow capture and flow-tap services on the same router simultaneously.
- When the dynamic flow capture process or an AS PIC configured for flow-tap processing restarts, all filters are deleted and the mediation devices are disconnected.
- Only the first fragment of an IPv4 fragmented packet stream is sent to the content destination.
- If the flow-tap application is configured, you cannot configure the filter action **then syslog** for any *firewall filter* running on the same platform.
- Running the flow-tap application over an IPsec tunnel on the same router can cause packet loops and is not supported.
- The flow-tap service [edit services flow-tap] on tunnel interfaces on MX Series routers (FlowTapLite) and the RADIUS flow-tap service [edit services radius-flow-tap] cannot run simultaneously on the router. Consequently, you cannot run both FlowTapLite and subscriber secure policy mirroring at the same time on the same router in the earlier releases. However, starting in Junos OS Release 17.3R1, FlowTapLite and subscriber secure policy mirroring are supported to run concurrently on the same MX Series router.

Release History Table

Release	Description
17.3R1	However, starting in Junos OS Release 17.3R1, FlowTapLite and subscriber secure policy mirroring are supported to run concurrently on the same MX Series router.

RELATED DOCUMENTATION

- [Understanding Flow-Tap Architecture | 321](#)
- [Configuring a Flow-Tap Interface on MX, M and T Series Routers | 325](#)
- [Configuring Flow-Tap Security Properties on MX, M and T Series Routers | 326](#)
- [Example: Flow-Tap Configuration on T and M Series Routers | 327](#)

Example: Flow-Tap Configuration on T and M Series Routers

The following example shows all the parts of a complete flow-tap configuration.

NOTE: The following example applies only to M Series and T Series routers, except M160 and TX Matrix routers. For MX Series routers, because the flow-tap application resides in the Packet Forwarding Engine rather than a service PIC or Dense Port Concentrator (DPC), the Packet Forwarding Engine must send the packet to a tunnel logical (vt-) interface to encapsulate the intercepted packet. In such a scenario, you need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use.

```
services {
    flow-tap {
        interface sp-1/2/0.100;
    }
}
interfaces {
    sp-1/2/0 {
        unit 100 {
            family inet;
        }
    }
}
system {
    services {
        flow-tap-dtcp {
            ssh {
                connection-limit 5;
                rate-limit 5;
            }
        }
    }
    login {
        class ft-class {
            permissions flow-tap-operation;
        }
        user ft-user1 {
            class ft-class;
            authentication {
                encrypted-password "xxxx";
            }
        }
    }
}
```

```
}
}
```

RELATED DOCUMENTATION

[Understanding Flow-Tap Architecture | 321](#)

[Configuring a Flow-Tap Interface on MX, M and T Series Routers | 325](#)

[Configuring Flow-Tap Security Properties on MX, M and T Series Routers | 326](#)

[Flow-Tap Application Restrictions | 327](#)

Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs

A lighter version of the flow-tap application is available on MX Series routers and also on M320 routers with Enhanced III Flexible PIC Concentrators (FPCs). All of the functionality resides in the Packet Forwarding Engine rather than in a service PIC or Dense Port Concentrator (DPC).

Starting in Junos OS Release 17.2R1, FlowTapLite supports the sampling of circuit cross connect (CCC) traffic.

Starting in Junos OS Release 19.3R1, you can configure FlowTapLite on MX240, MX480, and MX960 routers with an MPC10E line card.

NOTE: On M320 routers only, if the replacement of FPCs results in a mode change, you must restart the dynamic flow capture process manually by disabling and then re-enabling the CLI configuration.

FlowTapLite uses the same DTCP-SSH architecture to install the Dynamic Tasking Control Protocol (DTCP) filters and authenticate the users as the original flow-tap application and supports up to 3000 filters per chassis.

NOTE: The original flow-tap application and FlowTapLite cannot be used at the same time.

To configure FlowTapLite, include the `flow-tap` statement at the `[edit services]` hierarchy level:

```
flow-tap {
    tunnel-interface interface-name;
}
```

If you do not specify a family, FlowTapLite is applied only to IPv4 traffic. Starting in Junos OS release 17.2R1, FlowTapLite can be applied to circuit cross connect traffic (ccc).

For the Packet Forwarding Engine to encapsulate the intercepted packet, it must send the packet to a tunnel logical (vt-) interface. You need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use. To create the tunnel interface, include the following configuration:

```
chassis {
    fpc number {
        pic number {
            tunnel-services {
                bandwidth (1g | 10g);
            }
        }
    }
}
```

NOTE: Currently FlowTapLite supports only one tunnel interface per instance.

To configure the logical interfaces and assign them to the dynamic flow capture process, include the following configuration:

```
interfaces {
    vt-fpc/pic/port {
        unit 0 {
            family inet;
            family inet6;
        }
    }
}
```

NOTE: If a service PIC or DPC is available, you can use its tunnel interface for the same purpose.

NOTE: If you do not include the `family inet6` statement in the configuration, IPv6 flows are not intercepted.

NOTE: With FlowTapLite configured and traceoptions enabled, if you add more than two content destinations by including the X-JTAP- CDEST-DEST-ADDRESS line in the Dynamic Tasking Control Protocol (DTCP) parameter file and initiate a DTCP session by sending a DTCP ADD message, a 400 BAD request message is received. Although you can specify more than two content destinations in the DTCP file that is sent from the mediation device, this error message occurs when the DTCP ADD message is sent. This behavior is expected with more than two content destinations. You must specify only two content destinations per DTCP ADD message.

The FlowTapLite service [edit services flow-tap] and the RADIUS flow-tap service [edit services radius-flow-tap] cannot run simultaneously on the router. Consequently, you cannot run both FlowTapLite and subscriber secure policy mirroring at the same time on the same router. Starting in Junos OS Release 17.3R1, FlowTapLite and subscriber secure policy mirroring are supported to run concurrently on the same MX Series router.

Release History Table

Release	Description
19.3R1	Starting in Junos OS Release 19.3R1, you can configure FlowTapLite on MX240, MX480, and MX960 routers with an MPC10E line card.
17.3R1	Starting in Junos OS Release 17.3R1, FlowTapLite and subscriber secure policy mirroring are supported to run concurrently on the same MX Series router.
17.2R1	Starting in Junos OS Release 17.2R1, FlowTapLite supports the sampling of circuit cross connect (CCC) traffic.
17.2R1	Starting in Junos OS release 17.2R1, FlowTapLite can be applied to circuit cross connect traffic (ccc).

RELATED DOCUMENTATION

[Understanding Junos Packet Vision | 302](#)

[Configuring Junos Packet Vision on MX, M and T Series Routers | 303](#)

[Examples: Configuring Junos Packet Vision on M, T, and MX Series Routers | 306](#)

Subscriber Secure Policy Overview

3

PART

Inline Monitoring Services and Inband Network Telemetry

[Inline Monitoring Services | 334](#)

[Flow-Based Telemetry | 349](#)

[Inband Flow Analyzer 2.0 | 370](#)

[Juniper Resiliency Interface | 409](#)

Inline Monitoring Services

IN THIS CHAPTER

- [Inline Monitoring Services Configuration | 334](#)

Inline Monitoring Services Configuration

IN THIS SECTION

- [Understanding Inline Monitoring Services | 334](#)
- [Configuring Inline Monitoring Services | 342](#)

Understanding Inline Monitoring Services

IN THIS SECTION

- [Benefits of Inline Monitoring Services | 334](#)
- [Inline Monitoring Services Feature Overview | 335](#)
- [Inline Monitoring Services Configuration Overview | 339](#)
- [Supported and Unsupported Features with Inline Monitoring Services | 341](#)

Benefits of Inline Monitoring Services

Flexible—Inline monitoring services allow different inline-monitoring instances to be mapped to different firewall filter terms, unlike in traditional sampling technologies, where all the instances are mapped to

the Flexible PIC Concentrator (FPC). This provides you with the flexibility of sampling different streams of traffic at different rates on a single interface.

Packet format agnostic—Traditional flow collection technologies rely on packet parsing and aggregation by the network element. With inline monitoring services, the packet header is exported to the collector for further processing, but without aggregation. Thereby, you have the benefit of using arbitrary packet fields to process the monitored packets at the collector.

Inline Monitoring Services Feature Overview

Service providers and content providers typically require visibility into traffic flows to evaluate peering agreements, detect traffic anomalies and policy violations, and monitor network performance. To meet these requirements, you would traditionally export aggregate flow statistics information using JFlow or IPFIX variants.

As an alternative approach, you can sample the packet content, add metadata information, and export the monitored packets to an collector. The inline monitoring services enable you to do this on MX Series routers and on PTX routers that run Junos OS Evolved.

With inline monitoring services, you can monitor every IPv4 and IPv6 packet on both ingress and egress directions of an interface. The software encapsulates the monitored traffic in an IPFIX format and exports the actual packet up to the configured clip length to an collector for further processing. By default, Junos OS supports a maximum clip length of 126 bytes starting from the Ethernet header and Junos OS Evolved supports a maximum clip length of 256 bytes starting from the Ethernet header.

Figure 28 on page 335 illustrates the IPFIX format specification.

Figure 28: Inline Monitoring IPFIX Specification

Ethernet	
IP	
UDP	
IPFIX Header	
Set	
Information Elements	

ID	Length	Description	Details
10	4B	ingressInterface	SNMP index of incoming interface
14	4B	egressInterface	SNMP index of outgoing interface when flowDirection=Output, otherwise 0.
61	1B	flowDirection	Direction (0: Input , 1:Output)
312	2B	dataLinkFrameSize	Length of sampled data link frame N octet from data link frame of monitored packet.
315	Variable	dataLinkFrameSelection	Reports actual monitored packet starting from Layer 2 as [Ethernet header/802.1Q header(any)/IP header/Payload ...] up to configured maximum-clip-length

The IPFIX header and IPFIX payload are encapsulated using IP or UDP transport layer. The exported IPFIX format includes two data records and two data templates that are exported to every collector:

- Data record—Includes incoming and outgoing interface, flow direction, data link frame section, and data link frame size. This information is sent to the collector only when sampled packets are being exported.

Figure 29 on page 337 is a sample illustration of IPFIX data record packet.

- Option data record—Includes system level information, such as exporting process ID, and sampling interval. This information is sent to the collector periodically, irrespective of whether sampling packets are being exported are not.

Figure 30 on page 337 is a sample illustration of IPFIX option data record packet.

Table 49: Information Element fields in IPFIX Option Data Packet

Number	Information Element ID	Information Element Length	Details
1	144	4B	Observation domain ID - An unique identifier of exporting process per IPFIX device. Purpose of this field is to limit the scope of other information element fields.
2	34	4B	Sampling interval at which the packets are sampled. 1000 indicates that one of 1000 packets is sampled.

- Data template—Includes five information elements:
 - Ingress interface
 - Egress interface
 - Flow direction
 - Data link frame size
 - Variable data link frame selection

Figure 31 on page 338 is a sample illustration of IPFIX data template packet.

- Option data template—Includes flow exporter and sampling interval information.

Figure 32 on page 338 is a sample illustration of IPFIX option data template packet.

When there is a new or changed inline monitoring services configuration, periodic export of data template and option data template is immediately sent to the respective collectors.

Figure 29: IPFIX Data Record

```

Version: 10
Length: 160
> Timestamp: Feb 28, 2019 14:05:41.000000000 IST
FlowSequence: 474
Observation Domain Id: 1342242816
▼ Set 1 [id=2000] (1 flows)
  FlowSet Id: (Data) (2000)
  FlowSet Length: 144
  \[Template Frame: 9\]
▼ Flow 1
  InputInt: 553
  OutputInt: 0
  Direction: Ingress (0)
  Data Link Frame Size: 1496
  ▼ Data Link Frame Section: 80711f7ce252000001000e000000450005ca000000004011...
    String_len_short: 128

```

Figure 30: IPFIX Option Data Record

```

Version: 10
Length: 28
> Timestamp: Feb 28, 2019 14:21:10.000000000 IST
FlowSequence: 11
Observation Domain Id: 1342242816
▼ Set 1 [id=2600] (1 flows)
  FlowSet Id: (Data) (2600)
  FlowSet Length: 12
  \[Template Frame: 1\]
▼ Flow 1
  FlowExporter: 1
  Sampling interval: 1

```

Figure 31: IPFIX Data Template

```

Version: 10
Length: 44
▶ Timestamp: Feb 28, 2019 14:05:42.000000000 IST
FlowSequence: 474
Observation Domain Id: 1342242816
▼ Set 1 [id=2] (Data Template): 2000
  FlowSet Id: Data Template (V10 [IPFIX]) (2)
  FlowSet Length: 28
  ▼ Template (Id = 2000, Count = 5)
    Template Id: 2000
    Field Count: 5
    ▼ Field (1/5): INPUT_SNMP
      0... .. = Pen provided: No
      .000 0000 0000 1010 = Type: INPUT_SNMP (10)
      Length: 4
    ▼ Field (2/5): OUTPUT_SNMP
      0... .. = Pen provided: No
      .000 0000 0000 1110 = Type: OUTPUT_SNMP (14)
      Length: 4
    ▼ Field (3/5): DIRECTION
      0... .. = Pen provided: No
      .000 0000 0011 1101 = Type: DIRECTION (61)
      Length: 1
    ▼ Field (4/5): dataLinkFrameSize
      0... .. = Pen provided: No
      .000 0001 0011 1000 = Type: dataLinkFrameSize (312)
      Length: 2
    ▼ Field (5/5): dataLinkFrameSection
      0... .. = Pen provided: No
      .000 0001 0011 1011 = Type: dataLinkFrameSection (315)
      Length: 65535 [i.e.: "Variable Length"]

```

Figure 32: IPFIX Option Data Template

```

Version: 10
Length: 36
▶ Timestamp: Feb 28, 2019 14:21:10.000000000 IST
FlowSequence: 11
Observation Domain Id: 1342242816
▼ Set 1 [id=3] (Options Template): 2600
  FlowSet Id: Options Template (V10 [IPFIX]) (3)
  FlowSet Length: 20
  ▼ Options Template (Id = 2600) (Scope Count = 1; Data Count = 1)
    Template Id: 2600
    Total Field Count: 2
    Scope Field Count: 1
    ▼ Field (1/1) [Scope]: FLOW_EXPORTER
      0... .. = Pen provided: No
      .000 0000 1001 0000 = Type: FLOW_EXPORTER (144)
      Length: 4
    ▼ Field (1/1): SAMPLING_INTERVAL
      0... .. = Pen provided: No
      .000 0000 0010 0010 = Type: SAMPLING_INTERVAL (34)
      Length: 4
    Padding: 0000

```

Inline Monitoring Services Configuration Overview

You can configure a maximum of sixteen (Junos OS) or seven (Junos OS Evolved) inline-monitoring instances that support template and collector-specific configuration parameters. Each inline monitoring instance supports up to four collectors (maximum of 64 collectors in total), and, for Junos OS only, you can specify different sampling rates under each collector configuration. Because of this flexibility, the inline monitoring services overcome the limitations of traditional sampling technologies, such as JFlow, sFlow, and port mirroring.

To configure inline monitoring:

1. You must include the `inline-monitoring` statement at the `[edit services]` hierarchy level. Here you specify the template and inline monitoring instance parameters. You must specify the collector parameters under the inline-monitoring instance.
2. Specify arbitrary match conditions using a firewall filter term and an action to accept the configured inline-monitoring instance. This maps the inline-monitoring instance to the firewall term.
3. Map the firewall filter under the family `inet` or `inet6` statement using the `inline-monitoring-instance` statement at the `[edit firewall filter name then]` hierarchy level. Starting in Junos OS Release 21.1R1, you can also map the firewall filter under the family `any`, `bridge`, `ccc`, `mpls`, or `vpls` statements. For Junos OS Evolved, the `bridge` and `vpls` families are not supported; use the `ethernet-switch` family instead. Junos OS Evolved does support the `any`, `ccc`, `inet`, `inet6`, and `mpls` families as well. You can also alternatively apply the firewall filter to a forwarding table filter with `input` or `output` statement to filter ingress or egress packets, respectively.

Remember:

- The device must support a maximum packet length (clip length) of 126 bytes (Junos OS) or 256 bytes (Junos OS Evolved) to enable inline monitoring services.
- You cannot configure more than 16 (Junos OS) or 7 (Junos OS Evolved) inline-monitoring instances because of the scarcity of bits available in the packet in the forwarding path.
- Apply inline monitoring services only on a collector interface, that is, the interface on which the collector is reachable. You must not apply inline monitoring on IPFIX traffic as this generates another IPFIX packet for sampling, thereby creating a loop. This includes inline monitoring service-generated traffic, such as template and record packets, option templates, and option record packets.
- When inline monitoring service is enabled on aggregated Ethernet (AE) interfaces, the information element values are as follows:

Table 50: Information Element Values for Aggregated Ethernet Interfaces

Direction of inline monitoring service on AE interface	Information element-10 (Incoming interface)	Information element-14 (Outgoing interface)
Ingress	SNMP ID of AE	0
Egress	SNMP ID of AE	SNMP ID of member link

- When inline monitoring service is enabled on IRB interfaces, the information element values are as follows:

Table 51: Information Element Values for IRB Interfaces

Direction of inline monitoring service on IRB interface	Information element-10 (Incoming interface)	Information element-14 (Outgoing interface)
Ingress	SNMP ID of IRB	0
Egress	SNMP ID of IRB	SNMP ID of vlan-bridge encapsulated interface

- For XL-XM based devices (with Lookup chip (XL) and buffering ASIC (XM)), the length of the Data Link Frame Section information element in an exported packet can be shorter than the clip length even if the egress packet length is greater than clip length.

The length of the Data Link Frame Section information element is reduced by 'N' number of bytes where 'N' = (ingress packet Layer 2 encapsulation length - egress packet Layer 2 encapsulation length).

For instance, the Layer 2 encapsulation length for the ingress packet is greater than that of the egress packet when the ingress packet has MPLS labels and egress packet is of IPv4 or IPv6 type. When traffic flows from the provider edge (PE) device to the customer edge (CE) device, the ingress packet has VLAN tags and the egress packet is untagged.

In such cases, the clip length can go past the last address location of the packet head, generating a PKT_HEAD_SIZE system log message. This can result in degradation of packet forwarding for the device.

- In case of inline monitoring services in the ingress direction, the egressInterface (information element ID 14) does not report SNMP index of the output interface. This information element ID always

reports value zero in case of ingress direction. The receiving collector process should identify the validity of this field based on the `flowDirection` (information element ID 61).

Supported and Unsupported Features with Inline Monitoring Services

Inline monitoring services supports:

- Graceful Routing Engine switchover
- In-service software upgrade (ISSU), nonstop software upgrade (NSSU), and nonstop active routing (NSR)
- Ethernet interfaces and integrated routing and bridging (IRB) interfaces
- Junos node slicing
- Starting in Junos OS Evolved Release 22.4R1, configuring DSCP, forwarding class, or routing instances for collectors.
- Starting in Junos OS Evolved Release 22.4R1, configuring template IDs or option template IDs.

Inline monitoring services currently does not support:

- Configuring more than 16 (Junos OS) or 7 (Junos OS Evolved) inline-monitoring instances.
- Junos Traffic Vision
- Prior to Junos OS Release 21.1R1, the inline-monitoring-instance term action is supported only for `inet` and `inet6` family firewall filters. Starting in Junos OS Release 21.1R1, it is supported for the `any`, `bridge`, `ccc`, `mpls`, and `vpls` family firewall filters.
- IPv6 addressable collectors
- Virtual platforms
- Logical systems
- Configuring both the observation domain ID and observation cloud ID. You must choose only one of them.
- An inline monitoring instance action used for exception reporting cannot be used for any other purpose, such as a firewall re-direct action or a regular inline-monitoring action.
- An inline monitoring instance used for a firewall re-direct action cannot be used for any other purpose, such as exception reporting or a regular inline-monitoring action.
- Prior to Junos OS Evolved Release 22.4R1, configuring DSCP, forwarding class, or routing instances for collectors.

- Prior to Junos OS Evolved Release 22.4R1, configuring template IDs or option template IDs. The system generates these for you.
- Configuring port mirroring and inline monitoring services under the same firewall filter term (Junos OS Evolved).
- In the egress direction, configuring both SFlow and exception reporting; you must choose only one of them (Junos OS Evolved).

Configuring Inline Monitoring Services

The inline monitoring services can monitor both IPv4 and IPv6 traffic on both ingress and egress directions. You can enable inline monitoring on MX Series routers with MPCs (Junos OS) and on PTX routers that run Junos OS Evolved.

You can configure inline monitoring services to monitor different streams of traffic at different sampling rates on the same logical unit of the interface. You can also export the original packet size to a collector along with information on the interface origin for effective troubleshooting.

Before You Configure

When you configure inline monitoring services, you can:

- Configure up to 16 (Junos OS) or 7 (Junos OS Evolved) inline-monitoring instances. Under each instance, you can configure specific collector and template parameters.
- Configure up to 4 IPv4-addressable collectors under each inline-monitoring instance. In total, you can configure up to 64 collectors. The collectors can be remote, and at different locations.

For each collector, you can configure specific parameters, such as source and destination address, and so on. The default routing-instance name at the collector is `default.inet`.

- For Junos OS, you can configure the `inet` or `inet6` family firewall filter with the term action `inline-monitoring-instance` *inline-monitoring-instance-name*. Starting in Junos OS Release 21.1R1, you can configure `any`, `bridge`, `ccc`, `mpls`, or `vpls` family firewall filters with the term action `inline-monitoring-instance` *inline-monitoring-instance-name*. For Junos OS Evolved, you can configure the `any`, `ccc`, `ethernet-switch`, `inet`, `inet6`, or `mpls` family firewall filters with the term action `inline-monitoring-instance` *inline-monitoring-instance-name*.

Each term can support a different inline-monitoring instance.


- Attach the inline monitoring firewall filter under the family of the logical unit of the interface.

After successfully committing the configuration, you can verify the implementation of the inline monitoring services by issuing the ["show services inline-monitoring statistics fpc-slot"](#) on page 1707 command from the CLI.

NOTE: If a packet requires inline monitoring services to be applied along with any of the traditional sampling technologies (such as JFlow or SFlow), the Packet Forwarding Engine performs both inline monitoring services and the traditional sampling technology on that packet. Port mirroring currently must be configured under a different term for Junos OS Evolved.

Figure 33 on page 343 is a sample illustration of inline monitoring services, where traffic is monitored at two different sampling rates on the device interface, and exported to four remote collectors in an IPFIX encapsulation format. For Junos OS, you configure the sampling rate on each collector, allowing different rates for each collector. For Junos OS Evolved, you configure the sampling rate on the inline-monitoring instance, and it applies to all of the collectors configured for that instance.

Figure 33: Inline Monitoring Services



In this example, the et-1/0/0 interface of the device is configured with inline monitoring services. The details of the configurations are as follows:

- There are two inline-monitoring instances — Instance 1 and Instance 2.
- There are four collectors, two collectors under each inline monitoring instance.
 - Instance 1 has Collector-1 and Collector-2.
 - Instance 2 has Collector-101 and Collector-102.
- The collectors on Instance 1 have a sampling rate of 1:10000.
- The collectors on Instance 2 have a sampling rate of 1:1.
- Instance 1 collectors have a source and destination address of 10.1.1.1 and 10.2.2.1, respectively.
- Instance 2 collectors have a source and destination address of 10.11.1.1 and 10.12.2.1, respectively.

- The packets are exported to the collectors in an IPFIX encapsulated format.

To configure inline monitoring services:

1. Define a firewall filter for each inline-monitoring instance for servicing the inline monitoring services. You can configure a family firewall filter with the term action `inline-monitoring-instance`.

To define a firewall filter:

```
[edit firewall family family filter filter-name term term]
user@host# set from source-address source-IPv4-address
user@host# set from destination-address destination-IPv4-address
user@host# set then inline-monitoring-instance inline-monitoring-instance-name
user@host# set then action
```

In this example, Terms t1 and t2 are configured for Instance1 and Instance2, respectively.

```
[edit firewall family inet filter SAMPLE_FOR_1 term t1]
user@host# set from source-address 10.1.1.0/24
user@host# set from destination-address 10.2.2.0/24
user@host# set then inline-monitoring-instance Instance1
user@host# set then accept
user@host# set term t2 from source-address 10.11.1.0/24
user@host# set term t2 from destination-address 10.12.2.0/24
user@host# set term t2 then inline-monitoring-instance Instance2
user@host# set term t2 then accept
```

2. Enable inline monitoring services by configuring the associated template, instance, and collector parameters.
 - a. To configure the inline monitoring services template:

```
[edit services inline-monitoring template template-name]
user@host# set template-refresh-rate template-refresh-rate
user@host# set option-template-refresh-rate option-template-refresh-rate
user@host# set observation-domain-id observation-domain-id
```

In this example, templates template-1 and template-2 are configured.

```
[edit services inline-monitoring template template-1]
user@host# set template-refresh-rate 60
user@host# set option-template-refresh-rate 100
```

```

user@host# set observation-domain-id 1
[edit services inline-monitoring template template-2]
user@host# set template-refresh-rate 60
user@host# set option-template-refresh-rate 100
user@host# set observation-domain-id 2

```

- b. To configure inline monitoring instance and collector parameters:

For Junos OS:

```

[edit services inline-monitoring instance inline-monitoring-instance-name]
user@host# set template-name template-name
user@host# set maximum-clip-length maximum-clip-length
user@host# set collector collector-name source-address source-IPv4-address
user@host# set collector collector-name destination-address destination-IPv4-address
user@host# set collector collector-name destination-port destination-port
user@host# set collector collector-name sampling-rate sampling-rate

```

In this example for Junos OS, Instance1 has two collectors, collector-1 and collector-2, and Instance2 has two collectors, collector-101 and collector-102. Different sampling rates have been configured for both the instances.

```

[edit services inline-monitoring instance Instance1]
user@host# set template-name template-1
user@host# set maximum-clip-length 126
user@host# set collector collector-1 source-address 10.1.1.1
user@host# set collector collector-1 destination-address 10.2.2.1
user@host# set collector collector-1 destination-port 2055
user@host# set collector collector-1 sampling-rate 10000
user@host# set collector collector-2 source-address 10.1.1.1
user@host# set collector collector-2 destination-address 10.2.2.1
user@host# set collector collector-2 destination-port 2055
user@host# set collector collector-2 sampling-rate 10000

```

```

[edit services inline-monitoring instance Instance2]
user@host# set template-name template-2
user@host# set maximum-clip-length 126
user@host# set collector collector-101 source-address 10.11.1.1
user@host# set collector collector-101 destination-address 10.12.2.1
user@host# set collector collector-101 destination-port 2055

```

```

user@host# set collector collector-101 sampling-rate 1
user@host# set collector collector-102 source-address 10.11.1.1
user@host# set collector collector-102 destination-address 10.12.2.1
user@host# set collector collector-102 destination-port 2055
user@host# set collector collector-102 sampling-rate 1

```

For Junos OS Evolved:

```

[edit services inline-monitoring instance inline-monitoring-instance-name]
user@host# set template-name template-name
user@host# set maximum-clip-length maximum-clip-length
user@host# set sampling-rate sampling-rate
user@host# set collector collector-name source-address source-IPv4-address
user@host# set collector collector-name destination-address destination-IPv4-address
user@host# set collector collector-name destination-port destination-port

```

In this example, for Junos OS Evolved, Instance1 has two collectors, collector-1 and collector-2, and Instance2 has two collectors, collector-101 and collector-102. Different sampling rates have been configured for both the instances.

```

[edit services inline-monitoring instance Instance1]
user@host# set template-name template-1
user@host# set maximum-clip-length 126
user@host# set sampling-rate 10000
user@host# set collector collector-1 source-address 10.1.1.1
user@host# set collector collector-1 destination-address 10.2.2.1
user@host# set collector collector-1 destination-port 2055
user@host# set collector collector-2 source-address 10.1.1.1
user@host# set collector collector-2 destination-address 10.2.2.1
user@host# set collector collector-2 destination-port 2055

```

```

[edit services inline-monitoring instance Instance2]
user@host# set template-name template-2
user@host# set maximum-clip-length 126
user@host# set sampling-rate 1
user@host# set collector collector-101 source-address 10.11.1.1
user@host# set collector collector-101 destination-address 10.12.2.1
user@host# set collector collector-101 destination-port 2055
user@host# set collector collector-102 source-address 10.11.1.1

```

```
user@host# set collector collector-102 destination-address 10.12.2.1
user@host# set collector collector-102 destination-port 2055
```

3. Map the firewall filter under the family of the logical unit of the interface to apply inline monitoring in the ingress or egress direction.

Alternatively, you can apply inline monitoring by mapping the firewall filter to a forwarding table filter with an input or output statement to filter ingress or egress packets, respectively.

To attach the firewall filter:

```
[edit interfaces interface-name]
user@host# set unit 0 family family filter input filter
user@host# set unit 0 family family address ip-address
```

In this example, the inline monitoring filter is attached to family inet of unit 0 of et-1/0/0.

```
[edit interfaces et-1/0/0]
user@host# set unit 0 family inet filter input SAMPLE_FOR_1
user@host# set unit 0 family inet address 10.100.0.1/30
```

Release History Table

Release	Description
22.4R1-EVO	Inline monitoring services (PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers with either the JNP10K-LC1201 or JNP10K-LC1203 linecards) - Starting in Junos OS Evolved Release 22.4R1, you can configure inline monitoring services on the PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers to sample packets, add metadata, and export the packets up to the configured clip length to an IPFIX collector for further processing. You can also configure the any, ccc, ethernet-switch, inet, inet6, or mpls family firewall filters with the term action inline-monitoring-instance <i>inline-monitoring-instance-name</i> .
22.3R1	Inline monitoring services (MX304 routers) - Starting in Junos OS Release 22.3R1, you can configure inline monitoring services on the MX304 router.
22.2R1-EVO	Inline monitoring services (PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers with either the JNP10K-LC1201 or JNP10K-LC1203 linecards) - Starting in Junos OS Evolved Release 22.1R1, you can configure inline monitoring services on the PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers to report exceptions. You can also configure the any, ccc, ethernet-switch, inet, inet6, or mpls family firewall filters with the term action inline-monitoring-instance <i>inline-monitoring-instance-name</i> .

21.4R1	Inline monitoring services (LC9600 linecard for the MX10008 router) - Starting in Junos OS Release 21.4R1, you can configure inline monitoring services on MX10008 routers that contain the LC9600 linecard.
21.2R1	Support for Layer 2 and any firewall filter families for inline monitoring services (MX Series with MPC10E and MPC11E linecards)—Starting in Junos OS Release 21.2R1, you can configure the any, bridge, ccc, mpls, or vpls family firewall filters with the term action inline-monitoring-instance <i>inline-monitoring-instance-name</i> .
21.2R1	Inline monitoring services (LC480 linecard for MX10008 and MX10016 routers - Starting in Junos OS Release 21.2R1, you can configure inline monitoring services on MX10008 and MX10016 routers that contain the LC480 linecard.
21.1R1	Support for Layer 2 and any firewall filter families for inline monitoring services (MX Series with MPCs excluding MPC10E and MPC11E linecards)—Starting in Junos OS Release 21.1R1, you can configure the any, bridge, ccc, mpls, or vpls family firewall filters with the term action inline-monitoring-instance <i>inline-monitoring-instance-name</i> .
20.4R1	Inline monitoring services (MPC10E and MPC11E linecards for MX Series routers - Starting in Junos OS Release 20.4R1, you can configure inline monitoring services on MX Series routers that contain the MPC10E and MPC11E linecards.
19.4R1	Inline monitoring services (MX Series with MPCs excluding MPC10E and MPC11E linecards) - Starting in Junos OS Release 19.4R1, you can configure a new monitoring technology that provides the flexibility to monitor different streams of traffic at different sampling rates on the same interface. You can also export the packet up to the configured clip length to a collector in an IP Flow Information Export (IPFIX) format. The IPFIX format includes important metadata information about the monitored packets for further processing at the collector.

Flow-Based Telemetry

IN THIS CHAPTER

- [Flow-Based Telemetry \(EX4100, EX4100-F, and EX4400 Series\) | 349](#)
- [Flow-Based Telemetry for VXLANs \(QFX5120\) | 361](#)

Flow-Based Telemetry (EX4100, EX4100-F, and EX4400 Series)

SUMMARY

Flow based telemetry (FBT) enables per-flow-level analytics, using inline monitoring services to create flows, collect them, and export them to a collector using the open standard IPFIX template to organize the flow.

IN THIS SECTION

- [FBT Overview | 349](#)
- [Configure FBT \(EX4100, EX4100-F, and EX4400 Series\) | 356](#)

FBT Overview

IN THIS SECTION

- [Benefits of FBT | 350](#)
- [FBT Flow Export Overview | 351](#)
- [Limitations and Caveats | 353](#)
- [Licenses | 354](#)
- [Drop Vectors \(EX4100 and EX4100-F only\) | 354](#)

You can configure flow-based telemetry (FBT) for the EX4100, EX4100-F, and EX4400 Series switches. FBT enables per-flow-level analytics, using inline monitoring services to create flows, collect them, and

export them to a collector. With inline monitoring services, you can monitor every IPv4 and IPv6 packet on both ingress and egress directions of an interface. A flow is a sequence of packets that have the same source IP, destination IP, source port, destination port, and protocol on an interface. For each flow, the software collects various parameters and exports the actual packet up to the configured clip length to a collector using the open standard IPFIX template to organize the flow. Once there is no active traffic for a flow, the flow is aged out after the configured inactive-timeout period (configure the `flow-inactive-timeout` statement at the [edit services inline-monitoring template *template-name*] hierarchy level). The software exports a IPFIX packet periodically at the configured flow-export timer interval. The observation domain identifier is used in the IPFIX packet to identify which line card sent the packet to the collector. Once set, the software derives a unique identifier for each line card based upon the system value set here.

Benefits of FBT

With FBT, you can:

- Count packet, TTL, and TCP window ranges
- Track and count Denial of Service (DoS) attacks
- Analyze the load distribution of ECMP groups/link aggregation groups (LAG) over the member IDs (EX4100 and EX4100-F only)
- Track traffic congestion (EX4100 and EX4100-F only)
- Gather information about multimedia flows (EX4100 and EX4100-F only)
- Gather information on why packets are dropped (EX4100 and EX4100-F only)

FBT Flow Export Overview

See [Figure 34 on page 351](#) for a sample template, which shows the information element IDs, names, and sizes:

Figure 34: Sample FBT Information Element Template

```
[CINT] Output information elements (total 21):
  output info element ( 0): elem = 15(          Reserved) and size = 2
  output info element ( 1): elem = 14(      FlowtrackerGroup) and size = 2
  output info element ( 2): elem = 15(          Reserved) and size = 42
  output info element ( 3): elem = 5(        L4DstPort) and size = 2
  output info element ( 4): elem = 4(        L4SrcPort) and size = 2
  output info element ( 5): elem = 1(        DstIPv4) and size = 4
  output info element ( 6): elem = 0(        SrcIPv4) and size = 4
  output info element ( 7): elem = 6(      IPProtocol) and size = 1
  output info element ( 8): elem = 15(          Reserved) and size = 1
  output info element ( 9): elem = 40(    TimestampNewLearn) and size = 6
  output info element (10): elem = 7(        PktCount) and size = 4
  output info element (11): elem = 4(        L4SrcPort) and size = 2
  output info element (12): elem = 15(          Reserved) and size = 20
  output info element (13): elem = 47(      FlowtrackerCheck) and size = 4
  output info element (14): elem = 84(    IngDropReasonGroupIdVector) and size = 2
  output info element (15): elem = 83(          IngPort) and size = 1
  output info element (16): elem = 15(          Reserved) and size = 9
  output info element (17): elem = 47(      FlowtrackerCheck) and size = 4
  output info element (18): elem = 90(    EgrDropReasonGroupIdVector) and size = 2
  output info element (19): elem = 54(          EgrPort) and size = 1
  output info element (20): elem = 15(          Reserved) and size = 9
```

jr-000303

Figure 35 on page 352 shows the format of a sample IPFIX data template for FBT:

Figure 35: Sample FBT IPFIX Data Template

Apply a display filter ... <36/>

No.	Time	Source	Destination	Protocol	Length	Info	
1209	180.970315	10.6.6.1	10.6.6.2	CFLOW	210	IPFIX flow (180 bytes)	Obs-Domain-ID=167837712 [Data-Template:1200]
1270	170.929160			CFLOW	222	IPFIX flow (180 bytes)	Obs-Domain-ID=167837712 [Data-Template:1201]
1297	179.926594	10.6.6.1	10.6.6.2	CFLOW	186	IPFIX flow (144 bytes)	Obs-Domain-ID=167837696 [Data:1200]
1298	180.070515	10.6.6.1	10.6.6.2	CFLOW	210	IPFIX flow (168 bytes)	Obs-Domain-ID=167837696 [Data-Template:1200]
1299	180.070527	10.6.6.1	10.6.6.2	CFLOW	222	IPFIX flow (180 bytes)	Obs-Domain-ID=167837696 [Data-Template:1201]
1300	180.930733	10.6.6.1	10.6.6.2	CFLOW	210	IPFIX flow (168 bytes)	Obs-Domain-ID=167837712 [Data-Template:1200]

ExportTime: 1603965361

FlowSequence: 1

Observation Domain Id: 167837712

Set 1 [id=2] (Data Template): 1201

FlowSet Id: (Data Template (V10 [IPFIX]) (2)

FlowSet Length: 164

Template (Id = 1201, Count = 24)

Template Id: 1201

Field Count: 24

Field (1/24): 255 [pen: Juniper Networks, Inc.]

Field (2/24): 254 [pen: Juniper Networks, Inc.]

Field (3/24): 255 [pen: Juniper Networks, Inc.]

Field (4/24): IPV6_SRC_ADDR

Field (5/24): 255 [pen: Juniper Networks, Inc.]

Field (6/24): L4_DST_PORT

Field (7/24): L4_SRC_PORT

Field (8/24): PROTOCOL

Field (9/24): IPV6_DST_ADDR

Field (10/24): 255 [pen: Juniper Networks, Inc.]

Field (11/24): 1 [pen: Juniper Networks, Inc.]

Field (12/24): 2 [pen: Juniper Networks, Inc.]

Field (13/24): 19 [pen: Juniper Networks, Inc.]

Field (14/24): 17 [pen: Juniper Networks, Inc.]

Field (15/24): 4 [pen: Juniper Networks, Inc.]

Field (16/24): 15 [pen: Juniper Networks, Inc.]

Field (17/24): PKTS

Field (18/24): BYTES

Field (19/24): 16 [pen: Juniper Networks, Inc.]

Field (20/24): payloadLengthIPv6

Field (21/24): TCP_WINDOW_SIZE

Field (22/24): 255 [pen: Juniper Networks, Inc.]

Field (23/24): 254 [pen: Juniper Networks, Inc.]

Field (24/24): 255 [pen: Juniper Networks, Inc.]

jn-000304

Figure 36 on page 353 shows the format of a sample exported IPFIX flow for FBT:

Figure 36: Sample Exported IPFIX Flow for FBT

```

Version: 10
Length: 144
▼ Timestamp: Jan 1, 1970 05:30:00.000000000 IST
  ExportTime: 0
▼ FlowSequence: 21 (expected 1)
  ► [Expert Info (Warning/Sequence): Unexpected flow sequence for domain ID 167837696 (expected 1, got 21)]
    Observation Domain Id: 167837696
▼ Set 1 [id=1200] (1 flows)
  FlowSet Id: (Data) (1200)
  FlowSet Length: 128
  [Template Frame: 1]
▼ Flow 1
  Enterprise Private entry: (Juniper Networks, Inc.) Type 255: Value (hex bytes): 0c bc
  Enterprise Private entry: (Juniper Networks, Inc.) Type 254: Value (hex bytes): 00 00
  Enterprise Private entry: (Juniper Networks, Inc.) Type 255: Value (hex bytes): 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 _
  DstPort: 6068
  SrcPort: 15000
  DstAddr: 192.168.100.1
  SrcAddr: 192.168.200.1
  Protocol: TCP (6)
  Enterprise Private entry: (Juniper Networks, Inc.) Type 255: Value (hex bytes): 01
  Enterprise Private entry: (Juniper Networks, Inc.) Type 1: Value (hex bytes): 00 00 00 00 00 00
  Enterprise Private entry: (Juniper Networks, Inc.) Type 2: Value (hex bytes): 00 00 00 00 00 00
  Enterprise Private entry: (Juniper Networks, Inc.) Type 19: Value (hex bytes): 00 00 31 43
  Enterprise Private entry: (Juniper Networks, Inc.) Type 17: Value (hex bytes): 00 00 00 00
  Enterprise Private entry: (Juniper Networks, Inc.) Type 4: Value (hex bytes): 00 00 31 43
  Enterprise Private entry: (Juniper Networks, Inc.) Type 15: Value (hex bytes): 00 00 31 43
  Packets: 12611
  Octets: 3228416
  Enterprise Private entry: (Juniper Networks, Inc.) Type 16: Value (hex bytes): 01 00 00 00
  IPV4 Total Length: 238
  TCP Windows Size: 4096
  Enterprise Private entry: (Juniper Networks, Inc.) Type 255: Value (hex bytes): 01 85
  IP TTL: 255

```

jn-000305

When you create a new inline monitoring services configuration or change an existing one, the software immediately sends the periodic flow export of the data template to the respective collectors, instead of waiting until the next scheduled send time.

Limitations and Caveats

- IRB interfaces are supported; however, L2 firewall filters are not supported.
- Only 8 inline-monitoring instances and 8 collectors per instance are supported.
- Flow records are limited to 128 bytes in length.
- The collector must be reachable through either the loopback interface or a network interface, not only through a management interface.
- You cannot configure an option template identifier or a forwarding class.
- The IPFIX Option Data Record and IPFIX Option Data Template are not supported.
- Feature profiles are not supported on EX4400 switches.
- If you make any changes to the feature-profile configuration, you must reboot the device.

- (EX4100 and EX4100-F only) If you configure any of the congestion or egress features in the feature profile for an inline-monitoring instance, you cannot configure a counter profile for a template in that instance.
- (EX4100 and EX4100-F only) Because the congestion and egress features collect a lot of data, you can only configure 4 or 5 of these features per inline-monitoring instance.
- (EX4100 and EX4100-F only) For multicast flow tracking, one ingress copy can produce multiple egress copies. All copies may update the same entry. Therefore, you can track the aggregate results of all copies of the same multicast flow.

Licenses

You must get a permanent license to enable FBT. To check if you have a license for FBT, issue the `show system license` command in operational mode:

```
user@host> show system license
License usage:

Feature name           Licenses   Licenses   Licenses   Expiry
                       used      installed  needed
Flow Based Telemetry    1          1          0    permanent
Licenses installed:
License identifier: XXXXXXXXXXXXXXXX
License version: 4
Order Type: commercial
Valid for device: XXXXXXXXXXXXXXXX
Features:
  Flow Based Telemetry - License for activating Flow Based Telemetry
  Permanent
```

For the EX4100 and EX4100-F switches, you need license S-EX4100-FBT-P. For the EX4400 switches, you need license S-EX-FBT-P.

Drop Vectors (EX4100 and EX4100-F only)

FBT can report more than 100 drop reasons. Drop vectors are very large vectors, too large to be reasonably accommodated in a flow record. Therefore, the software groups and compresses the drop vectors into a 16-bit compressed drop vector, and then passes that drop vector to the flow table. The 16-bit compressed drop vector corresponds to a particular drop vector group. [Table 52 on page 355](#) and [Table 53 on page 355](#) describe how drop vectors are grouped together to form a particular 16-bit compressed drop vector.

Table 52: Ingress Drop Vector Groups (EX4100 and EX4100-F only)

Group ID	Drop Reason
1	MMU drop
2	TCAM, PVLAN
3	DoS attack or LAG loopback fail
4	Invalid VLAN ID, invalid TPID, or the port is not in the VLAN
5	Spanning Tree Protocol (STP) forwarding, bridge protocol data unit (BPDU), Protocol, CML
6	Source route, L2 source discard, L2 destination discard, L3 disable, and so on.
7	L3 TTL, L3 Header, L2 Header, L3 source lookup miss, L3 destination lookup miss
8	ECMP resolution, storm control, ingress multicast, ingress next-hop error

Table 53: Egress Drop Vector Groups (EX4100 and EX4100-F only)

Group ID	Drop Reason
1	MMU unicast traffic
2	MMU weighted random early detection (WRED) unicast traffic
3	MMU RQE
4	MMU multicast traffic

Table 53: Egress Drop Vector Groups (EX4100 and EX4100-F only) (Continued)

Group ID	Drop Reason
5	Egress TTL, stgblock
6	Egress field processor drops
7	IPMC drops
8	Egress quality of service (QoS) control drops

Configure FBT (EX4100, EX4100-F, and EX4400 Series)

FBT enables per-flow-level analytics, using inline monitoring services to create flows, collect them, and export them to a collector. A flow is a sequence of packets that have the same source IP, destination IP, source port, destination port, and protocol on an interface. For each flow, various parameters are collected and sent to a collector using the open standard IPFIX template to organize the flow. Once there is no active traffic for a flow, the flow is aged out after the configured inactive-timeout period (configure the `flow-inactive-timeout` statement at the [edit services inline-monitoring template *template-name*] hierarchy level). The software exports a IPFIX packet periodically at the configured flow-export timer interval. The observation domain identifier is used in the IPFIX packet to identify which line card sent the packet to the collector. Once set, the software derives a unique identifier for each line card based upon the system value set here.

To configure flow-based telemetry:

1. Define the IPFIX template.

To configure attributes of the template:

```
user@host# set services inline-monitoring template template_1 flow-inactive-timeout seconds
user@host# set services inline-monitoring template template_1 observation-domain-id identifier
user@host# set services inline-monitoring template template_1 template-refresh-rate template-
refresh-rate
user@host# set services inline-monitoring template template_1 template-identifier template-
identifier
```

In this example, the inactive-flow timeout period is set to 10 seconds, the observation domain ID is set to 25, the template refresh rate is set to 30 seconds, and you've configured a template identifier

```
user@host# set services inline-monitoring template template_1 flow-inactive-timeout 10
user@host# set services inline-monitoring template template_1 observation-domain-id 25
user@host# set services inline-monitoring template template_1 template-refresh-rate 30
user@host# set services inline-monitoring template template_1 template-identifier 32768
```

2. Attach a template to the instance and describe the collector.

To configure the instance and collector:

```
user@host# set services inline-monitoring instance instance-name template-name template-name
user@host# set services inline-monitoring instance instance-name collector collector-name
source-address IPv4-address
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-address IPv4-address
user@host# set services inline-monitoring instance instance-name collector collector-name
dscp dscp-bits
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-port port
```

In this example, you create a template with the name `template_1`, create an inline-monitoring instance `i1`, and create the configuration for the collector `c2`:

```
user@host# set services inline-monitoring instance i1 template-name template_1
user@host# set services inline-monitoring instance i1 collector c2 source-address 10.11.12.1
user@host# set services inline-monitoring instance i1 collector c2 destination-address
10.11.12.2
user@host# set services inline-monitoring instance i1 collector c2 dscp 21
user@host# set services inline-monitoring instance i1 collector c2 destination-port 2055
```

3. Create a firewall filter and configure the action inline-monitoring-instance.

To configure the firewall filter:

```
user@host# set firewall family inet filter filter-name term term-name from source-address
source-address
user@host# set firewall family inet filter filter-name term term-name from destination-
address destination-address
user@host# set firewall family inet filter filter-name term term-name then inline-monitoring-
```



```
instance instance-name
user@host# set firewall family inet filter filter-name term term-name then accept
```

In this example, you configure an IPv4 firewall filter named `ipv4_ingress`, with the term name `rule1` containing the action `inline-monitoring-instance`, and the inline monitoring instance `i1` is mapped to it:

```
user@host# set firewall family inet filter ipv4_ingress term rule1 from source-address
10.11.12.1
user@host# set firewall family inet filter ipv4_ingress term rule1 from destination-address
10.11.12.2
user@host# set firewall family inet filter ipv4_ingress term rule1 then inline-monitoring-
instance i1
user@host# set firewall family inet filter ipv4_ingress term rule1 then accept
```

4. Map the firewall filter to the family under the logical unit of the already-configured interface to apply inline monitoring in the ingress direction.

To map the firewall filter:

```
user@host# set interface interface-name unit 0 family inet filter input filter-name
```

In this example, you map the `ipv4_ingress` firewall filter to the `inet` family of logical interface 0 of the physical interface `et-0/0/1`:

```
user@host# set interface et-0/0/1 unit 0 family inet filter input ipv4_ingress
```

5. (Optional) Configure the sampling profile and rate, configure the profile for which counters to export to the collector, configure the flow rate and burst size, and enable security analytics for flow-based telemetry:

To configure the flow-monitoring properties:

```
user@host# set services inline-monitoring template template _1 flow-monitoring sampling-
profile profile-name
user@host# set services inline-monitoring template template _1 flow-monitoring sampling-rate
rate
user@host# set services inline-monitoring template template _1 flow-monitoring counter-
profile profile-identifier
user@host# set services inline-monitoring template template _1 flow-monitoring flow-rate kbps
burst-size bytes
user@host# set services inline-monitoring template template _1 flow-monitoring security-enable
```

In this example, the sampling profile is set to Random, the sampling rate is set to every 512 bytes, the counter profile is set to Per_flow_6_counters, the flow-rate is set to 100000 kbps, the burst-size is set to 2048 bytes, and security analytics are enabled:

```
user@host# set services inline-monitoring template template _1 flow-monitoring sampling-
profile Random
user@host# set services inline-monitoring template template _1 flow-monitoring sampling-rate
512
user@host# set services inline-monitoring template template _1 flow-monitoring counter-
profile Per_flow_6_counters
user@host# set services inline-monitoring template template _1 flow-monitoring flow-rate
100000 burst-size 2048
user@host# set services inline-monitoring template template _1 flow-monitoring security-enable
```

6. (Optional, EX4100 and EX4100-F switches only) Configure a feature profile to collect more data about packets as they move through the switch.

For example, you could monitor congestion or collect information about why packets are being dropped. You can enable security analytics either here or in the previous step. To configure a feature profile:

```
user@host# set services inline-monitoring feature-profile feature_1 features aggregate-intf-
member-id
user@host# set services inline-monitoring feature-profile feature_1 features chip-delay
user@host# set services inline-monitoring feature-profile feature_1 features egress-drop-
reason
user@host# set services inline-monitoring feature-profile feature_1 features flow-start-end-
time
user@host# set services inline-monitoring feature-profile feature_1 features ingress-drop-
reason
user@host# set services inline-monitoring feature-profile feature_1 features inter-arrival-
time
user@host# set services inline-monitoring feature-profile feature_1 features inter-departure-
time
user@host# set services inline-monitoring feature-profile feature_1 features queue-congestion-
level
user@host# set services inline-monitoring feature-profile feature_1 features security-enable
user@host# set services inline-monitoring feature-profile feature_1 features shared-pool-
congestion
```

You must reboot the system for the feature profile to take effect. Because the aggregate interface distribution monitoring, congestion, and egress features collect a lot of data, you can only configure 4

or 5 of these features per inline-monitoring instance. The statements that configure these features are:

- aggregate-intf-member-id
- egress-drop-reason
- inter-departure-time
- queue-congestion-level
- shared-pool-congestion

After you commit the configuration and reboot the system, use the `show services inline-monitoring feature-profile-mapping fpc-slot slot-number` command to verify that the features have been successfully configured.

7. After committing the configuration, monitor inline-monitoring statistics with the `show services inline-monitoring statistics fpc-slot slot-number` command.

Release History Table

Release	Description
22.2R1	You can now configure flow-based telemetry (FBT) for the EX4100 and EX4100-F Series switches, and configure additional items to track for a flow using the <code>feature-profile <i>name</i> features</code> statement at the <code>[edit inline-monitoring]</code> hierarchy level.
21.1R1	You can configure flow-based telemetry (FBT) for the EX4400 Series switches. FBT enables per-flow-level analytics, using inline monitoring services to create flows, collect them, and export them to a collector.

RELATED DOCUMENTATION

Inline Monitoring Services Configuration	 334
inline-monitoring	 1163
flow-monitoring (Inline Monitoring Services)	 1122
features	 1090

Flow-Based Telemetry for VXLANs (QFX5120)

SUMMARY

Flow based telemetry (FBT) for VXLANs in Junos OS enables per-flow-level analytics on IRB interfaces, using inline monitoring services to create flows, collect them, and export them to a collector using the open standard IP Flow Information Export (IPFIX) template to organize the flow.

IN THIS SECTION

- [FBT for VXLANs Overview | 361](#)
- [Configure FBT for VXLANs \(QFX5120\) | 366](#)

FBT for VXLANs Overview

IN THIS SECTION

- [Benefit of FBT for VXLANs | 361](#)
- [Flow Export Overview | 363](#)
- [Limitations and Caveats | 365](#)

You can configure flow-based telemetry (FBT) for VXLANs for the QFX5120 -32C and QFX5120-48y-8c switches. FBT for VXLANs enables inline telemetry data for VXLANs that have either centrally-routed bridging (CRB) or edge-routed bridging (ERB) overlays. FBT for VXLANs enables per-flow-level analytics for IRB interfaces, using inline monitoring services to create flows, collect them, and export them to a collector. With inline monitoring services, you can monitor every IPv4 and IPv6 packet on both ingress and egress directions of an interface.

Benefit of FBT for VXLANs

With FBT for VXLANs, you can enable inline telemetry data for EVPN-VXLAN architectures that have either CRB or ERB overlays, giving you an additional source of information about your network,

A VXLAN with a CRB overlay has core switches configured as Layer/Layer 3 VXLAN gateways where the Integrated Routing and Bridging (IRB) interfaces for the virtual networks are configured on the core switches. In contrast, core switches in a VXLAN with an ERB overlay provide transport of EVPN type-2 and type-5 routes and the IRB interfaces are configured on the distribution switches. The ERB design also enables faster server-to-server, intra-campus traffic. As a result, with an ERB overlay, routing happens much closer to the end systems than with a CRB overlay. [Figure 37 on page 362](#) and [Figure 38](#)

on page 363 show sample topologies for these overlays. To learn more about these EVPN-VXLAN architectures, see [Technology Primer: EVPN-VXLAN Fabrics for the Campus](#).

Figure 37: Centrally-Routed Bridging (CRB) Topology

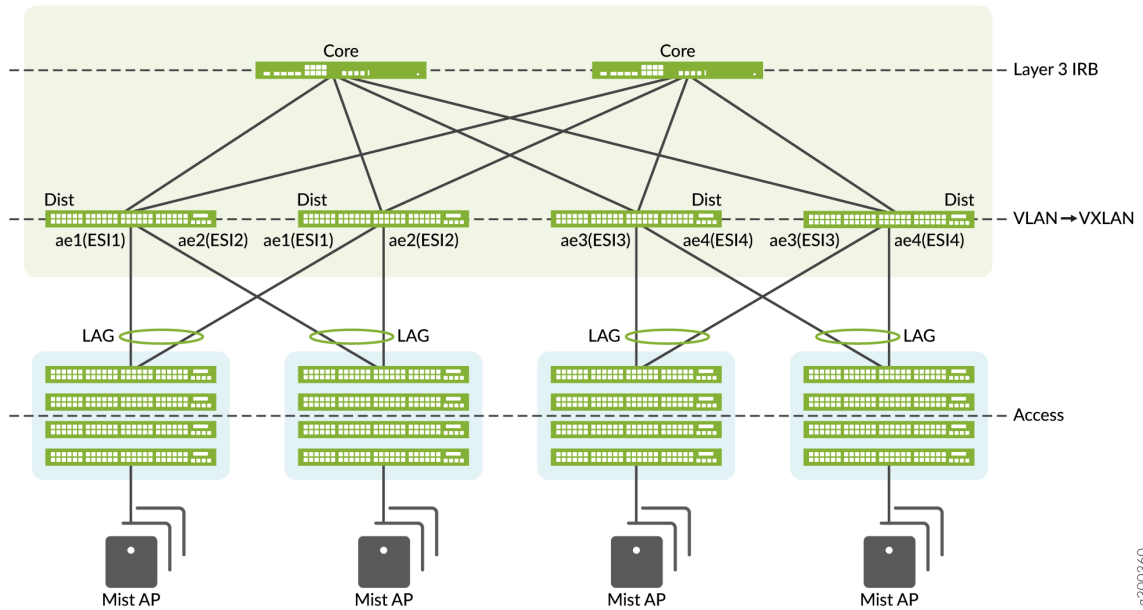



Figure 38: Edge-Routed Bridging (ERB) Topology



Flow Export Overview

FBT for VXLANs uses software-based IPFIX flow export. (IPFIX is defined in RFC 7011.) A flow is a sequence of packets that have the same core set of parameters on an interface, some of which are source IP, destination IP, source port, destination port, and protocol. This core set of parameters is called a flow key, and the software uses this key to learn about the flows. For each flow, the software collects various parameters and exports the actual packet up to the configured clip length to a collector using the open standard IPFIX template to organize the flow. Once there is no active traffic for a flow, the flow is aged out after the configured inactive-timeout period (configure the `flow-inactive-timeout` statement at the [edit services inline-monitoring template *template-name*] hierarchy level).

For FBT for VXLANs, the flow key differs depending on whether you are monitoring IPv4 or IPv6 traffic. The flow key for IPv4 traffic is explained in [Table 54 on page 364](#) and the flow key for IPv6 traffic is explained in [Table 55 on page 364](#). For both IPv4 and IPv6 traffic, in addition to the key fields, the flow contains fields for the ingress and egress ports, the flow start and end time, and the byte and packet count delta. The flow start time is the timestamp for when the software learned the flow. The flow stop time is the timestamp of the latest counter query. A sample IPFIX data template for IPv4 traffic is shown in [Figure 39 on page 365](#).

Table 54: IPv4 Flow Key

Field	Field size in bytes
Source IP address	4
Destination IP address	4
Protocol (TCP or UDP)	1
Source port (TCP or UDP)	2
Destination port (TCP or UDP)	2
Virtual routing and forwarding table (VRF) identifier	2
Ingress port	1
VXLAN network identifier (layer 2 segment ID)	3

Table 55: IPv6 Flow Key

Field	Field size in bytes
Source IP address	4
Destination IP address	4
Protocol (TCP or UDP)	1
Source port (TCP or UDP)	2
Destination port (TCP or UDP)	2

Table 55: IPv6 Flow Key (Continued)

Field	Field size in bytes
Virtual routing and forwarding table (VRF) identifier	2

Figure 39: Sample IPFIX Data Template for IPv4 Traffic

```

Version: 10
Length: 76
▼ Timestamp: Jan 24, 2022 06:18:29.000000000 Pacific Standard Time
FlowSequence: 1159
Observation Domain Id: 167837696
▼ Set 1 [id=2] (Data Template): 65000
  Flowset Id: Data Template (V10 [IPFIX] (2))
  FlowSet Length: 60
  ▼ Template (Id = 65000, Count = 13)
    Template Id: 65000
    Field Count: 13
    ▼ Field (1/13): IP_SRC_ADDR
    ▼ Field (2/13): IP_DST_ADDR
    ▼ Field (3/13): PROTOCOL
    ▼ Field (4/13): L4_SRC_PORT
    ▼ Field (5/13): L4_DST_PORT
    ▼ Field (6/13): ingressVRFID
    ▼ Field (7/13): layer2SegmentId
    ▼ Field (8/13): INPUT_SNMP
    ▼ Field (9/13): OUTPUT_SNMP
    ▼ Field (10/13): flowStartSeconds
    ▼ Field (11/13): flowEndSeconds
    ▼ Field (12/13): PKTS
    ▼ Field (13/13): BYTES

```

jn-000342

Limitations and Caveats

- FBT for VXLANs is supported only on Junos OS.
- Only IRB interfaces are supported. For EVPN-VXLAN networks with CRB overlays, you can only monitor the IRB interfaces on the spine. For EVPN-VXLAN networks with ERB overlays, you can only monitor the IRB interfaces on the leaves.
- Only one inline-monitoring instance and one collector are supported.
- The collector must be reachable through a network interface, not only through a management or loopback interface.
- You cannot configure an option template identifier or a forwarding class.
- The IPFIX Option Data Record and IPFIX Option Data Template are not supported.

- Flow learning and tracking is based on client traffic data only, not the outer tunnel header. Flow learning is software-based and takes up to 10 seconds per flow.
- Counters are not active until the software learns the flow and installs the flow in the flow table.
- The software does not use the TCP FIN/RST flag for flow aging.
- The software requires a layer 3 header in the packet, and supports only the TCP and UDP protocols.
- The reported egress port might not be correct with LAG, ECMP, broadcast, multicast, or unknown traffic, if the egress port is in a different VRF.

Configure FBT for VXLANs (QFX5120)

You can configure flow-based telemetry (FBT) for VXLANs for the QFX5120 -32C and QFX5120-48y-8c switches. FBT for VXLANs enables inline telemetry data for VXLANs that have either centrally-routed bridging (CRB) or edge-routed bridging (ERB) overlays. FBT for VXLANs enables per-flow-level analytics for IRB interfaces, using inline monitoring services to create flows, collect them, and export them to a collector. With inline monitoring services, you can monitor every IPv4 and IPv6 packet on both ingress and egress directions of an interface.

Before you can configure FBT for VXLANs, you must first enable software-based IPFIX flow export and must allocate exact-match memory in the unified forwarding table to learn the flows. To configure:

```
user@host# set system packet-forwarding-options ipfix-sw-mode
user@host# set chassis forwarding-options em-hw-profile
user@host# commit
```

After you commit the configuration, the system then prompts you to reboot the system.

To configure FBT for VXLANs:

1. Define the IPFIX template.

To configure attributes of the template:

```
user@host# set services inline-monitoring template template_1 flow-inactive-timeout seconds
user@host# set services inline-monitoring template template_1 template-refresh-rate template-  
refresh-rate
user@host# set services inline-monitoring template template_1 template-identifier template-  
identifier
user@host# set services inline-monitoring template template_1 template-type (ipv4-template |  
ipv6-template)
```

In this example, the inactive-flow timeout period is set to 10 seconds, the template refresh rate is set to 30 seconds, you've configured a template identifier, and you're using the IPv4 template:

```
user@host# set services inline-monitoring template template_1 flow-inactive-timeout 10
user@host# set services inline-monitoring template template_1 template-refresh-rate 10
user@host# set services inline-monitoring template template_1 template-identifier 1200
user@host# set services inline-monitoring template template_1 template-type ipv4-template
```

2. Attach a template to the instance and describe the collector.

FBT for VXLANs only supports IPv4 addresses for the collector. To configure the instance and collector:

```
user@host# set services inline-monitoring instance instance-name template-name template-name
user@host# set services inline-monitoring instance instance-name collector collector-name
source-address (IPv4-address)
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-address (IPv4-address)
user@host# set services inline-monitoring instance instance-name collector collector-name
dscp dscp-bits
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-port port
```

In this example, you create a template with the name `template_1`, create an inline-monitoring instance `i1`, and create the configuration for the collector `c2` using IPv4 addresses:

```
user@host# set services inline-monitoring instance i1 template-name template_1
user@host# set services inline-monitoring instance i1 collector c2 source-address 10.11.12.1
user@host# set services inline-monitoring instance i1 collector c2 destination-address
10.11.12.2
user@host# set services inline-monitoring instance i1 collector c2 dscp 21
user@host# set services inline-monitoring instance i1 collector c2 destination-port 2055
```

3. Create a firewall filter and configure the action `inline-monitoring-instance`.

To configure the firewall filter:

```
user@host# set firewall family inet filter filter-name term term-name from source-address
(IPv4-source-address | IPv6-source-address)
user@host# set firewall family inet filter filter-name term term-name from destination-
address (IPv4-destination-address | IPv6-destination-address)
user@host# set firewall family inet filter filter-name term term-name then inline-monitoring-
```

```
instance instance-name
user@host# set firewall family inet filter filter-name term term-name then accept
```

In this example, you configure an IPv4 firewall filter named `ipv4_ingress`, with the term name `rule1` containing the action `inline-monitoring-instance`, and the inline monitoring instance `i1` is mapped to it:

```
user@host# set firewall family inet filter ipv4_ingress term rule1 from source-address
10.11.12.1
user@host# set firewall family inet filter ipv4_ingress term rule1 from destination-address
10.11.12.2
user@host# set firewall family inet filter ipv4_ingress term rule1 then inline-monitoring-
instance i1
user@host# set firewall family inet filter ipv4_ingress term rule1 then accept
```

- 4. Map the firewall filter to the family under the logical unit of the already-configured interface to apply inline monitoring in the ingress direction.

To map the firewall filter:

```
user@host# set interfaces irb unit unit-number family inet filter input filter-name
```

In this example, you map the `ipv4_ingress` firewall filter to the `inet` family of unit 100:

```
user@host# set interface irb unit 100 family inet filter input ipv4_ingress
```

- 5. Commit the configuration.
- 6. Monitor inline-monitoring statistics with the `show services inline-monitoring statistics fpc-slot slot-number` command.

Release History Table

Release	Description
22.2R1	You can configure flow-based telemetry (FBT) for VXLANs for the QFX5120 -32C and QFX5120-48y-8c switches. FBT for VXLANs enables inline telemetry data for VXLANs that have either centrally-routed bridging (CRB) or edge-routed bridging (ERB) overlays. FBT for VXLANs enables per-flow-level analytics for IRB interfaces, using inline monitoring services to create flows, collect them, and export them to a collector.

RELATED DOCUMENTATION

| [Inline Monitoring Services Configuration](#) | 334

CHAPTER 11

Inband Flow Analyzer 2.0

IN THIS CHAPTER

- [Inband Flow Analyzer \(IFA\) 2.0 Probe for Real-Time Flow Monitoring | 370](#)

Inband Flow Analyzer (IFA) 2.0 Probe for Real-Time Flow Monitoring

SUMMARY

Inband Flow Analyzer (IFA) 2.0 collects data on a per-hop basis across the network. You export this data to external collectors to perform localized or end-to-end analytics.

IN THIS SECTION

- [Inband Flow Analyzer 2.0 | 370](#)
- [Configure Inband Flow Analyzer 2.0 | 384](#)
- [Example - Configure Inband Flow Analyzer 2.0 for Traffic Monitoring | 394](#)

Inband Flow Analyzer 2.0

IN THIS SECTION

- [Inband Flow Analyzer 2.0 Overview | 371](#)
- [Benefits | 371](#)
- [Inband Flow Analyzer Process | 372](#)
- [IFA Probe Packet Headers | 373](#)
- [Tailstamps for IFA Probe Packets \(QFX5220 only\) | 378](#)
- [Supported Features on IFA Nodes | 379](#)
- [Limitations of IFA 2.0 Configuration | 381](#)
- [Usage Considerations | 383](#)

Inband Flow Analyzer 2.0 Overview

Inband Flow Analyzer 2.0 (IFA 2.0) is a feature that you can use to monitor and analyze packets as they enter and exit the network. As the network administrator, you can use this feature to collect data related to the paths the packets take through the network and how long the packets spend at each hop. This data provides an indication of excessive latency and possible congestion. This feature helps you to get insights about complex networks by collecting per-hop flow data on the data plane.

IFA uses probe packets to collect network-wide flow data. IFA samples the flow of interest and generates probe packets. These packets are representative of the original flow, possessing the same characteristics as the original flow. This means that IFA packets traverse the same path in the network and the same queues in the networking element as the original packet would. As a result, IFA probe packets traverse the same network path as the original flow, experiencing similar latency and congestion.

You can use Inband Flow Analyzer 2.0 (IFA 2.0) to collect flow data information such as:

- Residence time (latency)
- Per-hop latency
- Per-hop ingress port number
- Per-hop egress port number
- Received packet timestamp value
- Queue ID
- Congestion notification
- Egress port speed

IFA 2.0 is defined in the IETF draft titled [Inband Flow Analyzer, draft-kumar-ippm-ifa-02](#).

Benefits

- IFA probe packets traverse the same network path as the original flow, helping you to monitor the network for faults and performance issues.
- Monitors live traffic and thus helps to perform packet-level latency analysis and queue-congestion monitoring to optimize the network performance.


Inband Flow Analyzer Process

IFA uses the following processing nodes (as shown in [Figure 40 on page 372](#)) to monitor and analyze flows:

- IFA initiator node (also known as ingress node)
- IFA transit node
- IFA terminating node (also known as egress node)

IFA 2.0 supports processing both Layer 3 (L3) and VXLAN flows, but you can't configure IFA for both L3 and VXLAN flows on the same device. The flow-type options are mutually exclusive. You use the `flow-type` configuration statement to set the flow type of interest —either L3 or VXLAN. You configure the `flow-type` statement only for the IFA initiator and IFA terminating nodes (generally leaf nodes). For an IFA transit node (generally a spine node), you don't need to configure the `flow-type` statement.

Figure 40: IFA Processing



[Table 56 on page 372](#) summarizes the different functions that the IFA processing nodes perform:

Table 56: IFA Node Functions

IFA Node	Function
IFA initiator node	Samples the flow traffic of interest (L3 or VXLAN) and creates an IFA copy by adding an IFA header to each sample.

Table 56: IFA Node Functions (*Continued*)

IFA Node	Function
IFA transit node	<p>Identifies IFA packets and appends their metadata to the metadata stack in the packet.</p> <ul style="list-style-type: none"> • If any packet comes with an IFA header, the node inserts the metadata into the metadata stack and forwards it. If the hop limit is 0, the node does not insert the metadata. • When a non-IFA device receives an IFA packet, the device forwards it without IFA processing. • The QFX5220 as an IFA transit node can not insert metadata into the metadata stack of the IFA probe packet header. Instead, the QFX5220 adds a timestamp to the end of the IFA probe packet that includes timestamps and other metadata.
IFA terminating node	<ul style="list-style-type: none"> • Inserts terminating node metadata into an IFA packet. • Formats the IFA packets in IP Flow Information Export (IPFIX) format and sends the packets to the configured collector. You can use any collector (or application) that supports the IPFIX format. <p>NOTE: IFA terminating functionality requires a valid Juniper IFA license.</p>

IFA Probe Packet Headers

An IFA 2.0 probe packet contains the following:

- IFA Header
- IFA Metadata Header
- IFA Metadata Stack

Figure 41 on page 374 shows the L3 IFA 2.0 packet format at the IFA initiator node:

Figure 41: Layer 3 IFA 2.0 Packet Format at the IFA Initiator Node


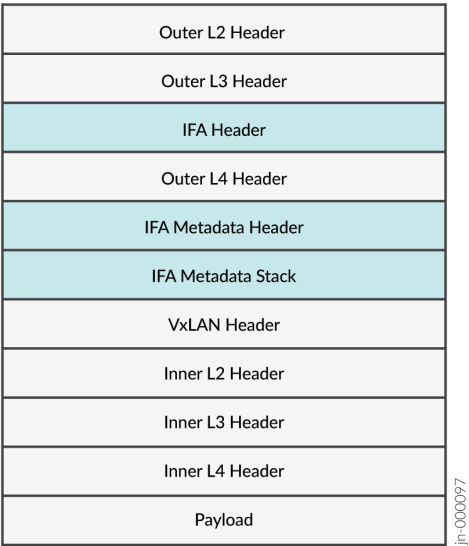


Figure 42 on page 374 shows the VXLAN IFA 2.0 packet format at the IFA initiator node.

Figure 42: VXLAN IFA 2.0 Packet Format at the IFA Initiator Node




NOTE: When VXLAN is used, then the IFA headers are added after VXLAN encapsulation using a three-pass mechanism.

IFA Header

IFA 2.0 defines an upper layer header (ULH), similar to how TCP, UDP, Generic Routing Encapsulation (GRE), and Spanning Tree Protocol (STP) define a ULH. The IFA ULH is always the first header after the IP header, even if there are some other IPv4 extension headers. The NextHdr field (that is, the Protocol Type field in the IFA header) carries the original IP header protocol field value. [Figure 43 on page 375](#) shows the IFA header format.

Figure 43: IFA Header



[Table 57 on page 375](#) provides details about the IFA header fields.
Table 57: IFA Header Fields

IFA Header Field	Description
IFA Version	Version of the IFA header. In the current implementation, the IFA version is 2.0.
GNS	Global namespace (GNS) for IFA metadata. The IFA initiator node sets the value for this field as 0xF.
Protocol Type	IP header protocol type. This value is copied from the IP header.
FLAGS	Unused.
MAX Length	Maximum allowed length of the metadata stack in multiples of four octets. The initiator node initializes this field. Each node in the path compares the current length with the maximum length. If the current length equals or exceeds the maximum length, the transit node stops inserting metadata. You can configure this maximum allowed length. The default value is 240 octets (for 30 hops).

IFA Metadata Header

IFA 2.0 defines a compact 4-byte metadata header as shown in [Figure 44 on page 376](#). The IFA initiator node adds this header to the probe packet.

Figure 44: IFA Metadata Header Format



[Table 58 on page 376](#) provides details about the IFA metadata header fields.
Table 58: IFA Metadata Header Fields

IFA Metadata Header Field	Description
Request Vector	Specifies the presence of fields as specified by the GNS. Unused.
Action Vector	Specifies the node-local or the end-to-end action on the IFA packets. Unused.
Hop Limit	Specifies the maximum number of allowed hops in an IFA zone. The initiator node initializes this field. The hop limit is decremented at each hop. If the hop limit of the incoming packet is 0, the current node does not insert metadata. You can configure this limit. The default value is 250. The terminating node does not perform the hop limit check.
Current Length	Specifies the current length of the metadata stack in multiples of 4 octets.

IFA Metadata Stack

Each IFA hop inserts hop-specific metadata into an IFA metadata stack as shown in [Figure 45 on page 377](#). The IFA initiator node adds the metadata header after the L4 header.

The QFX5220 as a transit node can not insert metadata into the metadata stack of the IFA probe packet header. Instead, the QFX5220 adds a timestamp to the end of the IFA probe packet that includes

timestamps and other metadata. For more information about these timestamps, see ["Timestamps for IFA Probe Packets \(QFX5220 only\)" on page 378](#).

Figure 45: IFA Metadata Stack Header




Table 59 on page 377 provides details about the IFA metadata stack header fields.

Table 59: IFA Metadata Stack Header Fields

IFA Metadata Stack Header Field	Description
LNS	Local namespace. You must set the LNS value to 1.
Device ID	User-configurable device ID. You can explicitly configure the device ID or configure the auto statement. If you configure auto, the device ID is internally generated from the router ID or the management IP address.
IP TTL	IP time-to-live (TTL) value at each hop.

Table 59: IFA Metadata Stack Header Fields *(Continued)*

IFA Metadata Stack Header Field	Description
Egress Port Speed	Encodings are 0–10Gbps, 1–25Gbps, 2–40Gbps, 3–50Gbps, 4–100Gbps, 5–200Gbps, 6–400Gbps. Egress port speed is mapped with IFA metadata. For example, when a egress port speed is 10Gbps, then the speed field of IFA packet is set to 0.
Congestion	Indicates whether the packet has experienced congestion. You must enable an explicit congestion notification (ECN) on the egress port.
Queue ID	Egress port queue ID.
Rx Timestamp Seconds	Received packet timestamp value (in seconds). It is the collector's responsibility to retrieve time-of-day (ToD) from these 20-bit values. 20-bit seconds will wrap around every 12 days. Collector has to periodically sync up ToD within the wraparound time and use it along with 20-bit from metadata to derive the 32-bit Rx Timestamp Seconds value.
Egress Port Number	Egress hardware (ASIC) port number.
Ingress Port Number	Ingress hardware port number.
Rx Timestamp Nano Seconds	Received timestamp value in nanoseconds.
Residence Time Nano Seconds	Per-hop latency in nanoseconds. For the QFX5120, the residence time is calculated as $0x3B9ACA00$ (1 second in nanoseconds) + TX_NSEC - RX_NSEC. (An extra second is added to every packet to avoid wraparound handling.) In contrast, for the QFX5130, QFX5220, and QFX5700, the residence time is updated as the actual value.

Tailstamps for IFA Probe Packets (QFX5220 only)

The QFX5220 as a transit node can not insert metadata into the metadata stack of the IFA probe packet header. Instead, the QFX5220 adds a tailstamp to the end of the IFA probe packet that includes

timestamps and other metadata. The QFX5220 adds a total of 28 bytes of metadata as a tailstamp. Upon receiving the IFA probe packet, the IFA termination node uses the TTL value in the metadata to identify the number of tailstamps (that is, the number of QFX5220 hops on the path between two QFX5120 or QFX5130 devices). Then the tailstamps are converted into the correct metadata format and inserted into the correct place in the metadata stack, so that the metadata appears in the order that the transit nodes added them. Once complete, the IFA termination node exports the data in IPFIX format to the configured external collector.

Due to this inability to insert metadata into the stack, the IFA metadata stack fields IP TTL , Egress Port Speed and Congestion for the QFX5220 are received with the value of 0 at the collector. You must configure the collector to ignore these unsupported fields from the QFX5220.

The tailstamp includes 14 bytes of ingress (Rx) tailstamp and 14 bytes of egress (Tx) tailstamp. [Figure 46 on page 379](#) and [Figure 47 on page 379](#) provide details about the format of these timestamps.

Figure 46: Ingress (Rx) Tailstamp Format

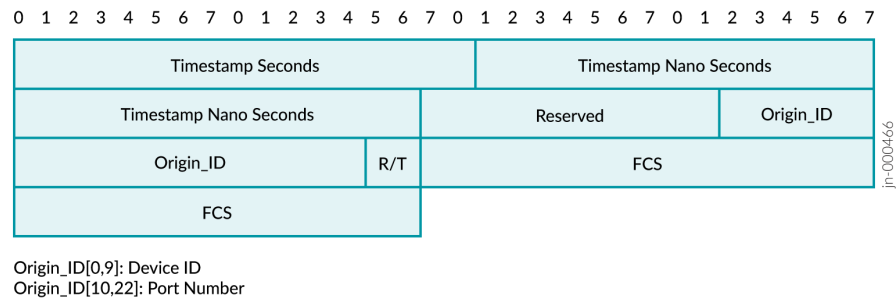



Figure 47: Egress (Tx) Tailstamp Format



Supported Features on IFA Nodes

[Table 60 on page 380](#) lists the features supported by IFA nodes.

Table 60: Supported Features on IFA Nodes

IFA Node	Supported Features
IFA initiator	<p>Traffic and interface types:</p> <ul style="list-style-type: none"> • IPv4 and IPv6 traffic. • VXLAN traffic. • UDP and TCP. • Tagged and untagged packets. • Aggregation links (LAG) and multichassis LAG (MC-LAG). In case of LAG egress, the original packet and the IFA probe copy use the same port to exit. • IRB interfaces. • ECMP traffic. In case of ECMP traffic, the original packet and the IFA probe copy use the same port to exit. • Interface speeds, such as 10 Gbps, 25 Gbps, 40 Gbps, 50 Gbps, and 100 Gbps.
IFA transit	Identifies IFA packets, appends their metadata, and forwards it.
IFA terminating	<ul style="list-style-type: none"> • Support to export IFA data to any configured IPv4 collector in IPFIX format. • Support to combine multiple IFA packets into a single IPFIX export.

Supported IFA 2.0 IPFIX Format (Terminating Node)

The terminating node formats the IFA 2.0 packets in IPFIX format, updates the egress port information, and sends the packet to the configured collector. The IFA 2.0 IPFIX template is the same for L3 traffic

and VXLAN traffic. [Figure 48 on page 381](#) shows the IPFIX template in which the terminating node formats the IFA 2.0 data and sends it to a collector.

Figure 48: IFA 2.0 IPFIX Template

```

Version: 10
Length: 48
> Timestamp: Jan 11, 1970 19:58:15.000000000 IST
FlowSequence: 1128891
Observation Domain Id: 305419896
< Set 1 [id=2] (Data Template): 257
  FlowSet Id: Data Template (V10 [IPFIX]) (2)
  FlowSet Length: 32
  < Template (Id = 257, Count = 3)
    Template Id: 257
    Field Count: 3
    < Field (1/3): 100 [pen: Reserved]
      1... .... = Pen provided: Yes
      .000 0000 0110 0100 = Type: 100 [pen: Reserved]
      Length: 8
      PEN: Reserved (0)
    < Field (2/3): 101 [pen: Reserved]
      1... .... = Pen provided: Yes
      .000 0000 0110 0101 = Type: 101 [pen: Reserved]
      Length: 65535 [i.e.: "Variable Length"]
      PEN: Reserved (0)
    < Field (3/3): 102 [pen: Reserved]
      1... .... = Pen provided: Yes
      .000 0000 0110 0110 = Type: 102 [pen: Reserved]
      Length: 65535 [i.e.: "Variable Length"]
      PEN: Reserved (0)

```

[Figure 49 on page 381](#) shows a sample VXLAN IFA 2.0 packet received by the configured collector in IPFIX format.

Figure 49: VXLAN IFA 2.0 IPFIX Sample Packet

```

Version: 10
Length: 274
> Timestamp: Jan 11, 1970 19:58:12.000000000 IST
FlowSequence: 1125782
Observation Domain Id: 305419896
< Set 1 [id=257] (1 flows)
  FlowSet Id: (Data) (257)
  FlowSet Length: 258
  [Template Frame: 330 (received after this frame)]
  < Flow 1
    Enterprise Private entry: (Reserved) Type 100: Value (hex bytes): 2f 11 00 f0 00 00 f9 10
    < Enterprise Private entry: (Reserved) Type 101: Value (hex bytes): 10 07 d0 40 0a fc 21 5a 00 01 00 1f 14 d8 f6 80 ...
      String_len_short: 255
      String_len_short: 64
    < Enterprise Private entry: (Reserved) Type 102: Value (hex bytes): 54 4b 8c 1a 05 95 54 4b 8c 19 e7 95 81 00 0f ff ...
      String_len_short: 178

```

Limitations of IFA 2.0 Configuration

Before you configure IFA 2.0 on a device running Junos OS, you must be aware of the following limitations:

- **Protocol Number**—IFA 2.0 uses the experimental protocol number 253. If the switch receives any traffic with protocol number 253, those packets will hit the IFA transit filter. In this case the QFX5220 adds a 28-byte timestamp to those packets. For the QFX5130 and QFX5700 switches, even

though the packets hit the filter, IFA metadata is not added to the packets. However, the IFA transit statistics do increment.

- **Filter Resource Allocation**—If filter hardware resources are already exhausted in the system, the IFA feature does not work because it needs filter resources. You can monitor the system log (syslog) for filter space exhaustion errors.
- **Layer 2 and BUM Traffic**—IFA 2.0 is not supported on Layer 2 switched traffic and broadcast, unknown unicast, and multicast (BUM) traffic.
- **IFA Layer 3 and VXLAN Flows**
 - IFA 2.0 supports processing both L3 and VXLAN flows, but you can't configure IFA for both L3 and VXLAN flows on the same device. The `flow-type` options are mutually exclusive. You use the `flow-type` configuration statement to set the flow type of interest —either L3 or VXLAN. This restriction is only applicable for IFA initiator and terminating nodes (generally leaf nodes). For IFA transit nodes (generally spine nodes), it is not required to configure the flow type.
 - For VXLAN IFA flow, the egress port-related metadata for the terminating node (including egress port number, speed, queue ID, and congestion) are incorrect. It is recommended that you ignore the termination node egress-port-related metadata for VXLAN flows.
 - An IFA flow-type (L3 or VXLAN) change requires IFA filter removal and reconfiguration. In case of a flow-type mismatch (for example, `flow-type` configured as VXLAN, whereas the incoming traffic is L3 or vice versa), we can't guarantee IFA behavior (IFA probe packets could be initiated with invalid fields).
- **IFA Initiator Node**
 - L4 header (UDP/TCP) is mandatory for IFA initiation.
 - IFA initiation for VXLAN flow does not work if the egress port is configured to function as a link aggregation group (LAG) (links connecting leaf to spine).
 - You cannot configure different sample rates for different flows on a port for an IFA initiator. All flows within a port should have the same sample rate.
- **IFA Transit Nodes**—Devices running Junos OS and Junos OS Evolved do not support the maximum length check for the metadata stack. Configure the `hop-limit` option to limit the insertion of metadata on transit nodes. The QFX5220 cannot perform the hop-limit check to insert the timestamp. The QFX5220 also cannot insert metadata into the metadata stack in the IFA probe packet header; instead, the QFX 5220 appends a timestamp to the end of the IFA probe packet.

QFX5220 supports only 18 bits for the Rx Seconds Timestamp value. The QFX5130 and QFX5700 support a 20-bit Rx Seconds Timestamp value.

The Residence Time Nano Seconds field is updated as the actual value on the QFX5220, QFX5130, and QFX5700 transit nodes, but on the QFX5120 transit node, 1 second (1000000000 ns) is added along with the actual residence time.

- **IFA Terminating Node**

- You can configure only a single IPv4 collector at the terminating node.
- The terminating node metadata has the queue ID 47. This queue ID is reserved for IFA packet export.
- The terminating node does not perform a hop-limit check. Even if the incoming IFA packet has hop-limit set to 0, the terminating node inserts the metadata and reduces the hop limit by 1, which resets the hop-limit value to 255.

Usage Considerations

Following are the IFA 2.0 related usage considerations:

- Sampled IFA packets have an additional 40 bytes (4-byte IFA header + 4-byte IFA metadata header + 32-byte metadata) when it egresses on the initiator node. On subsequent IFA nodes, 32-byte IFA metadata is inserted at every hop. Due to insertion of per-hop metadata into IFA packets, the packet size grows after every hop. You must configure the interface's maximum transmission unit (MTU) accordingly along the network path. In case of an IFA zone with a large number of transit nodes, you must take care of the MTU. Alternatively, you can configure the hop-limit option at the initiator node to ensure that the size of the IFA packets never exceeds the specified MTU value.
- To select the flow of interest, you can use any combination of source IP address, destination IP address, source port, destination port, and protocol match qualifiers. IFA 2.0 doesn't support any other match qualifiers.
- You must configure a unique device ID for each hop within an IFA zone. If you've configured the auto option for the device ID, then the device ID is generated from the last 20 bits of the router ID or management IP address.
- If you've configured the sampling rate as aggressive, the egress ports might experience congestion due to more IFA copies. This port congestion could create congestion on terminating nodes when IFA copies are sent to the chip processor for IPFIX export. We recommend that you select the sampling rate accordingly.
- When you configure an IFA 2.0 initiator, an internal mirror session is created for the loopback port. As a result, the number of user-configurable mirror sessions reduces from 4 to 3.
- The terminating node accepts an IFA packet size up to 9000 bytes (including IFA headers). On the terminating node, multiple IFA received packets are combined into a single IPFIX export packet. You can combine a maximum of 10 IFA records in a single IPFIX export packet. By default, a maximum of

256 bytes of the original flow packet are exported as part of the IPFIX export, along with IFA headers. The maximum size of a single IPFIX packet is 9000 bytes. You must configure the MTU properly on the collector port. Because the maximum size of a single IPFIX packet is 9000 bytes, the maximum clip length for the IPFIX packet is equal to or less than: 9000 bytes - (IFA header length + IFA metadata header length + IFA metadata stack length).

- We recommend that you use only IFA-aware (supported) devices within the IFA zone. We cannot guarantee proper IFA behavior with IFA-unaware devices.

Configure Inband Flow Analyzer 2.0

IN THIS SECTION

- [Configure IFA Initiator Node | 388](#)
- [Configure IFA Transit Node | 391](#)
- [Configure IFA Terminating Node | 391](#)
- [View Inband Flow Analyzer Statistics | 393](#)

IFA is a type of Inband Network Telemetry (INT) that allows you to collect information about the network state by the data plane.

To configure IFA 2.0 for monitoring the network for faults, performance issues, and collect the data for analysis, you need to configure the IFA roles first. You can configure the IFA roles on a Junos OS device that supports IFA feature. The following QFX switches support the IFA 2.0 feature:

- QFX5120-32C, QFX5120-48Y, QFX5120-48T, and QFX5120-48YM, running Junos OS
- QFX5130-32CD, running Junos OS Evolved (transit node role only)
- QFX5220-32CD and QFX5220-128C, running Junos OS Evolved (transit node role only)
- QFX5700, running Junos OS Evolved (transit node role only)

See the release history table at the end of this topic for information on when devices were first supported in Junos OS.

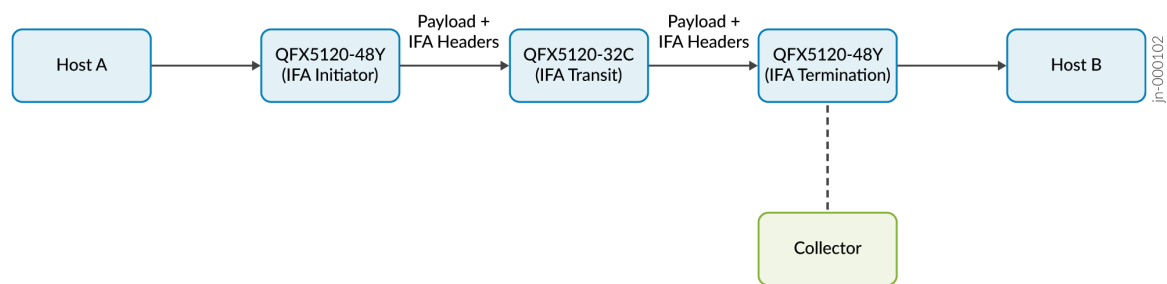
Following are some of the guidelines for configuring a Junos OS device for an IFA role:

- You can use the same model switches or different switches to play the IFA roles (initiator, transit, terminating) for a particular IFA flow.
- You can use the same device to perform all three different IFA roles for different flows.

- In an IFA flow, the transit IFA role is optional.

Figure 50 on page 385 illustrates a sample scenario for configuring IFA nodes on Junos OS devices. In this scenario, different Junos OS devices that support the IFA feature play different IFA roles in a single IFA flow.

Figure 50: Sample Inband Flow Analyzer Scenario



Following are some of the guidelines for configuring IFA nodes:

- You can enable the IFA configuration on the interface only through the firewall filter configuration.
- You can apply IFA filter only on ingress direction on the port.

Table 61 on page 385 summarizes the configurations for IFA initiator, transit, and terminating nodes.

Table 61: IFA Configurations for IFA Roles

IFA Configuration Parameter	Configuration Statement	IFA Role
(Mandatory) Configure Device ID	<code>user@host# set services inband-flow-telemetry device-id (<1 - 1048575> auto)</code>	Mandatory configuration for IFA initiator, transit, and terminating nodes.
(Optional, QFX5120-48YM or QFX5220 only) Configure a more accurate clock source	<code>user@host# set services inband-flow-telemetry clock-source (ntp ptp)</code>	IFA initiator, transit, and terminating nodes.

Table 61: IFA Configurations for IFA Roles *(Continued)*

IFA Configuration Parameter	Configuration Statement	IFA Role
(Optional) IFA maximum metadata stack length	<pre>user@host# set services inband- flow-telemetry meta-data-stack- length <8 - 255></pre> <p>Default value : 240 (for 30 hops)</p>	IFA initiator node
(Optional) IFA maximum hop limit	<pre>user@host# set services inband- flow-telemetry hop-limit <1 - 250></pre> <p>Default value : 250</p>	IFA initiator node
(Optional) No IPv6 address match	<pre>user@host# set services inband- flow-telemetry no-ipv6-address- match</pre>	IFA initiator/terminating node
(Mandatory) IFA flow type	<pre>user@host# set services inband- flow-telemetry flow-type (13 vxlan)</pre>	Mandatory configuration for IFA initiator and terminating node. This configuration is not required for IFA transit node.
IFA sampling	<pre>user@host# set services inband- flow-telemetry profile ifa- profile-name sample-rate <1-16777215></pre>	IFA initiator node

Table 61: IFA Configurations for IFA Roles *(Continued)*

IFA Configuration Parameter	Configuration Statement	IFA Role
Collector information	<pre> user@host# set services inband- flow-telemetry profile <i>ifa- profile-name</i> collector source- address <i>IP-address</i> user@host# set services inband- flow-telemetry profile <i>ifa- profile-name</i> collector destination-address <i>IP-address</i> user@host# set services inband- flow-telemetry profile <i>ifa- profile-name</i> collector destination-port <i>port-number</i> user@host# set services inband- flow-telemetry profile <i>ifa- profile-name</i> collector maximum- clip-length <i>length</i> user@host# set services inband- flow-telemetry profile <i>ifa- profile-name</i> collector mtu <i>size</i> </pre>	IFA terminating node
IFA filter for L3 flow	<p>For example:</p> <pre> user@host# set firewall family inet filter f1 term t1 from <i>match- condition</i> user@host# set firewall family inet filter f1 term t1 then inband-flow-telemetry-init p1 user@host# set firewall family inet filter f1 term t2 from <i>match- condition</i> user@host# set firewall family inet filter f1 term t2 then inband-flow-telemetry-terminate p2 user@host# set interfaces (<i>interface-name</i> <i>wildcard</i>) unit 0 family inet filter input f1 </pre>	IFA initiator/terminating node

Table 61: IFA Configurations for IFA Roles *(Continued)*

IFA Configuration Parameter	Configuration Statement	IFA Role
IFA filter for VXLAN flow	<p>For example:</p> <pre> user@host# set firewall family ethernet-switching filter f1 term term1 from match-condition user@host# set firewall family ethernet-switching filter f1 term t1 then inband-flow-telemetry- init p1 user@host# set firewall family ethernet-switching filter f1 term t2 from match-condition user@host# set firewall family ethernet-switching filter f1 term t2 then inband-flow-telemetry- terminate p2 user@host# set interfaces (interface-name wildcard) unit 0 family ethernet-switching filter input f1 </pre>	IFA initiator/terminating node

Configure IFA Initiator Node

To configure your device as IFA 2.0 initiator:

1. Configure the device ID. You can also configure the value auto for device-id. If the device-id is configured as auto, the device-id is internally generated from the router ID or the management IP address.

```
user@host# set services inband-flow-telemetry device-id (id-number | auto)
```

In this example, the device id for IFA initiator node is configured as 10000.

```
user@host# set services inband-flow-telemetry device-id 10000
```

2. Configure the flow type. You can configure either of two flow types, l3 or vxlan. You cannot configure L3 and VXLAN flows together in the same device.

```
user@host# set services inband-flow-telemetry flow-type (l3 | vxlan)
```

In this example, the flow type is configured as l3. If you configure l3 flow-type in the initiator node, then you must choose l3 flow-type for the terminating node also.

```
user@host# set services inband-flow-telemetry flow-type l3
```

3. (Optional) Configure the maximum metadata stack length. Each IFA hop inserts hop-specific metadata into the IFA metadata stack.

```
user@host# set services inband-flow-telemetry meta-data-stack-length value
```

In this example, the metadata stack length is configured as 80.

```
user@host# set services inband-flow-telemetry meta-data-stack-length 80
```

4. Configure the hop limit.

```
user@host# set services inband-flow-telemetry hop-limit value
```

In this example, hop-limit is configured as 10. The hop limit is decremented at each hop. If the incoming hop limit is 0, the current node does not insert metadata.

```
user@host# set services inband-flow-telemetry hop-limit 10
```

5. Configure IFA sampling. The sampling rate is the average number of samples obtained in one second. You cannot have different sample rate for different flows on an IFA initiator node enabled on a port. All flows within a port should have same sample rate.

```
user@host# set services inband-flow-telemetry profile ifa-profile-name sample-rate value
```


In this example, the sample rate is configured as 1000; meaning out of 1000 packets, 1 packet will be sampled per second.

```
user@host# set services inband-flow-telemetry profile p1 sample-rate 1000
```

6. Configure IFA firewall filters. You can configure firewall filter with any of the below match conditions:

- Source IP address
- Destination IP address
- Source port
- Destination port
- Protocol

Create a firewall and configure the action `inband-flow-telemetry-init`.

```
user@host# set firewall family inet filter filter-name term term-name from source-address
ipv4-address
user@host# set firewall family inet filter filter-name term term-name from destination-
address ipv4-address
user@host# set firewall family inet filter filter-name term term-name then inband-flow-
telemetry-init ifa-profile-name
```

In this example, you configure a firewall filter named `f1`, with the term name `t1` containing the action `inband-flow-telemetry-init`, and the inband flow telemetry initiator profile `p1` mapped to it:

```
user@host# set firewall family inet filter f1 term t1 from source-address 10.30.1.4/32
user@host# set firewall family inet filter f1 term t1 from destination-address 10.40.1.4/32
user@host# set firewall family inet filter f1 term t1 then inband-flow-telemetry-init p1
```

7. Map the firewall filter to the family under the logical unit of the already-configured interface to apply the action `inband-flow-telemetry-init` in the ingress direction.

To map the firewall filter:

```
user@host# set interfaces interface-name unit 0 family inet filter input filter-name
```

In this example, you map the f1 firewall filter to the `inet` family of logical interface 0 of the physical interface `et-0/0/0`:

```
user@host# set interfaces et-0/0/0 unit 0 family inet filter input f1
```

Configure IFA Transit Node

To configure your device as IFA transit node:

Configure the device ID. You can also configure the value `auto` for `device-id`. If the `device-id` is configured as `auto`, then the `device-id` is internally generated from the router ID or the management IP address.

```
user@host# set services inband-flow-telemetry device-id (id-number | auto)
```

For example:

```
user@host# set services inband-flow-telemetry device-id 10001
```

Configure IFA Terminating Node

To configure your device as IFA terminating node:

1. Configure the device ID. You can also configure the value `auto` for `device-id`. If the `device-id` is configured as `auto`, then the `device-id` is internally generated from the router ID or the management IP address.

```
user@host# set services inband-flow-telemetry device-id (id-number | auto)
```

For example:

```
user@host# set services inband-flow-telemetry device-id 10002
```

2. Configure the flow type. You can configure either of two flow types, `l3` or `vxlan`. You cannot configure L3 and VXLAN flows together in the same device.

```
user@host# set services inband-flow-telemetry flow-type (l3 | vxlan)
```

If you configure 13 flow-type in the initiator node, then you must choose 13 flow-type for the terminating node also.

```
user@host# set services inband-flow-telemetry flow-type 13
```

3. Configure IFA profile with the collector information for the terminating node.

```
user@host#
user@host# set services inband-flow-telemetry profile ifa-profile-name collector source-
address ipv4-address
user@host# set services inband-flow-telemetry profile ifa-profile-name collector destination-
address ipv4-address
user@host# set services inband-flow-telemetry profile ifa-profile-name collector destination-
port port-number
```

For example:

```
user@host# set services inband-flow-telemetry profile p2 collector source-address 10.50.1.1
user@host# set services inband-flow-telemetry profile p2 collector destination-address
10.60.1.1
user@host# set services inband-flow-telemetry profile p2 collector destination-port 2055
```

4. You can configure firewall filter with any of the below match conditions:

- Source IP address
- Destination IP address
- Source port
- Destination port
- Protocol

Create a firewall and configure the action inband-flow-telemetry-terminate.

```
user@host# set firewall family inet filter filter-name term term-name from source-address
ipv4-address
user@host# set firewall family inet filter filter-name term term-name from destination-
address ipv4-address
user@host# set firewall family inet filter filter-name term term-name then inband-flow-
telemetry-terminate ifa-profile-name
```

```
user@host# set firewall interfaces interfaces-name unit logical-unit-number family inet
filter input filter-name
```

In this example, you configure a firewall filter named f2, with the term name t1 containing the action inband-flow-telemetry-terminate, and the inband-flow-telemetry-terminate profile p2 mapped to it:

```
user@host# set firewall family inet filter f2 term t1 from source-address 10.30.1.4/32
user@host# set firewall family inet filter f2 term t1 from destination-address 10.40.1.4/32
user@host# set firewall family inet filter f2 term t1 then inband-flow-telemetry-terminate p2
```

5. Map the firewall filter to the family under the logical unit of the already-configured interface to apply the inband-flow-telemetry-terminate action in the egress direction.

To map the firewall filter:

```
user@host# set interfaces interface-name unit 0 family inet filter input filter-name
```

In this example, you map the f2 firewall filter to the inet family of the logical interface 0 of the physical interface et-0/0/0:

```
user@host# set interfaces et-0/0/0 unit 0 family inet filter input f2
```

View Inband Flow Analyzer Statistics

You can view the following IFA related information:

- IFA statistics using the `show services inband-flow-telemetry stats operational mode` command.
- IFA global parameters using the `show services inband-flow-telemetry global operational mode` command.
- IFA-configured profiles using the `show services inband-flow-telemetry profile operational mode` command.

You can clear the IFA statistics using `clear inband-flow-telemetry stats operational mode` command.

IFA statistics are retrieved directly from the PFE and are not maintained in the Routing Engine. Therefore, a PFE-process restart clears the IFA statistics and a Routing-Engine process restart does not impact the IFA statistics.

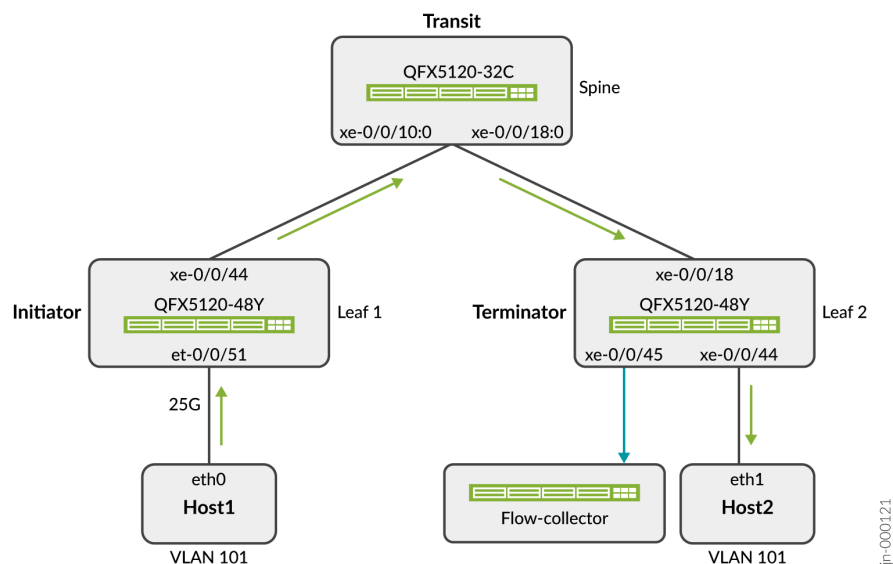
Example - Configure Inband Flow Analyzer 2.0 for Traffic Monitoring

IN THIS SECTION

- Requirements | 395
- Before you Begin | 395
- Overview | 395
- Configuration | 395
- CLI Quick Configuration | 396
- Step-by-Step Procedure | 397
- Results | 400
- Verification | 404

Use this example to configure the IFA 2.0 nodes on your QFX Series switches that enable analyzing of Layer 3 or VXLAN traffic flows. [Figure 51 on page 394](#) shows the topology where IFA 2.0 is configured on QFX Series switches that support the IFA 2.0 feature. In this topology, VXLAN traffic is monitored at the initiator and data is collected at the terminating node for analysis.

Figure 51: Topology for Analyzing VXLAN Traffic Flow using IFA 2.0



jn-000121

Requirements

This example uses the following hardware and software components:

- One QFX5120-32C switch as a spine node
- Two QFX5120-48Y switches as the leaf nodes
- Junos OS Release 21.4R1

This example assumes that you already have an EVPN-VXLAN based network and want to enable traffic monitoring on QFX switches.

Before you Begin

- Make sure you understand how EVPN and VXLAN works. See [Example: Configuring IRB Interfaces in an EVPN-VXLAN Environment to Provide Layer 3 Connectivity for Hosts in a Data Center](#) and [Bridged Overlay Design and Implementation](#) to understand EVPN-VXLAN in detail.
- For IFA terminating node configurations to take effect you need to have a valid IFA license in place.

Overview

In this example, you'll configure one of the QFX5120-48Y switches (Leaf 1) as an initiator node, the QFX5120-32C switch as a transit node, and the second QFX5120-48Y switch (Leaf 2) as a terminating node. The VXLAN traffic flows from Host 1 to Host 2. Configuring IFA on the ingress and egress nodes allows you to monitor network operation and identify the performance issues.

The QFX5120-32C functions as a spine to connect the QFX5120-48Y leaf nodes. At the terminating node, you collect the sampled traffic in IPFIX format using an IPv4 collector application.

Configuration

In this example, you'll configure the following functionality on the switches:

1. Configure Leaf 1 as an initiator node and configure initiator related attributes, like global device identifier and the sampling rate. Configure an IFA profile and firewall filter with the action as `inband-flow-telemetry-init`, and bind the IFA firewall filter to the interfaces.
2. Configure the QFX5120-32C spine switch as a transit node with a global device identifier. When you configure a global device identifier, the spine device adds the IFA metadata and forwards the IFA probe packets.
3. Configure Leaf 2 as a terminating node. Configure the IFA profile with the collector information and firewall filter with the action as `inband-flow-telemetry-terminate`, and bind the IFA firewall filter to the interfaces.

CLI Quick Configuration

To quickly configure this example on your QFX series devices, 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.

Configuration on QFX5120-48Y Switch (Leaf 1 — IFA Initiator Node)

NOTE: Recall that in this example you add IFA to a pre-configured EVPN-VXLAN baseline. The configuration shown here focuses on the delta needed to add IFA to the baseline. We show some of the existing configuration to best show how the IFA delta relates to the baseline.

```
set services inband-flow-telemetry device-id 15000
set services inband-flow-telemetry meta-data-stack-length 100
set services inband-flow-telemetry hop-limit 4
set services inband-flow-telemetry flow-type vxlan
set services inband-flow-telemetry profile ifa_profile_host1 sample-rate 1

set interfaces et-0/0/51:0 unit 0 family ethernet-switching filter input f_init

set firewall family ethernet-switching filter f_init term t1 from ip-protocol udp
set firewall family ethernet-switching filter f_init term t1 from ip-protocol tcp
set firewall family ethernet-switching filter f_init term t1 then inband-flow-telemetry-init
ifa_profile_host1
set firewall family ethernet-switching filter f_init term t1 then count ifa_stats
set firewall family ethernet-switching filter f_init term t1 then accept
set firewall family ethernet-switching filter f_init term t2 then count non_ifa_stats
set firewall family ethernet-switching filter f_init term t2 then accept
```

Configuration on QFX5120-32C Switch (IFA Transit Node)

```
set services inband-flow-telemetry device-id 15001
```

Configuration on QFX5120-48Y Switch (Leaf 2 — IFA Terminating Node)

```
set services inband-flow-telemetry device-id 15002
set services inband-flow-telemetry meta-data-stack-length 100
set services inband-flow-telemetry hop-limit 5
set services inband-flow-telemetry flow-type vxlan
```

```

set services inband-flow-telemetry profile p_term collector source-address 172.16.3.1
set services inband-flow-telemetry profile p_term collector destination-address 172.16.3.2
set services inband-flow-telemetry profile p_term collector destination-port 3055

set interfaces xe-0/0/18 unit 0 family inet filter input f_term

set interfaces xe-0/0/45 description To_Collector
set interfaces xe-0/0/45 unit 0 family inet address 172.16.3.1/24

set firewall family inet filter f_term term ifa then inband-flow-telemetry-terminate p_term
set firewall family inet filter f_term term ifa then count ifa_term
set firewall family inet filter f_term term other then count non_ifa_term
set firewall family inet filter f_term term other then accept

```

Step-by-Step Procedure

Configure QFX5120-48Y Switch (Leaf 1) as an Initiator Node

An IFA initiator node performs the following functions for a flow:

- Samples the flow traffic of interest based on the configuration.
- Converts the traffic into an IFA flow by adding an IFA header to each sample.
- Updates the packet with initiator node metadata.

1. Configure the IFA initiator node attributes. The traffic flow type is configured as VXLAN for initiator node. Note that you must configure the same flow type for both the initiator and the terminating node, either L3 or VXLAN. As in this example, if the VXLAN traffic flow type is configured for the initiator node, ensure that you configure VXLAN traffic flow type for the terminating node as well.

```

[edit]
user@host# set services inband-flow-telemetry device-id 15000
user@host# set services inband-flow-telemetry meta-data-stack-length 100
user@host# set services inband-flow-telemetry hop-limit 4
user@host# set services inband-flow-telemetry flow-type vxlan
user@host# set services inband-flow-telemetry profile ifa_profile_host1 sample-rate 1

```

When sample-rate is configured with value as 1, every packet that is received in the ingress port is sampled. If you prefer less aggressive sampling, increase the sample-rate value.

2. Bind the filter to the initiator node ingress interface.

```
[edit]
user@host# set interfaces et-0/0/51:0 unit 0 family ethernet-switching filter input f_init
```

3. Create a firewall to control IFA sampling. You begin by defining the types of host traffic that should be sampled. In this example you want to perform analysis on UDP and TCP traffic flows. In this example, you configure an firewall filter named `f_init`, with the term name `term1`.

```
[edit]
user@host# set firewall family ethernet-switching filter f_init term t1 from ip-protocol udp
user@host# set firewall family ethernet-switching filter f_init term t1 from ip-protocol tcp
user@host# set firewall family ethernet-switching filter f_init term t1 then accept
```

You configure the filter to perform IFA sampling by adding the action modifier `inband-flow-telemetry-init` to the `t1` term. Note that the inband flow telemetry profile `ifa_profile_host1` is linked to the filter:

```
user@host# set firewall family ethernet-switching filter f_init term t1 then inband-flow-
telemetry-init ifa_profile_host1
user@host# set firewall family ethernet-switching filter f_init term t1 then count ifa_stats
user@host# set firewall family ethernet-switching filter f_init term t2 then count
non_ifa_stats
user@host# set firewall family ethernet-switching filter f_init term t2 then accept
```

Configure QFX5120-32C Switch as a Transit Node

An IFA transit node inserts transit node metadata in the IFA packets in the specified VXLAN flow.

Configure the global device identifier for the transit node, QFX5120-32C switch.

```
user@host# set services inband-flow-telemetry device-id 15001
```

Configure QFX5120-48Y Switch (Leaf 2) as a Terminating Node

An IFA terminating node performs the following for a flow:

- Inserts terminating node metadata in IFA packets.
- Performs a local analytics function on one or more segments of metadata, for example, threshold breach for residence time, congestion notifications, and so on.

- Filters an IFA flow in case of cloned traffic.
 - Sends a copy or report of the packet to collector.
 - Removes the IFA headers and forwards the packet in case of live traffic.
1. Configure the terminating node related attributes, like global device identifier and flow type.

```
user@host# set services inband-flow-telemetry device-id 15002
user@host# set services inband-flow-telemetry meta-data-stack-length 100
user@host# set services inband-flow-telemetry hop-limit 5
user@host# set services inband-flow-telemetry flow-type vxlan
```

Configure an IFA profile with the collector related information.

```
user@host# set services inband-flow-telemetry profile p_term collector source-address
172.16.3.1
user@host# set services inband-flow-telemetry profile p_term collector destination-address
172.16.3.2
user@host# set services inband-flow-telemetry profile p_term collector destination-port 3055
```

2. Configure the collector interface for terminating node Leaf 2.

```
user@host# set interfaces xe-0/0/45 unit 0 family inet address 172.16.3.1/24
```

Apply the firewall filter to the pre-configured interface to activate inband flow telemetry egress processing at Leaf 2.

In this example, you map the f-term firewall filter to the inet family of logical interface 0 of the physical interface xe-0/0/18:

```
user@host# set interfaces xe-0/0/18 unit 0 family inet filter input f_term
```

3. Create a firewall filter and configure the action inband-flow-telemetry-terminate.

In this example, you configure a firewall filter named `f-term`, with the term name `t1` containing the action `inband-flow-telemetry-terminate`, with the inband flow telemetry terminate profile `p_term` mapped to it:

```
user@host# set firewall family inet filter f_term term t1 then count ifa_term
user@host# set firewall family inet filter f_term term t1 then inband-flow-telemetry-
terminate p_term
user@host# set firewall family inet filter f_term term t1 then accept
user@host# set firewall family inet filter f_term term other then count non_ifa_term
user@host# set firewall family inet filter f_term term other then accept
```

Results

Results on QFX5120-48Y Switch (Leaf 1 — IFA Initiator Node)

From operational mode, confirm your configuration by entering the `show configuration services`, `show configuration interfaces`, and `show configuration firewall` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

NOTE: The output shows portions of the pre-existing EVPN-VXLAN baseline to provide the context for the configuration delta needed to add IFA.

```
[edit]
user@host> show configuration services
inband-flow-telemetry {
    device-id {
        15000;
    }
    meta-data-stack-length 100;
    hop-limit 4;
    flow-type vxlan;
    profile {
        ifa_profile_host1 {
            sample-rate 1;
        }
    }
}
```

```

    }
}

```

```

[edit]
user@host> show configuration interfaces
[output truncated]
xe-0/0/44 {
    description Connected_to_Spine1;
    unit 0 {
        family inet {
            address 10.100.13.1/24;
        }
    }
}
et-0/0/51:0 {
    description Connected_to_Host1_vlan_101;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members 101;
            }
            filter {
                input f_init;
            }
        }
    }
}
[output truncated]

```

```

[edit]
user@host> show configuration firewall
family ethernet-switching {
    filter f_init {
        term t1 {
            from {
                ip-protocol [ udp tcp ];
            }
            then {
                accept;
            }
        }
    }
}

```

```

        inband-flow-telemetry-init ifa_profile_host1;
        count ifa_stats;
    }
}
term t2 {
    then {
        accept;
        count non_ifa_stats;
    }
}
}
}
}

```

When you are done configuring the feature on your device, enter `commit` from configuration mode.

Results on QFX5120-32C Switch (IFA Transit Node)

From operational mode, confirm your configuration by entering the `show configuration services`, and `show configuration interfaces` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@host> show configuration services
inband-flow-telemetry {
    device-id {
        15001;
    }
}
}

```

When you are done configuring the feature on your device, enter `commit` from configuration mode.

Results on QFX5120-48Y Switch (Leaf 1 — IFA Terminating Node)

From operational mode, confirm your configuration by entering the `show configuration services`, `show configuration interfaces`, and `show configuration firewall` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@host> show configuration services
inband-flow-telemetry {
    device-id {
        15002;
    }
}

```

```

meta-data-stack-length 100;
hop-limit 5;
flow-type vxlan;
profile {
    p_term {
        collector {
            source-address 172.16.3.1;
            destination-address 172.16.3.2;
            destination-port 3055;
        }
    }
}
}

```

[edit]

user@host> **show configuration interfaces**

```

[edit]
user@host> show configuration interfaces
[output truncated]
xe-0/0/18 {
    description Connected_to_Spine1;
    unit 0 {
        family inet {
            filter {
                input f_term;
            }
            address 10.100.12.1/24;
        }
    }
}
xe-0/0/44 {
    description Connected_to_Host2_vlan_101;
    unit 0 {
        family ethernet-switching {
            interface-mode trunk;
            vlan {
                members 101;
            }
        }
    }
}

```

```

}
xe-0/0/45 {
    description To_Collector;
    mtu 9200;
    unit 0 {
        family inet {
            address 172.16.3.1/24;
        }
    }
}
[output truncated]

```

```

[edit]
user@host> show configuration firewall
family inet {
    filter f_term {
        term t1 {
            then {
                count ifa_term_c;
                inband-flow-telemetry-terminate p_term;
                accept;
            }
        }
        term other {
            then {
                count non_ifa_term;
                accept;
            }
        }
    }
}
}

```

When you are done configuring the feature on your device, enter `commit` from configuration mode.

Verification

Verification on QFX5120-48Y Switch (Leaf 1 – IFA Initiator Node)

Verify IFA Statistics

Purpose

Display the IFA statistics on the initiator node.

Action

From operational mode, enter the `show services inband-flow-telemetry stats` command.

```
IFA Init Packets      : 70989449712
IFA Transit Packets   : 0
IFA Terminate Rx Packets : 0
IFA Terminate Tx Packets : 0
```

Verify IFA Global Configuration**Purpose**

Display the IFA global parameters configured on the initiator node.

Action

From operational mode, enter the `show services inband-flow-telemetry global` command.

```
Global Device ID      : 15000
Meta-data Stack Length : 100
Hop Limit             : 4
Flow Type             : vxlan
```

Verify IFA Profile**Purpose**

Display the IFA profile configured on the initiator node.

Action

From operational mode, enter the `show services inband-flow-telemetry profile` command.

```
Profile Name          : ifa_profile_host1
Sample rate           : 1
Source Address         : 0.0.0.0
Destination Address    : 0.0.0.0
Destination Port       : 0
```

Verification on QFX5120-32C Switch (IFA Transit Node)**Verify IFA Statistics****Purpose**

Display the IFA statistics on the transit node.

Action

From operational mode, enter the `show services inband-flow-telemetry stats` command.

```
IFA Init Packets           : 0
IFA Transit Packets        : 26057387140
IFA Terminate Rx Packets   : 0
IFA Terminate Tx Packets   : 0
```

Verify IFA Global Configuration

Purpose

Display the IFA global parameters configured on the transit node.

Action

From operational mode, enter the `show services inband-flow-telemetry global` command.

```
Global Device ID          : 15001
Meta-data Stack Length    : 240
Hop Limit                  : 250
Flow Type                  : NA
```

Verification on QFX5120-48Y Switch (Leaf 2 — IFA Terminating Node)

Verify IFA Statistics

Purpose

Display the IFA statistics on the terminating node.

Action

From operational mode, enter the `show services inband-flow-telemetry stats` command.

```
IFA Init Packets           : 0
IFA Transit Packets        : 373569
IFA Terminate Rx Packets   : 374448690
IFA Terminate Tx Packets   : 41605188
```

Verify IFA Global Configuration

Purpose

Display the IFA global parameters configured on the terminating node.

Action

From operational mode, enter the `show services inband-flow-telemetry global` command.

```
Global Device ID      : 15002
Meta-data Stack Length : 100
Hop Limit             : 5
Flow Type              : vxlan
```

Verify IFA Profile

Purpose

Display the IFA profile configured on the terminating node.

Action

From operational mode, enter the `show services inband-flow-telemetry profile` command.

```
Profile Name          : p_term
Sample rate           : 0
Source Address         : 172.16.3.1
Destination Address    : 172.16.3.2
Destination Port       : 3055
```

SEE ALSO

- [inband-flow-telemetry | 1157](#)
- [clear inband-flow-telemetry stats | 1560](#)
- [show services inband-flow-telemetry | 1701](#)

Release History Table

Release	Description
22.4R1-EVO	Inband Flow Analyzer (IFA) 2.0 transit node support (QFX Series switches)—In Junos OS Evolved 22.4R1, we've extended support for the IFA 2.0 transit node role to the QFX5130-32CD, QFX5220-32CD, QFX5220-128C, and QFX5700 switches.

22.2R1	Inband Flow Analyzer (IFA) 2.0 (QFX Series switches)—In Junos OS Release 22.2R1, we've extended support for IFA 2.0 to the QFX5120-48YM and QFX5120-48T switches. We've also added support for configuring the MTU and maximum clip length for IFA packets, and for the QFX5120-48YM switch, setting the IFA clock source.
21.4R1	Inband Flow Analyzer (IFA) 2.0 (QFX5120-48Y and QFX5120-32C)—In Junos OS Release 21.4R1, we've introduced support for IFA 2.0 on QFX Series switches. IFA 2.0 monitors and analyzes packets when they enter and exit the network. You can use IFA 2.0 to monitor the network for faults and performance bottlenecks. IFA 2.0 supports both Layer 3 and VXLAN flows.

Juniper Resiliency Interface

IN THIS CHAPTER

- [Juniper Resiliency Interface | 409](#)

Juniper Resiliency Interface

SUMMARY

For MX Series routers with MPC line cards and PTX Series routers with the JNP10K-LC1201 or JNP10K-LC1203 linecards running Junos OS Evolved, you can configure the Juniper Resiliency Interface (JRI) to detect, correlate, and mitigate exceptions and thereby reduce the mean time to repair (MTTR) for issues. For forwarding exceptions, JRI also extends the inline monitoring services feature with Juniper-specific IPFIX information elements (IEs) for exception data and introduces the concept of an observation cloud, which is a set of observation domains. You can send the IPFIX packets to either an on-box or an off-box collector.

IN THIS SECTION

- [Understand Juniper Resiliency Interface | 409](#)
- [Configure JRI for Operating System and Routing Exceptions | 412](#)
- [Configure JRI for Forwarding Exceptions | 413](#)

Understand Juniper Resiliency Interface

Packets that need to be forwarded to the adjacent network element or a neighboring device along a routing path might be dropped by a router owing to several factors. Every network encounters issues, such as packet loss, from time to time. Some of the causes for such a loss of traffic or a block in transmission of data packets include: overloaded system conditions, profiles and policies that restrict the bandwidth or priority of traffic, network outages, or disruption by physical cable faults. Packet loss also happens because of incorrect stitching of the forwarding path or a mismatch between the control plane state and the data plane state. You could use counters and metrics from `show` commands to diagnose and debug network performance, but doing so can be tedious and time-consuming. JRI reports exception

data from entities in the system which encounter packet drops, enabling you to automate the workflow involved in detecting, reporting and mitigating adverse exceptions.

For operating system and routing exceptions, the exception data is reported in telemetry key-value pairs.

For forwarding exceptions, the exception data is reported in IPFIX packets. The IEs in the IPFIX primary data record packet capture the following data:

- Exception reason (for example, firewall discard)
- Packet direction (ingress or egress)
- First N bytes of the packet
- Ingress interface
- Egress interface
- Next-hop identifier (Junos OS only)

[Table 62 on page 410](#) shows the format of the IPFIX Primary Data Record with the Juniper-specific IEs.

Table 62: IPFIX Primary Data Record

IE Name	IE Identifier	Description	Length (in Bytes)
forwardingClassandDropPriority	Observation Cloud Common Property ID (CPID)—IE 137, a set of common properties that is locally unique per Observation Cloud	Forwarding class and drop priority ID	4
forwardingExceptionCode	Observation Cloud CPID—IE 137	Exception code that causes packet drops OR is zero when the exception is not met or set	2
forwardingNextHopId	Observation Cloud CPID—IE 137	(Junos OS only) Unicast next-hop Index used for forwarding	4

Table 62: IPFIX Primary Data Record *(Continued)*

IE Name	IE Identifier	Description	Length (in Bytes)
egressInterfaceIndex	Observation Cloud CPID—IE 137	Index of egress logical interface when flowDirection=output, otherwise 0.	4
underlyingIngressInterfaceIndex	Observation Cloud CPID—IE 137	(Junos OS only) Index of underlying layer 2 ingress logical interface, wherever applicable (for example, AE and IRB cases—see "primary-data-record-fields" on page 1310 for more information)	4
ingressInterfaceIndex	Observation Cloud CPID—IE 137	Index of ingress logical interface	4
ingressInterface	IE 10	SNMP index of ingress logical interface	4
egressInterface	IE 14	SNMP index of egress logical interface when flowDirection=output, otherwise 0.	4
flowDirection	IE 61	Direction (0: input, 1:output)	1
dataLinkFrameSize	IE 312	Length of sampled data link frame	2
dataLinkFrameSection	IE 315	N octets from the data link frame of the monitored packet	variable

Limitations:

- Exceptions are collected and exported on a best-effort basis.

- Any limitations or caveats for inline monitoring services also apply to JRI, because JRI uses inline monitoring services to sample and collect the packets.
- All dropped packets cannot be sampled and profiled. Classes of exceptions are sampled at the default sampling rate, unless you configure this rate with the `sampling-rate` statement at either the `[edit services inline-monitoring instance instance-name collector collector-name]` hierarchy level (Junos OS) or at the `[edit services inline-monitoring instance instance-name]` hierarchy level (Junos OS Evolved). Junos OS allows the sampling rate to be configured per collector, allowing different rates for each collector; Junos OS Evolved allows one sampling rate per inline-monitoring instance.
- For exception reporting in the egress direction, the layer 2 header or any encapsulation header is not included in IE-315, `dataLinkFrameSelection`, because exceptions happen before layer 2 or tunnel encapsulation.
- For exception reporting in the egress direction, the receiver of the IPFIX packet must ignore IE-312, `dataLinkFrameSize`, because the field does not have the correct value.
- For the egress direction, you cannot configure both sFlow and exception reporting on the same interface.
- Inline-monitoring instance actions and firewall re-direct instance actions are not supported in the same term of the firewall filter. (Junos OS Evolved)
- Inline-monitoring instance actions and port-mirroring instance actions are not supported in the same term of the firewall filter. (Junos OS Evolved)
- For collectors, you cannot configure routing instances, DSCP bits, or forwarding class. (Junos OS Evolved)
- For more information about the Juniper-specific IEs, including caveats and limitations, see ["primary-data-record-fields" on page 1310](#).

Configure JRI for Operating System and Routing Exceptions

To configure JRI for operating system and routing exceptions:

1. Subscribe to the Junos Telemetry Interface XPaths:

Notifications are exported using gRPC/gNMI to an off-box collector.

For Junos OS:

```
/junos/exception-profiles/routing-profile
/junos/exception-profiles/os-profile/
```

For Junos OS Evolved (routing exceptions only):

```
/junos/exception-profiles/routing-profile
```

2. (Optional) Additionally, if you prefer to use the on-box collector instead of sending the data to an off-box collector, then configure an on-box storage location for the exception data.

To configure:

```
user@host# set system resiliency exceptions exception-type
user@host# set system resiliency store file file-name
user@host# set system resiliency store file size file-size
```

In this example, you configure the file in which to store the exception data:

For Junos OS:

```
user@host# set system resiliency exceptions routing
user@host# set system resiliency exceptions os
user@host# set system resiliency store file file1
user@host# set system resiliency store size 1g
```

For Junos OS Evolved:

```
user@host# set system resiliency exceptions routing
user@host# set system resiliency store file file1
user@host# set system resiliency store size 1g
```

Configure JRI for Forwarding Exceptions

To configure JRI for forwarding exceptions:

1. Define the IPFIX template.

To configure attributes of the template:

For Junos OS:

```
user@host# set services inline-monitoring template template_1 template-refresh-rate template-  
refresh-rate
user@host# set services inline-monitoring template template_1 template-id template-identifier
```



```
user@host# set services inline-monitoring template template_1 primary-data-record-fields
primary-data-record-field-name
```

In this example, the template refresh rate is set to 30 seconds, you've configured a template identifier, and you've configured the fields of the primary data record:

```
user@host# set services inline-monitoring template template_1 template-refresh-rate 30
user@host# set services inline-monitoring template template_1 template-id 1024
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
ingress-interface-index
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
underlying-interface-index
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
egress-interface-index
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
forwarding-nexthop-id
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
forwarding-exception-code
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
forwarding-class-drop-priority
user@host# set services inline-monitoring template template_1 primary-data-record-fields
ingress-interface-snmp-id
user@host# set services inline-monitoring template template_1 primary-data-record-fields
egress-interface-snmp-id
user@host# set services inline-monitoring template template_1 primary-data-record-fields
direction
```

For Junos OS Evolved, the system generates the template ID and the software supports most of the fields of the primary data record:

```
user@host# set services inline-monitoring template template_1 template-refresh-rate template-
refresh-rate
user@host# set services inline-monitoring template template_1 primary-data-record-fields
primary-data-record-field-name
```

In this example, the template refresh rate is set to 30 seconds and you've configured the fields of the primary data record:

```
user@host# set services inline-monitoring template template_1 template-refresh-rate 30
user@host# set services inline-monitoring template template_1 template-id 1024
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
```

```

ingress-interface-index
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
egress-interface-index
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
forwarding-exception-code
user@host# set services inline-monitoring template template_1 primary-data-record-fields cpid-
forwarding-class-drop-priority
user@host# set services inline-monitoring template template_1 primary-data-record-fields
ingress-interface-snmp-id
user@host# set services inline-monitoring template template_1 primary-data-record-fields
egress-interface-snmp-id
user@host# set services inline-monitoring template template_1 primary-data-record-fields
direction

```

2. Attach the template to the instance and describe the collector.

Junos OS and Junos OS Evolved differ in how to achieve this step. To configure the instance and collector:

For Junos OS:

```

user@host# set services inline-monitoring instance instance-name template-name template-name
user@host# set services inline-monitoring instance instance-name collector collector-name
source-address IPv4-address
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-address IPv4-address
user@host# set services inline-monitoring instance instance-name collector collector-name
dscp dscp-bits
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-port port

```

In this example, you create a template with the name `template_1`, create an inline-monitoring instance `i1`, and create the configuration for the on-box collector `c2`. For an on-box collector for Junos OS, the destination address must be a local address and the destination port must be port 4739. For an off-box collector for Junos OS, specify the destination address and port for that collector.

For Junos OS:

```

user@host# set services inline-monitoring instance i1 template-name template_1
user@host# set services inline-monitoring instance i1 collector c2 source-address 10.11.12.1
user@host# set services inline-monitoring instance i1 collector c2 destination-address
10.11.12.2

```

```

user@host# set services inline-monitoring instance i1 collector c2 dscp 21
user@host# set services inline-monitoring instance i1 collector c2 destination-port 4739

```

For Junos OS Evolved, you cannot configure the DSCP bits, but the process is otherwise the same as in Junos OS for an off-box collector:

```

user@host# set services inline-monitoring instance instance-name template-name template-name
user@host# set services inline-monitoring instance instance-name collector collector-name
source-address IPv4-address
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-address IPv4-address
user@host# set services inline-monitoring instance instance-name collector collector-name
destination-port port

```

For Junos OS Evolved, for an on-box collector, you configure the `controller re` statement instead of a local destination address and port, and you cannot configure the DSCP bits:

```

user@host# set services inline-monitoring instance instance-name template-name template-name
user@host# set services inline-monitoring instance instance-name collector collector-name
source-address IPv4-address
user@host# set services inline-monitoring instance instance-name controller re

```

In this example, for Junos OS Evolved, you create a template with the name `template_1`, create an inline-monitoring instance `i1`, and create the configuration for the on-box collector `c2`. For an on-box collector, you specify the `controller re` statement instead of a local destination address and port:

```

user@host# set services inline-monitoring instance i1 template-name template_1
user@host# set services inline-monitoring instance i1 collector c2 source-address 10.11.12.1
user@host# set services inline-monitoring instance i1 controller re

```

3. Configure the observation cloud identifier.

An observation cloud is the largest set of observation domains. According to RFC 5101, an observation domain is the largest set of observation points for which flow information can be aggregated by a metering process. For example, a router line card may be an observation domain if it is composed of several interfaces, each of which is an observation point. By configuring an observation cloud, you allow inline-monitoring services to report on a set of common properties that

is locally unique per observation cloud. For more information about observation clouds, see ["inline-monitoring" on page 1163](#). To configure the observation cloud identifier:

```
user@host# set services inline-monitoring observation-cloud-id identifier
```

In this example, you have configured the identifier as 1:

```
user@host# set services inline-monitoring observation-cloud-id 1
```

4. Subscribe to various exception types and configure exception reporting for a particular PFE and specify the inline-monitoring instance. For Junos OS, you must specify a particular exception category name, such as forwarding-state. For Junos OS Evolved, you simply specify all as the category name.

By default, the exception data is sent to an off-box collector. To configure:

```
user@host# set chassis fpc slot-number pfe identifier exception-reporting category category-name inline-monitoring-instance inline-monitoring-instance-name
```

For Junos OS:

In this example, you subscribe to forwarding exceptions and configure FPC 0 to send forwarding exceptions to the inline-monitoring instance i1:

```
user@host# set chassis fpc 0 pfe 0 exception-reporting category forwarding-state inline-monitoring-instance i1
```

For Junos OS Evolved:

In this example, you subscribe to all exception categories and configure FPC 0 to send exceptions to the inline-monitoring instance i1:

```
user@host# set chassis fpc 0 pfe 0 exception-reporting category all inline-monitoring-instance i1
```

5. (Optional) Additionally, if you prefer to use the on-box collector instead of sending the data to an off-box collector, then configure an on-box storage location for the exception data.

To configure:

```
user@host# set system resiliency exceptions forwarding
user@host# set system resiliency store fwding-file file-name
user@host# set system resiliency store fwding-file size file-size
```

In this example, you configure the file in which to store the forwarding exception data:

```
user@host# set system resiliency exceptions forwarding
user@host# set system resiliency store fwding-file file1
user@host# set system resiliency store fwding-file size 1g
```

Release History Table

Release	Description
22.2R1-EVO	Support for the Juniper Resiliency Interface (PTX10001-36MR, PTX10004, PTX10008, and PTX10016 routers with the JNP10K-LC1201 or JNP10K-LC1203 linecards)—Starting in Junos OS Evolved Release 22.2R1, you can use the Juniper Resiliency Interface (JRI) to detect, correlate, and mitigate exceptions.
21.2R1	Support for the Juniper Resiliency Interface (MX480, MX960, MX2010, MX2020 and vMX)—Starting in Junos OS Release 21.2R1, you can use our new Juniper Resiliency Interface (JRI) to detect, correlate, and mitigate exceptions. JRI extends the inline monitoring services feature with Juniper-specific IPFIX information elements (IEs) for exception data and introduces the concept of an Observation Cloud, which is a set of Observation Domains. You can send the IPFIX packets to either an on-box or an off-box collector.

RELATED DOCUMENTATION

| [Inline Monitoring Services Configuration](#)

4

PART

Sampling and Discard Accounting Services

[Sampling Data Using Traffic Sampling and Discard Accounting | 420](#)

[Sampling Data Using Inline Sampling | 437](#)

[Sampling Data Using Flow Aggregation | 576](#)

Sampling Data Using Traffic Sampling and Discard Accounting

IN THIS CHAPTER

- [Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)
- [Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches | 433](#)
- [Configuring Discard Accounting | 435](#)

Configuring Traffic Sampling on MX, M and T Series Routers

IN THIS SECTION

- [Configuring Firewall Filter for Traffic Sampling | 421](#)
- [Configuring Traffic Sampling on a Logical Interface | 422](#)
- [Disabling Traffic Sampling | 424](#)
- [Sampling Once | 424](#)
- [Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets | 425](#)
- [Configuring Traffic Sampling Output | 426](#)
- [Tracing Traffic Sampling Operations | 428](#)
- [Traffic Sampling Examples | 429](#)

Traffic sampling enables you to copy traffic to a Physical Interface Card (PIC) that performs flow accounting while the router forwards the packet to its original destination. You can configure the router to perform sampling in one of the following three locations:

- On the Routing Engine, using the sampled process. To select this method, use a filter (input or output) with a matching term that contains the `then sample` statement.

- On the Monitoring Services, Adaptive Services, or Multiservices PIC.
- On an inline data path without the need for a services Dense Port Concentrator (DPC). To do this inline active sampling, you define a sampling instance with specific properties. One Flexible PIC Concentrator (FPC) can support only one instance; for each instance, either services PIC-based sampling or inline sampling is supported per family. Inline sampling supports version 9 and IPFIX flow collection templates.

NOTE: Routing Engine based sampling is not supported on VPN routing and forwarding (VRF) instances.

Configuring Firewall Filter for Traffic Sampling

To configure firewall filter for traffic sampling, you must perform the following tasks:

- Create a firewall filter to apply to the logical interfaces being sampled by including the filter statement at the [edit firewall family *family-name*] hierarchy level. In the filter then statement, you must specify the action modifier `sample` and the action `accept`.

```
filter filter-name {
  term term-name {
    then {
      sample;
      accept;
    }
  }
}
```

For more information about firewall filter actions and action modifiers, see the [Routing Policies, Firewall Filters, and Traffic Policers User Guide](#).

- Apply the filter to the interfaces on which you want to sample traffic by including the address and filter statements at the [edit interfaces *interface-name* unit *logical-unit-number* family *family-name*] hierarchy level:

```
address address {
}
filter {
```



```
input filter-name;
}
```

The following prerequisites apply to M, MX, and T Series routers when you configure traffic sampling on interfaces and in firewall filters:

- If you configure a sample action in a firewall filter for an inet or inet6 family on an interface without configuring the forwarding-options settings, operational problems might occur if you also configure port mirroring or flow-tap functionalities. In such a scenario, all the packets that match the firewall filter are incorrectly sent to the service PIC.
- If you include the then sample statement at the [edit firewall family inet filter *filter-name* term *term-name*] hierarchy level to specify a sample action in a firewall filter for IPv4 packets, you must also include the family inet statement at the [edit forwarding-options sampling] hierarchy level or the instance *instance-name* family inet statement at the [edit forwarding-options sampling] hierarchy level. Similarly, if you include the then sample statement at the [edit firewall family inet6 filter *filter-name* term *term-name*] hierarchy level to specify a sample action in a firewall filter for IPv6 packets, you must also include family inet6 statement at the [edit forwarding-options sampling] hierarchy level or the instance *instance-name* family inet6 statement at the [edit forwarding-options sampling] hierarchy level. Otherwise, a commit error occurs when you attempt to commit the configuration.
- Also, if you configure traffic sampling on a logical interface by including the sampling input or sampling output statements at the [edit interface *interface-name* unit *logical-unit-number*] hierarchy level, you must also include the family inet | inet6 statement at the [edit forwarding-options sampling] hierarchy level, or the instance *instance-name* family inet | inet6 statement at the [edit forwarding-options sampling] hierarchy level.

Configuring Traffic Sampling on a Logical Interface

To configure traffic sampling on any logical interface, enable sampling and specify a non zero sampling rate by including the sampling statement at the [edit forwarding-options] hierarchy level:

```
sampling {
  input {
    rate number;
    run-length number;
    max-packets-per-second number;
    maximum-packet-length bytes;
  }
}
```


When you use Routing Engine-based sampling, specify the threshold traffic value by including the max-packets-per-second statement. The value is the maximum number of packets to be sampled, beyond which

the sampling mechanism begins dropping packets. The range is from 0 through 65,535. A value of 0 instructs the Packet Forwarding Engine not to sample any packets. The default value is 1000.

NOTE: When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the output statement, or when you configure inline sampling, the `max-packets-per-second` value is ignored.

Specify the sampling rate by setting the values for `rate` and `run-length` (see [Figure 52 on page 423](#)).

Figure 52: Configuring Sampling Rate



1746

NOTE: Do not configure ingress sampling on `ms-` logical interfaces on which PIC-based flow monitoring is enabled, which causes undesired flow monitoring behavior and might result in repeated sampling of a single packet. Starting in Junos OS Release 15.1, a commit error occurs when you try to configure ingress traffic sampling on that interface. In Junos OS Release 14.2 and earlier, the commit error does not occur, but you should not configure ingress traffic sampling on that interface.

If PIC-based flow monitoring is enabled on an `ms-fpc/pic/port.logical-unit` interface, a commit check error occurs when you attempt to configure ingress traffic sampling on that interface. This error occurs because a combination of ingress sampling and PIC-based flow monitoring operations on an `ms-` logical interface causes undesired flow monitoring behavior and might result

in repeated sampling of a single packet. You must not configure ingress sampling on `ms-` logical interfaces on which PIC-based flow monitoring is enabled.

The `rate` statement specifies the ratio of packets to be sampled. For example, if you configure a rate of 10, x number of packets out of every 10 is sampled, where $x = \text{run length} + 1$. By default, the rate is 0, which means that no traffic is sampled.

The `run-length` statement specifies the number of matching packets to sample following the initial one-packet trigger event. By default, the run length is 0, which means that no more traffic is sampled after the trigger event. The range is from 0 through 20. Configuring a run length greater than 0 allows you to sample packets following those already being sampled.

NOTE: The `run-length` and `maximum-packet-length` configuration statements are not supported on MX80 routers.

If you do not include the `input` statement, sampling is disabled.

To collect the sampled packets in a file, include the `file` statement at the `[edit forwarding-options sampling output]` hierarchy level. Output file formats are discussed later in the chapter.

Disabling Traffic Sampling

To explicitly disable traffic sampling on the router, include the `disable` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
disable;
```

Sampling Once

To explicitly sample a packet for active monitoring only once, include the `sample-once` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
sample-once;
```

Setting this option avoids duplication of packets in cases where sampling is enabled at both the ingress and egress interfaces and simplifies analysis of the sampled traffic.

Preserving Prerewrite ToS Value for Egress Sampled or Mirrored Packets

Starting in Junos OS Release 14.1, you can preserve the prenormalized type-of-service (ToS) value in egress sampled or mirrored packets. Include the `pre-rewrite-tos` statement at the `[edit forwarding-options sampling]` hierarchy level.

On MPC-based interfaces, you can configure ToS rewrite either using class-of-service (CoS) configuration by including the `rewrite-rules dscp rule_name` statement at the `[edit class-of-service interfaces interface-name unit logical-unit-number]` hierarchy level or using firewall filter configuration by including the `dscp` statement at the `[edit firewall family family-name filter filter-name term term-name then]` hierarchy level. If ToS rewrite is configured, the egress mirrored or sampled copies contain the post-rewrite ToS values by default. With the `pre-rewrite-tos` configuration, you can retain the prerewrite ToS value in the sampled or mirrored packets.

NOTE:

- If ToS rewrite is configured on the egress interface by using both CoS and firewall filter configuration, and if the `pre-rewrite-tos` statement is also configured, then the egress sampled packets contain the DSCP value set using the firewall filter configuration. However, if the `pre-rewrite-tos` statement is not configured, the egress sampled packets contain the DSCP value set by the CoS configuration.
- With the `pre-rewrite-tos` statement, you can configure retaining prenormalization ToS values only for sampling done under `family inet` and `family inet6`.
- This feature cannot be configured at the `[edit logical-systems]` hierarchy level. It can be configured only at the global level under the `forwarding-option` configuration.
- When ToS rewrite is configured by using a firewall filter on both ingress and egress interfaces, the egress sampled packets contain the DSCP value set by the ingress ToS rewrite configuration if the `pre-rewrite-tos` statement is configured. However, if the `pre-rewrite-tos` statement is not configured, the egress sampled packets contain the DSCP value set by the ToS rewrite configuration for the egress firewall filter.
- If the `pre-rewrite-tos` statement is configured, and a deactivate or delete operation is performed at the `[edit forwarding-options]` hierarchy level, `pre-rewrite-tos` configuration still remains active. To disable the `pre-rewrite-tos` configuration for such a case, you must explicitly deactivate or delete the `pre-rewrite-tos` statement at the `[edit forwarding-options sampling]` hierarchy level before performing a deactivate or delete operation at the `[edit forwarding-options]` hierarchy level.

Configuring Traffic Sampling Output

To configure traffic sampling output, include the following statements at the [edit forwarding-options sampling family (inet | inet6 | mpls) output] hierarchy level:

```

aggregate-export-interval seconds;
flow-active-timeout seconds;
flow-inactive-timeout seconds;
extension-service service-name;
  flow-server hostname {
    aggregation {
      autonomous-system;
      destination-prefix;
      protocol-port;
      source-destination-prefix {
        caida-compliant;
      }
      source-prefix;
    }
    autonomous-system-type (origin | peer);
    (local-dump | no-local-dump);
    port port-number;
    source-address address;
    version format;
    version9 {
      template template-name;
    }
  }
interface interface-name {
  engine-id number;
  engine-type number;
  source-address address;
}
file {
  disable;
  filename filename;
  files number;
  size bytes;
  (stamp | no-stamp);
  (world-readable | no-world-readable);
}

```

To configure inline flow monitoring on MX Series routers, include the `inline-jflow` statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output]` hierarchy level. Inline sampling exclusively supports a new format called IP_FIX that uses UDP as the transport protocol. When you configure inline sampling, you must include the `version-ipfix` statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output flow-server address]` hierarchy level and also at the `[edit services flow-monitoring]` hierarchy level. For more information about configuring inline flow monitoring, see ["Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250" on page 74](#).

To direct sampled traffic to a flow-monitoring interface, include the `interface` statement. The `engine-id` and `engine-type` statements specify the identity and type numbers of the interface; they are dynamically generated based on the Flexible PIC Concentrator (FPC), PIC, and slot numbers and the chassis type. The `source-address` statement specifies the traffic source.

Starting in Junos OS Release 19.3R1, to configure inline flow monitoring on Juniper Sky Advanced Threat Prevention (ATP), include the `flow-server` statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output]` hierarchy level. Inline sampling exclusively supports a new format called IP_FIX that uses UDP as the transport protocol. When you configure inline sampling, you must include the `version-ipfix` statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output flow-server address]` hierarchy level and also at the `[edit services flow-monitoring]` hierarchy level.

To configure flow sampling version 9 output, you need to include the `template` statement at the `[edit forwarding-options sampling output version9]` hierarchy level. For information on cflowd, see ["Enabling Flow Aggregation" on page 577](#).

The `aggregate-export-interval` statement is described in ["Configuring Discard Accounting" on page 435](#), and the `flow-active-timeout` and `flow-inactive-timeout` statements are described in ["Configuring Flow Monitoring" on page 5](#).

Traffic sampling results are automatically saved to a file in the `/var/tmp` directory. To collect the sampled packets in a file, include the `file` statement at the `[edit forwarding-options sampling family inet output]` hierarchy level:

```
file {
  disable;
  filename filename;
  files number;
  size bytes;
  (stamp | no-stamp);
  (world-readable | no-world-readable);
}
```

Traffic Sampling Output Format

Traffic sampling output is saved to an ASCII text file. The following is an example of the traffic sampling output that is saved to a file in the **/var/tmp** directory. Each line in the output file contains information for one sampled packet. You can optionally display a timestamp for each line.

The column headers are repeated after each group of 1000 packets.

```
# Apr  7 15:48:50
Time                Dest          Src Dest Src Proto TOS Pkt Intf  IP   TCP
                  addr          addr port port          len num frag flags
Apr 7 15:48:54 192.168.9.194 192.168.9.195  0    0    1   0x0 84  8   0x0  0x0
Apr 7 15:48:55 192.168.9.194 192.168.9.195  0    0    1   0x0 84  8   0x0  0x0
Apr 7 15:48:56 192.168.9.194 192.168.9.195  0    0    1   0x0 84  8   0x0  0x0
Apr 7 15:48:57 192.168.9.194 192.168.9.195  0    0    1   0x0 84  8   0x0  0x0
Apr 7 15:48:58 192.168.9.194 192.168.9.195  0    0    1   0x0 84  8   0x0  0x0
```

To set the timestamp option for the file `my-sample`, enter the following:

```
[edit forwarding-options sampling output file]
user@host# set filename my-sample files 5 size 2m world-readable stamp;
```

Whenever you toggle the timestamp option, a new header is included in the file. If you set the stamp option, the `Time` field is displayed.

```
# Apr  7 15:48:50
# Time            Dest          Src Dest Src Proto TOS  Pkt Intf  IP   TCP
#                addr          addr port port          len  num frag flags
# Feb  1 20:31:21
#                Dest          Src Dest Src Proto TOS  Pkt Intf  IP   TCP
#                addr          addr port port          len  num frag flags
```

Tracing Traffic Sampling Operations

Tracing operations track all traffic sampling operations and record them in a log file in the **/var/log** directory. By default, this file is named **/var/log/sampled**. The default file size is 128K, and 10 files are created before the first one gets overwritten.

To trace traffic sampling operations, include the `traceoptions` statement at the `[edit forwarding-options sampling]` hierarchy level:

```
traceoptions {
  no-remote-trace;
  file filename <files number> <size bytes> <match expression> <world-readable | no-world-readable>;
}
```

Traffic Sampling Examples

Example: Sampling a Single SONET/SDH Interface

The following configuration gathers statistical sampling information from a small percentage of all traffic on a single SONET/SDH interface and collects it in a file named `sonet-samples.txt`.

Create the filter:

```
[edit firewall family inet]
filter {
  input sample-sonet {
    then {
      sample;
      accept;
    }
  }
}
```

Apply the filter to the SONET/SDH interface:

```
[edit interfaces]
so-0/0/1 {
  unit 0 {
    family inet {
      filter {
        input sample-sonet;
      }
      address 10.127.68.254/32 {
        destination 172.16.74.7;
      }
    }
  }
}
```



```

    }
  }
}

```

Finally, configure traffic sampling:

```

[edit forwarding-options]
sampling {
  input {
    family inet {
      rate 100;
      run-length 2;
    }
  }
  family inet {
    output {
      file {
        filename sonet-samples.txt;
        files 40;
        size 5m;
      }
    }
  }
}

```

Example: Sampling All Traffic from a Single IP Address

The following configuration gathers statistical information about every packet entering the router on a specific Gigabit Ethernet port originating from a single source IP address of 172.16.92.31, and collects it in a file named `samples-172-16-92-31.txt`.

Create the filter:

```

[edit firewall family inet]
filter one-ip {
  term get-ip {
    from {
      source-address 172.16.92.31;
    }
    then {
      sample;
    }
  }
}

```

```

        accept;
    }
}

```

Apply the filter to the Gigabit Ethernet interface:

```

[edit interfaces]
ge-4/1/1 {
    unit 0 {
        family inet {
            filter {
                input one-ip;
            }
            address 10.45.92.254;
        }
    }
}

```

Finally, gather statistics on all the candidate samples; in this case, gather all statistics:

```

[edit forwarding-options]
sampling {
    input {
        family inet {
            rate 1;
        }
    }
    family inet {
        output {
            file {
                filename samples-172-16-92-31.txt;
                files 100;
                size 100k;
            }
        }
    }
}

```

Example: Sampling All FTP Traffic

The following configuration gathers statistical information about a moderate percentage of packets using the FTP data transfer protocol in the output path of a specific T3 interface, and collects the information in a file named `t3-ftp-traffic.txt`.

Create a filter:

```
[edit firewall family inet]
filter ftp-stats {
  term ftp-usage {
    from {
      destination-port [ftp ftp-data];
    }
    then {
      sample;
      accept;
    }
  }
}
```

Apply the filter to the T3 interface:

```
[edit interfaces]
t3-7/0/2 {
  unit 0 {
    family inet {
      filter {
        input ftp-stats;
      }
      address 10.35.78.254/32 {
        destination 10.35.78.4;
      }
    }
  }
}
```

Finally, gather statistics on 10 percent of the candidate samples:

```
[edit forwarding-options]
sampling {
```

```
input {
    family inet {
        rate 10;
    }
}
family inet {
    output {
        file {
            filename t3-ftp-traffic.txt;
            files 50;
            size 1m;
        }
    }
}
```

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, you can preserve the prenormalized type-of-service (ToS) value in egress sampled or mirrored packets. Include the pre-rewrite-tos statement at the [edit forwarding-options sampling] hierarchy level.

RELATED DOCUMENTATION

- Traffic Sampling, Forwarding, and Monitoring Overview*
- [Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches | 433](#)

Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches

You can configure active sampling by defining a sampling instance that specifies a name for the sampling parameters and bind the instance name to an FPC, MPC, or DPC. This configuration enables you to define multiple named sampling parameter sets associated with multiple destinations and protocol families per sampling destination. With the cflowd version 5 and version 8 and flow aggregation version 9, you can use templates to organize the data gathered from sampling.

To implement this feature, you include the `instance` statement at the `[edit forwarding-options sampling]` hierarchy level.

The following considerations apply to the sampling instance configuration:

- This configuration is supported on the IP version 4 (`inet`), IP version 6 (`ipv6`), and MPLS protocol families.
- You can configure the router to perform sampling in either of two locations:
 - On the Routing Engine, using the sampled process. To select this method, use a filter (input or output) with a matching term that contains the `then sample` statement.
 - On the Monitoring Services, Adaptive Services, or Multiservices PIC. Specify the interface name at the `[forwarding-options sampling instance instance-name family inet output interface]` hierarchy level. You can configure the same or different services PICs in a set of sampling instances.
- You can configure the rate and run-length options at the `[edit forwarding-options sampling input]` hierarchy level to apply common values for all families on a global basis. Alternatively, you can configure these options at the `[edit forwarding-options sampling instance instance-name input]` hierarchy level to apply specific values for each instance or at the `[edit forwarding-options sampling instance instance-name family family input]` hierarchy level to apply specific values for each protocol family you configure.
- Starting in Junos OS Release 16.1, for inline active flow monitoring, you can configure a Differentiated Services Code Point (DSCP) mapping and a forwarding class to apply to exported packets. Use the `dscp` and `forwarding-class` options at the `[edit forwarding-options sampling instance-name family (inet | inet6) output flow-server hostname]` hierarchy level.
- For MX Series devices with Modular Port Concentrators (MPCs), port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 through 255 bytes. Only the values 1 to 255 are valid for packet truncation on these devices. For other devices, the range is from 0 through 9216. A maximum-packet-length value of zero (0) represents that truncation is disabled, and the entire packet is mirrored or sampled.

NOTE: The `run-length` and `maximum-packet-length` configuration statements are not supported on MX80 routers.

To associate the defined instance with a particular FPC, MPC, or DPC, you include the `sampling-instance` statement at the `[edit chassis fpc number]` hierarchy level, as in the following example:

```
chassis {
  fpc 2 {
```

```

        sampling-instance samp1;
    }
}

```

Starting in Junos OS Release 14.1, you can associate a sampling instance with an FPC in the MX Series Virtual Chassis primary or backup router. Use the **sampling-instance *instance-name*** statement at the **[edit chassis member *member-number* fpc slot *slot-number*]** hierarchy level, where *member-number* is 0 (for the primary router) or 1 (for the backup router), and *slot-number* is a number in the range 0 through 11.

Release History Table

Release	Description
16.1	Starting in Junos OS Release 16.1, for inline active flow monitoring, you can configure a Differentiated Services Code Point (DSCP) mapping and a forwarding class to apply to exported packets.
14.1	Starting in Junos OS Release 14.1, you can associate a sampling instance with an FPC in the MX Series Virtual Chassis primary or backup router.

RELATED DOCUMENTATION

<i>Traffic Sampling, Forwarding, and Monitoring Overview</i>
Monitoring, Sampling, and Collection Services Interfaces User Guide
Configuring Active Flow Monitoring 42
<i>Directing Traffic Sampling Output to a Server Running the cflowd Application</i>
Configuring Traffic Sampling on MX, M and T Series Routers 420
Example: Sampling Instance Configuration 133
sampling (Forwarding Options) 1380
Inline Flow Monitoring for Virtual Chassis Overview

Configuring Discard Accounting

Discard accounting is similar to traffic sampling, but varies from it in two ways:

- In discard accounting, the packet is intercepted by the monitoring PIC and is not forwarded to its destination.

- Traffic sampling allows you to limit the number of packets sampled by configuring the `max-packets-per-second`, `rate`, and `run-length` statements. Discard accounting does not provide these options, and a high packet count can potentially overwhelm the monitoring PIC.

A discard instance is a named entity that specifies collector information under the `accounting name` statement. Discard instances are referenced in firewall filter `term` statements by including the `then discard accounting name` statement.

Most of the other statements are also found at the `[edit forwarding-options sampling]` hierarchy level. For information on `cflowd`, see ["Enabling Flow Aggregation" on page 577](#). The `flow-active-timeout` and `flow-inactive-timeout` statements are described in ["Configuring Flow Monitoring" on page 5](#).

To direct sampled traffic to a flow-monitoring interface, include the `interface` statement. The `engine-id` and `engine-type` statements specify the accounting interface used on the traffic, and the `source-address` statement specifies the traffic source.

You cannot use rate-limiting with discard accounting; however, you can specify the duration of the interval for exporting aggregated accounting information by including the `aggregate-export-interval` statement in the configuration. This enables you to put a boundary on the amount of traffic exported to a flow-monitoring interface.

RELATED DOCUMENTATION

[Enabling Flow Aggregation | 577](#)

[Configuring Flow Monitoring | 5](#)

Sampling Data Using Inline Sampling

IN THIS CHAPTER

- [Understand Inline Active Flow Monitoring | 437](#)
- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 527](#)
- [Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers | 537](#)
- [Configuring Inline Active Flow Monitoring on PTX Series Routers | 540](#)
- [Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers | 550](#)
- [Inline Active Flow Monitoring on IRB Interfaces | 558](#)
- [Example: Configuring Inline Active Flow Monitoring on MX Series and T4000 Routers | 567](#)

Understand Inline Active Flow Monitoring

IN THIS SECTION

- [Benefits of Inline Active Flow Monitoring | 439](#)
- [Inline Active Flow Monitoring Configuration Overview | 440](#)
- [Inline Active Flow Monitoring Limitations and Restrictions | 441](#)
- [IPFIX and Version 9 Templates | 442](#)

Inline active flow monitoring is implemented on the Packet Forwarding Engine. The Packet Forwarding Engine performs functions such as creating flows, updating flows, and exporting flow records to a flow collector. The flow records are sent out in industry-standard IPFIX or version 9 format. IPFIX and version 9 templates use UDP as the transport protocol.

You can configure inline active flow monitoring for IPv4, IPv6, MPLS, MPLS-IPv4, VPLS, and bridge traffic. Starting in Junos OS Release 18.1R1, you can configure inline active flow monitoring for MPLS-over-UDP traffic for PTX3000 and PTX5000 Series routers. Starting in Junos OS Release 18.2R1, you

can configure inline active flow monitoring for MPLS, MPLS-IPv4, and MPLS-IPv6 traffic for PTX3000 and PTX5000 Series routers. Starting in Junos OS Release 18.2R1, you can configure inline active flow monitoring for bridge traffic for MX Series routers.

Starting in Junos OS Release 18.4R1, you can configure inline active flow monitoring for MPLS-IPv6 traffic for MX Series routers.

Starting with Junos OS Release 19.4R1 on the PTX10002-60C router, you can perform flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel. MPLS IPv4 and IPv6 payloads and IPFIX and version 9 templates are supported. Only ingress sampling is supported.

Starting with Junos OS Release 21.2R1 on the QFX10002-60C switch, you can perform inline active flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel. MPLS IPv4 and IPv6 payloads and IPFIX and version 9 templates are supported. Only ingress sampling is supported.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

For PTX Series, starting with Junos OS Evolved Release 21.2R1 and Junos OS Release 21.3R1, no flows are maintained. Every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, the number of records sent to the collector is higher than before. See [Table 63 on page 439](#). The IPFIX and version 9 Options Template Data Record now contains 0 in the Flow Active Timeout (Element ID 36) and Flow Inactive Timeout (Element ID 37) fields. Therefore, the Options Template Data Record is not compliant with IPFIX RFC 7011. The `show services accounting flow inline-jflow fpc-slot slot` operational mode command now displays 0 for all of the Active Flows and Timed Out fields. The various Total Flows fields are now equal to their respective Flow Packets fields. The various Flows Inactive Timed Out fields are now equal to their respective Flow Packets fields. The effect of the `nexthop-learning` statement at the `[edit services flow-monitoring version version template template-name]` hierarchy level on this no-flow behavior varies depending upon the operating system. For Junos OS Evolved, we do not recommend that you configure the `nexthop-learning` statement, as it reduces the number of packets that can be processed. For Junos OS, you can configure the `nexthop-learning` statement to change this default no-flow behavior and once again create and maintain flows, then attach the template to all sampling instances associated with FPCs that require the previous behavior.

Table 63: Inline Active Flow Monitoring Behavior Comparison for PTX Series

Actions	Prior to Junos OS Evolved Release 21.2R1 and Junos OS Release 21.3R1	Starting in Junos OS Evolved Release 21.2R1 and Junos OS Release 21.3R1
Flow creation	Flows are created and maintained.	No flows are created. Every packet is considered as a new flow for accounting purposes.
Active timeout	Active timeout configuration is honored. Active flows are timed out if the traffic is continuous. An export record is created for the timed-out flow and exported to the collector.	Active timeout configuration is ignored. No flows are timed out.
Inactive timeout	Inactive timeout configuration is honored. Inactive flows are timed out and are deleted at that time. An export record is created for the timed-out flow and exported to the collector.	Inactive timeout configuration is ignored. All flows are inactively timed out immediately.
Export records creation	Export records are created only during timeouts.	Export records are created for every sampled packet.
Packet export to collector	The configured active and inactive timeouts determine the packet export rates to the collector.	The packet export rate to the collector is directly proportional to sampling rate (in packets per second) at that given point in time. Because each packet results in an export record, the number of packets sent out to the collector does increase in comparison to what it was before.

Benefits of Inline Active Flow Monitoring

Inline active flow monitoring is implemented on the Packet Forwarding Engine rather than on a services card. This enables:

- Lower cost—You do not need to invest in additional hardware.

- Higher scalability—You do not need to dedicate a PIC slot for a services PIC, so you can make full use of the available slots for handling traffic on the device.
- Better performance—Inline flow monitoring performance is not dependent on the capacity of a services card.

Inline Active Flow Monitoring Configuration Overview

The inline active flow monitoring configuration can be broadly classified into four categories:

1. Configurations at the `[edit services flow-monitoring]` hierarchy level—At this level, you configure the template properties for inline flow monitoring.
2. Configurations at the `[edit forwarding-options]` hierarchy level—At this level, you configure a sampling instance and associate the template (configured at the `[edit services flow-monitoring]` hierarchy level) with the sampling instance. At this level, you also configure the flow-server IP address and port number as well as the flow export rate, and specify the collectors.

You cannot change the source IP address for collectors under the same family. Also, the template mapped across collectors under a family should be the same.

3. Configurations at the `[edit chassis]` hierarchy level—At this level, you associate the sampling instance with the FPC on which the media interface is present. If you are configuring sampling of IPv4 flows, IPv6 flows, or VPLS flows, you can configure the flow hash table size for each family.
4. Configurations at the `[edit firewall]` hierarchy level—At this level you configure a firewall filter for the family of traffic to be sampled. You must attach this filter to the interface on which you want to sample the traffic.

Before you configure inline active flow monitoring, ensure that you have adequately-sized hash tables for IPv4, IPv6, MPLS, and VPLS flow sampling. (VPLS flow sampling is Junos OS only.) These tables can use from one up to fifteen 256K areas. Starting with Junos OS Release 16.1R1 and 15.1F2, the IPv4 table is assigned a default value of 1024. Prior to Junos OS Release 16.1 and 15.1F2, the IPv4 table is assigned a default value of fifteen 256K areas. The IPv6 table is assigned a default value of 1024, and the VPLS table is assigned a default value of 1024. Allocate larger tables when anticipated traffic volume makes it necessary.

You can configure flow collectors to be reachable through non-default VPN routing and forwarding (VRF) instances by including the routing-instance *instance-name* statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output flow-server hostname]` hierarchy level for inline flow monitoring. You cannot configure a flow collector to be reachable through non-default VRF instances for version 5 and version 8 flows. You must configure the routing instance to be a VRF instance by including the instance-type vrf statement at the `[edit routing-instances instance-name]` hierarchy level.

Inline Active Flow Monitoring Limitations and Restrictions

The following limitations and restrictions apply to the inline active flow monitoring feature:

- Inline active flow monitoring is not supported for input or output traffic on MS-MPC or MS-MIC-16G interfaces.
- In Junos OS release 15.1 and earlier, you can apply version 9 flow templates to IPv4 traffic. Starting in Junos OS Release 16.1, you can also apply version 9 flow templates to MPLS and MPLS-IPv4 traffic. Starting in Junos OS Release 18.1R1, you can also apply version 9 flow templates to IPv6 traffic.
- In Junos OS Release 15.1 and earlier, you can apply IPFIX flow templates to IPv4, IPv6, and VPLS traffic. Starting in Junos OS release 16.1, you can also apply IPFIX flow templates to MPLS and MPLS-IPv4 traffic.
- Starting with Junos OS Release 17.2R1, you can apply IPFIX flow templates to unicast IPv4 and IPv6 traffic on QFX10002 switches. Starting with Junos OS Release 17.4R1, you can apply IPFIX flow templates to unicast IPv4 and IPv6 traffic on QFX10008 and QFX10016 switches.
- Inline active flow monitoring is not supported when you enable Next Gen Services on an MX Series router.
- You can configure only one sampling instance on a Flexible PIC Concentrator (FPC).
- You can configure only one type of sampling—either services-card-based sampling or inline sampling—per family in a sampling instance. However, you can configure services-card-based and inline sampling for different families in a sampling instance.
- The following considerations apply to the inline sampling instance configuration:
 - Sampling run-length and clip-size are not supported.
 - In Junos OS Release 16.2 and in Junos OS Release 16.1R3 and earlier, you can configure only one collector under a family for inline active flow monitoring. Starting with Junos OS Release 16.1R4 and 17.2R1, you can configure up to four collectors under a family for inline active flow monitoring. Starting with Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring. To configure a collector under a family for inline active flow monitoring, configure the `flow-server` statement at the `edit forwarding-options sampling-instance instance-name family (inet | inet6) output hierarchy level`. To specify up to four collectors, include up to four `flow-server` statements.

- The user-defined sampling instance gets precedence over the global instance. When a user-defined sampling instance is attached to the FPC, the global instance is removed from the FPC and the user-defined sampling instance is applied to the FPC.
- Flow records and templates cannot be exported if the flow collector is reachable through any management interface.
- If the destination of the sampled flow is reachable through multiple paths, the IP_NEXT_HOP (Element ID 15) and OUTPUT_SNMP (Element ID 14) in the IPv4 and IPv6 flow records are not reported correctly unless you enable learning of next hop addresses by using the `nexthop-learning enable` statement. (Starting in Junos OS Evolved Release 21.2R1 for PTX Series, we do not recommend that you enable learning of next-hop addresses, as it reduces the number of packets that can be processed. However, starting in Junos OS Release 21.3R1 for PTX Series, you can configure the `nexthop-learning` statement to change the default no-flow behavior and once again create and maintain flows, then attach the template to all sampling instances associated with FPCs that require the previous behavior.) If you do *not* configure `nexthop-learning enable`:
 - For IPv4 flow records, the IP_NEXT_HOP and OUTPUT_SNMP are set to the Gateway Address and SNMP Index of the first path seen in the forwarding table.
 - For IPv6 flow records, the IP_NEXT_HOP and OUTPUT_SNMP are set to 0.
- The Incoming Interface (IIF) and Outgoing Interface (OIF) should be part of the same VRF. If OIF is in a different VRF, DST_MASK (Element ID 13), DST_AS (Element ID 17), IP_NEXT_HOP (Element ID 15), and OUTPUT_SNMP (Element ID 14) are set to 0 in the flow records.
- Each lookup chip maintains and exports flows independent of other lookup chips. Traffic received on a media interface is distributed across all lookup chips in a multi-lookup chip platform. It is likely that a single flow is processed by multiple lookup chips. Therefore, each lookup chip creates a unique flow and exports it to the flow collector. This can cause duplicate flow records to go to the flow collector. The flow collector should aggregate PKTS_COUNT and BYTES_COUNT for duplicate flow records to derive a single flow record.

IPFIX and Version 9 Templates

Fields Included in the IPFIX Bridge Template for MX Series

[Table 64 on page 443](#) shows the fields that are included in the IPFIX Bridge template. The fields are shown in the order in which they appear in the template.

Table 64: IPFIX Bridge Template Fields for MX, M, and T Series

Field	Element ID
Destination MAC	80
Source MAC	56
Ethernet Type	256
Input SNMP	10
Output SNMP	14
Flow End Reason	136
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX IPv4 Template for MX, M, and T Series

[Table 65 on page 443](#) shows the fields that are included in the IPFIX IPv4 template. The fields are shown in the order in which they appear in the template.

Table 65: IPFIX IPv4 Template Fields for MX, M, and T Series

Field	Element ID
IPv4 Source Address	8

Table 65: IPFIX IPv4 Template Fields for MX, M, and T Series *(Continued)*

Field	Element ID
IPv4 Destination Address	12
IPv4 ToS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
VLAN ID	58
IPv4 Source Mask	9
IPv4 Destination Mask	13
Source AS	16
Destination AS	17
IPv4 Next Hop Address	15
TCP Flags	6
Output Interface	14

Table 65: IPFIX IPv4 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Minimum TTL	52
Maximum TTL	53
Flow End Reason	136
IP Protocol Version	60
BGP IPv4 Next Hop Address	18
Flow Direction (Starting in Junos OS Release 16.1)	61
802.1Q VLAN identifier (dot1qVlanId)	243
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX IPv4 Template for PTX3000 Series, PTX5000 Series, and the PTX10001-20C Router

[Table 66 on page 446](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 66: IPFIX IPv4 Template Fields for PTX3000 Series, PTX5000 Series, and the PTX10001-20C Router

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 TOS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
BGP IPv4 Next Hop Address	18
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22

Table 66: IPFIX IPv4 Template Fields for PTX3000 Series, PTX5000 Series, and the PTX10001-20C Router (Continued)

Field	Element ID
Time the flow ended with respect to system up time (FPC up time)	21
IPv4 Next Hop Address	15
IPv4 Source Mask	9
IPv4 Destination Mask	13
TCP Flags	6
IP Protocol Version	60
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
The type of interface where packets are being received. This field can have the following values: <ul style="list-style-type: none"> • 1—Other (default value) • 131—De-encapsulated GRE traffic is reported as <i>tunnel</i> 	368

Fields Included in the IPFIX IPv4 Template for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series

[Table 67 on page 448](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 67: IPFIX IPv4 Template Fields for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 TOS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
BGP IPv4 Next Hop Address	18
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22

Table 67: IPFIX IPv4 Template Fields for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series (Continued)

Field	Element ID
Time the flow ended with respect to system up time (FPC up time)	21
IPv4 Next Hop Address	15
IPv4 Source Mask	9
IPv4 Destination Mask	13
TCP Flags	6
IP Protocol Version	60
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
The type of interface where packets are being received. This field can have the following values: <ul style="list-style-type: none"> • 1—Other (default value) • 131—De-encapsulated GRE traffic is reported as <i>tunnel</i> 	368
Forwarding Class Name (first two bytes)	32767

Table 67: IPFIX IPv4 Template Fields for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series (*Continued*)

Field	Element ID
Packet Loss Priority; this field can have the following values: <ul style="list-style-type: none"> • 0x00: Low • 0x01: Medium-low • 0x02: Medium-high • 0x03: High • 0xFF: Unknown 	32766

Fields Included in the IPFIX IPv4 Template for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers

[Table 68 on page 450](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 68: IPFIX IPv4 Template Fields for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 TOS	5
IPv4 Protocol	4
Source Port	7

Table 68: IPFIX IPv4 Template Fields for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers (Continued)

Field	Element ID
Destination Port	11
ICMP Type and Code	32
Input SNMP Index	10
Source AS	16
Destination AS	17
BGP Next Hop Address	18
Output SNMP Index	14
Number of Bytes	1
Number of Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv4 Next Hop	15
IPv4 Source Mask	9
IPv4 Destination Mask	13
TCP Flags	6

Table 68: IPFIX IPv4 Template Fields for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers (Continued)

Field	Element ID
IP Protocol Version	60
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
<p>The type of interface where packets are being received. This field can have the following values:</p> <ul style="list-style-type: none"> • 1—Other (default value) • 131—De-encapsulated GRE traffic is reported as <i>tunnel</i> 	368
Source AS Path List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	16 (list of this type)
Destination AS Path List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	17 (list of this type)
BGP Source Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	484
BGP Destination Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	485
BGP Source Extended Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	487
BGP Destination Extended Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	488

Table 68: IPFIX IPv4 Template Fields for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers (Continued)

Field	Element ID
BGP Source Large Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	490
BGP Destination Large Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	491

Fields Included in the IPFIX IPv6 Template for MX, M, and T Series

[Table 69 on page 453](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 69: IPFIX IPv6 Template Fields for MX, M, and T Series

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	139

Table 69: IPFIX IPv6 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Input Interface	10
VLAN ID	58
IPv6 Source Mask	29
IPv6 Destination Mask	30
Source AS	16
Destination AS	17
IPv6 Next Hop Address	62
IPv6 BGP Next Hop Address	63
TCP Flags	6
Output Interface	14
Minimum Hop Limits	52
Maximum Hop Limits	53
Flow End Reason	136
Flow Direction (Starting in Junos OS Release 16.1)	61
802.1Q VLAN identifier (dot1qVlanId)	243

Table 69: IPFIX IPv6 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
IPv6 Option Headers	64
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX IPv6 Template for PTX3000 Series, PTX5000 Series, and the PTX10001-20C Router

[Table 70 on page 455](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 70: IPFIX IPv6 Template Fields for PTX3000 Series, PTX5000 Series, and the PTX10001-20C Router

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 TOS	5

Table 70: IPFIX IPv6 Template Fields for PTX3000 Series, PTX5000 Series, and the PTX10001-20C Router (Continued)

Field	Element ID
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code (IPv6)	139
Input Interface	10
IPv6 Source Mask	29
IPv6 Destination Mask	30
TCP Flags	6
Source AS	16
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Table 70: IPFIX IPv6 Template Fields for PTX3000 Series, PTX5000 Series, and the PTX10001-20C Router (Continued)

Field	Element ID
IPv6 Next Hop Address	62
IPv6 BGP NextHop Address	63
IP Protocol Version	60
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
The type of interface where packets are being received. This field can have the following values: <ul style="list-style-type: none"> • 1—Other (default value) • 131—De-encapsulated GRE traffic is reported as <i>tunnel</i> 	368

Fields Included in the IPFIX IPv6 Template for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series

Table 71 on page 457 shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 71: IPFIX IPv6 Template Fields for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28

Table 71: IPFIX IPv6 Template Fields for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series (Continued)

Field	Element ID
IPv6 TOS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code (IPv6)	139
Input Interface	10
IPv6 Source Mask	29
IPv6 Destination Mask	30
TCP Flags	6
Source AS	16
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22

Table 71: IPFIX IPv6 Template Fields for PTX1000, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series (*Continued*)

Field	Element ID
Time the flow ended with respect to system up time (FPC up time)	21
IPv6 Next Hop Address	62
IPv6 BGP Next Hop Address	63
IP Protocol Version	60
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
<p>The type of interface where packets are being received. This field can have the following values:</p> <ul style="list-style-type: none"> • 1—Other (default value) • 131—De-encapsulated GRE traffic is reported as <i>tunnel</i> 	368
Forwarding Class Name (first two bytes)	32767
<p>Packet Loss Priority; this field can have the following values:</p> <ul style="list-style-type: none"> • 0x00: Low • 0x01: Medium-low • 0x02: Medium-high • 0x03: High • 0xFF: Unknown 	32766

Fields Included in the IPFIX IPv6 Template for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers

Table 72 on page 460 shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 72: IPFIX IPv6 Template Fields for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 TOS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code (IPv6)	139
Input SNMP Index	10
Source AS	16
Destination AS	17
IPv6 BGP Next Hop Address	63
Output SNMP Index	14

Table 72: IPFIX IPv6 Template Fields for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers (Continued)

Field	Element ID
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv6 Next Hop Address	62
IPv6 Source Mask	29
IPv6 Destination Mask	30
TCP Flags	6
IP Protocol Version	60
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
<p>The type of interface where packets are being received. This field can have the following values:</p> <ul style="list-style-type: none"> • 1—Other (default value) • 131—De-encapsulated GRE traffic is reported as <i>tunnel</i> 	368
Source AS Path List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	16 (list of this type)

Table 72: IPFIX IPv6 Template Fields for PTX10001-36MR, PTX10003-160C, PTX10003-80C, PTX10004, and PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3) Routers (Continued)

Field	Element ID
Destination AS Path List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	17 (list of this type)
BGP Source Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	484
BGP Destination Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	485
BGP Source Extended Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	487
BGP Destination Extended Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	488
BGP Source Large Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	490
BGP Destination Large Community List (when configured on the data-record-fields statement at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level)	491

Fields Included in the IPFIX MPLS-IPv4 Template for MX, M, and T Series

Starting in Junos OS Release 16.1, the IPFIX MPLS-IPv4 template is supported. [Table 73 on page 463](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 73: IPFIX MPLS-IPv4 Template Fields for MX, M, and T Series

Field	Element ID
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Top Label IP Address	47
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 ToS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
VLAN ID	58
IPv4 Source Mask	9
IPv4 Destination Mask	13

Table 73: IPFIX MPLS-IPv4 Template Fields for MX, M, and T Series *(Continued)*

Field	Element ID
Source AS	16
Destination AS	17
IPv4 Next Hop Address	15
TCP Flags	6
Output Interface	14
Minimum TTL	52
Maximum TTL	53
Flow End Reason	136
IP Protocol Version	60
BGP IPv4 Next Hop Address	18
Flow Direction	61
802.1Q VLAN identifier (dot1qVlanId)	243
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
Number of Flow Bytes	1

Table 73: IPFIX MPLS-IPv4 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Number of Flow Packets	2
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX MPLS-IPv4 Template for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), and PTX10016 Series

Starting in Junos OS Release 18.2R1, the IPFIX MPLS-IPv4 template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-IPv4 template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-IPv4 template is supported for the QFX10002-60C switch. [Table 74 on page 465](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 74: IPFIX MPLS-IPv4 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 ToS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11

Table 74: IPFIX MPLS-IPv4 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
BGP IPv4 Next Hop Address	18
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
IPv4 Next Hop Address	15
IPv4 Source Mask	9
IPv4 Destination Mask	13

Table 74: IPFIX MPLS-IPv4 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
TCP Flags	6
IP Protocol Version	60
Ingress Interface Type	368
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Top Label IPv6 Address	140

Fields Included in the IPFIX MPLS-IPv4 Template for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR

[Table 75 on page 467](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 75: IPFIX MPLS-IPv4 Template Fields for PTX Series, for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12

Table 75: IPFIX MPLS-IPv4 Template Fields for PTX Series, for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR
(Continued)

Field	Element ID
IPv4 ToS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
BGP IPv4 Next Hop Address	18
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152

Table 75: IPFIX MPLS-IPv4 Template Fields for PTX Series, for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR (Continued)

Field	Element ID
Time the flow ended with respect to Epoch time	153
IPv4 Next Hop Address	15
IPv4 Source Mask	9
IPv4 Destination Mask	13
TCP Flags	6
IP Protocol Version	60
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Top Label IPv6 Address	140

Fields Included in the IPFIX MPLS-IPv6 Template for MX, M, and T Series

Starting in Junos OS Release 18.4R1, the IPFIX MPLS-IPv6 template is supported for the MX Series. [Table 76 on page 470](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 76: IPFIX MPLS-IPv6 Template Fields for MX, M, and T Series

Field	Element ID
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Top Label IP Address (Only IPv4 top label addresses are exported. IPv6 top label addresses report a value of zero.)	47
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code (IPv6)	139
Input Interface	10
VLAN ID	58
IPv6 Source Mask	29
IPv6 Destination Mask	30

Table 76: IPFIX MPLS-IPv6 Template Fields for MX, M, and T Series *(Continued)*

Field	Element ID
Source AS	16
Destination AS	17
IPv6 Next Hop Address	62
IPv6 BGP Next Hop Address	63
TCP Flags	6
Output Interface	14
Minimum TTL	52
Maximum TTL	53
Flow End Reason	136
Flow Direction	61
802.1Q VLAN identifier (dot1qVlanId)	243
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
IPv6 Option Headers	64
Number of Flow Bytes	1

Table 76: IPFIX MPLS-IPv6 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Number of Flow Packets	2
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX MPLS-IPv6 Template for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Starting in Junos OS Release 18.2R1, the IPFIX MPLS-IPv6 template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-IPv6 template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-IPv6 template is supported for the QFX10002-60C switch. [Table 77 on page 472](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 77: IPFIX MPLS-IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11

Table 77: IPFIX MPLS-IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
ICMP Type and Code (IPv6)	139
Input Interface	10
Source AS	16
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
IPv6 Source Mask	29
IPv6 Destination Mask	30
IPv6 Next Hop Address	62

Table 77: IPFIX MPLS-IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
IPv6 BGP Next Hop Address	63
TCP Flags	6
IP protocol version of IP payload on MPLS VPN	60
Ingress Interface Type	368
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
MPLS Label 3	72
MPLS Top Label IPv6 Address	140

Fields Included in the IPFIX MPLS-IPv6 Template for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR

[Table 78 on page 474](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 78: IPFIX MPLS-IPv6 Template Fields for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR

Field	Element ID
IPv6 Source Address	27

Table 78: IPFIX MPLS-IPv6 Template Fields for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR (Continued)

Field	Element ID
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code (IPv6)	139
Input Interface	10
Source AS	16
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152

Table 78: IPFIX MPLS-IPv6 Template Fields for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR (Continued)

Field	Element ID
Time the flow ended with respect to Epoch time	153
IPv6 Source Mask	29
IPv6 Destination Mask	30
IPv6 Next Hop Address	62
IPv6 BGP Next Hop Address	63
TCP Flags	6
IP protocol version of IP payload on MPLS VPN	60
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
MPLS Label 3	72
MPLS Top Label IPv6 Address	140

Fields Included in the IPFIX MPLS Template for MX, M, and T Series

Starting in Junos OS Release 16.1, the IPFIX MPLS template is supported. [Table 79 on page 477](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 79: IPFIX MPLS Template Fields for MX, M, and T Series

Field	Element ID
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
Input Interface	10
Output Interface	14
Flow End Reason	136
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX MPLS Template for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Starting in Junos OS Release 18.2R1, the IPFIX MPLS template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS template is supported for the QFX10002-60C switch. [Table 80 on page 478](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 80: IPFIX MPLS Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Field	Element ID
Input Interface	10
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
Ingress Interface Type	368
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72

Fields Included in the IPFIX MPLS Template for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR

[Table 81 on page 479](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 81: IPFIX MPLS Template Fields for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR

Field	Element ID
Input Interface	10
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72

Fields Included in the IPFIX MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Within an IP Network and Having an IPv4 Payload

Starting in Junos OS Release 18.1R1, the IPFIX MPLS-over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

[Table 82 on page 480](#) shows the fields that are available in the IPFIX template for MPLS-over-UDP flows that are within an IP network and have an IPv4 payload. The fields are shown in the order in which they appear in the template.

Table 82: IPFIX MPLS-over-UDP Carried on IP Network Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch

Field	Element ID
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12
UDP source port for tunnel endpoint	7
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15
BGP IPv4 next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14

Table 82: IPFIX MPLS-over-UDP Carried on IP Network Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch *(Continued)*

Field	Element ID
MPLS label 1—VPN bottom of stack label	70
IP protocol version of IP payload on MPLS VPN	60
IPv4 source address of tunnel payload	8
IPv4 destination address of tunnel payload	12
IP protocol of tunnel payload	4
IP TOS	5
Source transport port	7
Destination transport port	11
ICMP type	32
TCP flags	6
Number of flow bytes	1
Number of flow packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152

Table 82: IPFIX MPLS-over-UDP Carried on IP Network Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch (Continued)

Field	Element ID
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Encapsulated in an RSVP-TE LSP and Having an IPv4 Payload

Starting in Junos OS Release 18.1R1, the IPFIX MPLS-over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

[Table 83 on page 482](#) shows the fields that are available in the IPFIX template for MPLS-over-UDP flows that are encapsulated in an RSVP-TE LSP in the inner MPLS network and have an IPv4 payload. The fields are shown in the order in which they appear in the template.

Table 83: IPFIX MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch

Field	Element ID
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12
UDP source port for tunnel endpoint	7

Table 83: IPFIX MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch *(Continued)*

Field	Element ID
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15
BGP IPv4 next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14
MPLS label 1—VPN bottom of stack label	70
IP protocol version of IP payload on MPLS VPN	60
IPv4 source address of tunnel payload	8
IPv4 destination address of tunnel payload	12
IP protocol of tunnel payload	4
IP TOS	5

Table 83: IPFIX MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch (Continued)

Field	Element ID
Source transport port	7
Destination transport port	11
ICMP type	32
TCP flags	6
Number of flow bytes	1
Number of flow packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Within an IP Network Having an IPv6 Payload

Starting in Junos OS Release 18.1R1, the IPFIX MPLS-Over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

Table 84 on page 485 shows the fields that are available in the IPFIX template for MPLS-over-UDP flows that are within an IP network and have an IPv6 payload. The fields are shown in the order in which they appear in the template.

Table 84: IPFIX MPLS-over-UDP Carried on IP Network Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch

Field	Element ID
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12
UDP source port for tunnel endpoint	7
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15
BGP next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14
MPLS label 1—VPN bottom of stack label	70

Table 84: IPFIX MPLS-over-UDP Carried on IP Network Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch *(Continued)*

Field	Element ID
IP protocol version of IP payload on MPLS VPN	60
IPv6 source address of tunnel payload	27
IPv6 destination address of tunnel payload	28
IP protocol of tunnel payload	4
IP TOS	5
Source transport port	7
Destination transport port	11
ICMP type V6	139
TCP flags	6
Number of flow bytes	1
Number of flow packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Encapsulated in an RSVP-TE LSP and Having an IPv6 Payload

Starting in Junos OS Release 18.1R1, the IPFIX MPLS-Over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

[Table 85 on page 487](#) shows the fields that are available in the IPFIX template for MPLS-over-UDP flows that are encapsulated in an RSVP-TE LSP in the inner MPLS network and have an IPv6 payload. The fields are shown in the order in which they appear in the template.

Table 85: IPFIX MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch

Field	Element ID
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12
UDP source port for tunnel endpoint	7
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16

Table 85: IPFIX MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch *(Continued)*

Field	Element ID
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15
BGP next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14
MPLS label 1—VPN bottom of stack label	70
IP protocol version of IP payload on MPLS VPN	60
IPv6 source address of tunnel payload	27
IPv6 destination address of tunnel payload	28
IP protocol of tunnel payload	4
IP TOS	5
Source transport port	7
Destination transport port	11
ICMP type V6	139
TCP flags	6

Table 85: IPFIX MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch (Continued)

Field	Element ID
Number of flow bytes	1
Number of flow packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the IPFIX VPLS Template for MX, M, and T Series

Starting in Junos OS Release 16.1, the IPFIX VPLS template is supported. [Table 86 on page 489](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 86: IPFIX VPLS Template Fields for MX, M, and T Series

Field	Element ID
Destination MAC	80
Source MAC	56
Ethernet Type	256
Input Interface	10

Table 86: IPFIX VPLS Template Fields for MX, M, and T Series *(Continued)*

Field	Element ID
Output Interface	14
Flow End Reason	136
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to Epoch time	152
Time the flow ended with respect to Epoch time	153

Fields Included in the Version 9 Bridge Template for MX Series

[Table 87 on page 490](#) shows the fields that are included in the version 9 Bridge template. The fields are shown in the order in which they appear in the template.

Table 87: Version 9 Bridge Template Fields for MX

Field	Element ID
Destination MAC	80
Source MAC	56
Ethernet Type	256
Input SNMP	10
Output SNMP	14

Table 87: Version 9 Bridge Template Fields for MX (Continued)

Field	Element ID
Flow End Reason	136
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to Epoch time	22
Time the flow ended with respect to Epoch time	21

Fields Included in the Version 9 IPv4 Template for MX, M, and T Series

[Table 88 on page 491](#) shows the fields that are included in the version 9 IPv4 template. The fields are shown in the order in which they appear in the template.

Table 88: Version 9 IPv4 Template Fields for MX, M, and T Series

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 ToS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11

Table 88: Version 9 IPv4 Template Fields for MX, M, and T Series *(Continued)*

Field	Element ID
ICMP Type and Code	32
Input Interface	10
VLAN ID	58
IPv4 Source Mask	9
IPv4 Destination Mask	13
Source AS	16
Destination AS	17
IPv4 Next Hop Address	15
TCP Flags	6
Output Interface	14
Minimum TTL	52
Maximum TTL	53
Flow End Reason	136
Internet Protocol Version	60
BGP IPv4 Next Hop Address	18

Table 88: Version 9 IPv4 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Flow Direction	61
802.1Q VLAN identifier (dot1qVlanId)	243
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 IPv4 Template for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

[Table 89 on page 493](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 89: Version 9 IPv4 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12

Table 89: Version 9 IPv4 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
IPv4 TOS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
BGP IPv4 Next Hop Address	18
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv4 Next Hop Address	15

Table 89: Version 9 IPv4 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
IPv4 Source Mask	9
IPv4 Destination Mask	13
TCP Flags	6
IP Protocol Version	60

Fields Included in the Version 9 IPv4 Template for PTX10003-160C, PTX10003-80C, PTX10004, PTX10008 (with the JNP10008-SF3), and PTX10001-36MR Routers

[Table 90 on page 495](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 90: Version 9 IPv4 Template Fields for PTX10003-160C, PTX10003-80C, PTX10004, PTX10008 (with the JNP10008-SF3), and PTX10001-36MR Routers

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 TOS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11

Table 90: Version 9 IPv4 Template Fields for PTX10003-160C, PTX10003-80C, PTX10004, PTX10008 (with the JNP10008-SF3), and PTX10001-36MR Routers (Continued)

Field	Element ID
ICMP Type and Code	32
Input SNMP Index	10
IPv4 Source Mask	9
IPv4 Destination Mask	13
Source AS	16
Destination AS	17
IPv4 Next Hop Address	15
BGP IPv4 Next Hop Address	18
TCP Flags	6
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IP Protocol Version	60
Output SNMP Index	14

Fields Included in the Version 9 IPv6 Template for MX, M, and T Series

Starting in Junos OS Release 18.1R1, the version 9 IPv6 template is supported. [Table 91 on page 497](#) shows the fields in the template. The fields are shown in the order in which they appear in the template.

Table 91: Version 9 IPv6 Template Fields for MX, M, and T Series

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	139
Input Interface	10
VLAN ID	58
IPv6 Source Mask	29
IPv6 Destination Mask	30
Source AS	16
Destination AS	17

Table 91: Version 9 IPv6 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
IPv6 Next Hop Address	62
IPv6 BGP Next Hop Address	63
TCP Flags	6
Output Interface	14
Minimum TTL	52
Maximum TTL	53
Flow End Reason	136
Flow Direction	61
802.1Q VLAN identifier (dot1qVlanId)	243
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
IPv6 Option Headers	64
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22

Table 91: Version 9 IPv6 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 IPv6 Template for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

[Table 92 on page 499](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 92: IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3SIB), PTX10016 Series

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 TOS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
Source AS	16

Table 92: IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3SIB), PTX10016 Series (*Continued*)

Field	Element ID
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv6 Next Hop Address	62
IPv6 BGP NextHop Address	63
IPv6 Source Mask	29
IPv6 Destination Mask	30
TCP Flags	6
IP Protocol Version	60

Fields Included in the Version 9 IPv6 Template for PTX10003-160C, PTX10003-80C, PTX10004, PTX10008 (with the JNP10008-SF3), and PTX10001-36MR routers

[Table 93 on page 501](#) shows the fields that are available in the template. The fields are shown in the order in which they appear in the template.

Table 93: IPv6 Template Fields for PTX10003-160C, PTX10003-80C, PTX10004, PTX10008 (with the JNP10008-SF3), and PTX10001-36MR routers

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 TOS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input SNMP Index	10
Output SNMP Index	14
IPv6 Source Mask	29
IPv6 DestinationMask	30
Source AS	16
Destination AS	17
IPv6 Next Hop Address	62
TCP Flags	6

Table 93: IPv6 Template Fields for PTX10003-160C, PTX10003-80C, PTX10004, PTX10008 (with the JNP10008-SF3), and PTX10001-36MR routers (Continued)

Field	Element ID
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IP Protocol Version	60

Fields Included in the Version 9 MPLS-IPv4 Template for MX, M, and T Series

Starting in Junos OS Release 16.1, the version 9 MPLS-IPv4 template is supported. [Table 94 on page 502](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 94: Version 9 MPLS-IPv4 Template Fields for MX, M, and T Series

Field	Element ID
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Top Label IP Address	47
IPv4 Source Address	8

Table 94: Version 9 MPLS-IPv4 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
IPv4 Destination Address	12
IPv4 ToS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
VLAN ID	58
IPv4 Source Mask	9
IPv4 Destination Mask	13
Source AS	16
Destination AS	17
IPv4 Next Hop Address	15
TCP Flags	6
Output Interface	14

Table 94: Version 9 MPLS-IPv4 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Minimum TTL	52
Maximum TTL	53
Flow End Reason	136
IP Protocol Version	60
BGP IPv4 Next Hop Address	18
Flow Direction	61
802.1Q VLAN identifier (dot1qVlanId)	243
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 MPLS-IPv4 Template for PTX Series and the QFX10002-60C Switch

Starting in Junos OS Release 18.2R1, the version 9 MPLS-IPv4 template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the version 9 MPLS-IPv4 template is supported for the

PTX1002-60C router. Starting in Junos OS Release 21.2R1, the version 9 MPLS-IPv4 template is supported for the QFX10002-60C switch. [Table 95 on page 505](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 95: Version 9 MPLS-IPv4 Template Fields for PTX Series and the QFX10002-60C Switch

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 ToS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
BGP IPv4 Next Hop Address	18
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2

Table 95: Version 9 MPLS-IPv4 Template Fields for PTX Series and the QFX10002-60C Switch
(Continued)

Field	Element ID
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv4 Next Hop Address	15
IPv4 Source Mask	9
IPv4 Destination Mask	13
TCP Flags	6
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Top Label IP Address	47

Fields Included in the Version 9 MPLS-IPv6 Template for MX, M, and T Series

Starting in Junos OS Release 18.4R1, the version 9 MPLS-IPv6 template is supported for the MX Series. [Table 96 on page 507](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 96: Version 9 MPLS-IPv6 Template Fields for MX, M, and T Series

Field	Element ID
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
MPLS Top Label IP Address (Only IPv4 top label addresses are exported. IPv6 top label addresses report a value of zero.)	47
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code (IPv6)	139
Input Interface	10
VLAN ID	58
IPv6 Source Mask	29
IPv6 Destination Mask	30

Table 96: Version 9 MPLS-IPv6 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Source AS	16
Destination AS	17
IPv6 Next Hop Address	62
IPv6 BGP Next Hop Address	63
TCP Flags	6
Output Interface	14
Minimum TTL	52
Maximum TTL	53
Flow End Reason	136
Flow Direction	61
802.1Q VLAN identifier (dot1qVlanId)	243
802.1Q Customer VLAN identifier (dot1qCustomerVlanId)	245
IP Identifier	54
IPv6 Option Headers	64
Number of Flow Bytes	1

Table 96: Version 9 MPLS-IPv6 Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 MPLS-IPv6 Template for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Starting in Junos OS Release 18.2R1, the version 9 MPLS-IPv6 template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the version 9 MPLS-IPv6 template is supported for the PTX1002-60C router. Starting in Junos OS Release 21.2R1, the version 9 MPLS-IPv6 template is supported for the QFX10002-60C switch. [Table 97 on page 509](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 97: Version 9 MPLS-IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11

Table 97: Version 9 MPLS-IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
ICMP Type and Code (IPv6)	32
Input Interface	10
Source AS	16
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv6 Source Mask	29
IPv6 Destination Mask	30
IPv6 Next Hop Address	62
IPv6 BGP NextHop Address	63
TCP Flags	6

Table 97: Version 9 MPLS-IPv6 Template Fields for PTX3000, PTX5000, PTX1000, PTX10001-20C, PTX10002-60C, QFX10002-60C, PTX10008 (without the JNP10008-SF3), PTX10016 Series
(Continued)

Field	Element ID
IP protocol version of IP payload on MPLS VPN	60
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
MPLS Label 3	72
MPLS Top Label IP Address	47

Fields Included in the Version 9 MPLS-IPv6 Template for PTX10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-LC1202 line card and the JNP10008-SF3), and PTX10001-36MR

[Table 98 on page 511](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 98: Version 9 MPLS-IPv6 Template Fields for PTX 10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-1202-36MR line card and the JNP10008-SF3), and PTX10001-36MR Routers

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 ToS	5
IPv6 Protocol	4

Table 98: Version 9 MPLS-IPv6 Template Fields for PTX 10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-1202-36MR line card and the JNP10008-SF3), and PTX10001-36MR Routers (Continued)

Field	Element ID
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code (IPv6)	32
Input Interface	10
Source AS	16
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv6 Source Mask	29
IPv6 Destination Mask	30
IPv6 Next Hop Address	62
TCP Flags	6

Table 98: Version 9 MPLS-IPv6 Template Fields for PTX 10003, PTX10004, PTX10008 (with the JNP10K-LC1201 or JNP10K-1202-36MR line card and the JNP10008-SF3), and PTX10001-36MR Routers (Continued)

Field	Element ID
IP protocol version of IP payload on MPLS VPN	60
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
MPLS Label 3	72
MPLS Top Label IP Address	47

Fields Included in the Version 9 MPLS Template for MX, M, and T Series

Starting in Junos OS Release 16.1, the version 9 MPLS template is supported. [Table 99 on page 513](#) shows the fields that are included in the template. The fields are shown in the order in which they appear in the template.

Table 99: Version 9 MPLS Template Fields for MX, M, and T Series

Field	Element ID
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72
Input Interface	10
Output Interface	14

Table 99: Version 9 MPLS Template Fields for MX, M, and T Series (Continued)

Field	Element ID
Flow End Reason	136
Number of Flow Bytes	1
Number of Flow Packets	2
First Switched	ww
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 MPLS Template for PTX Series and the QFX10002-60C Switch

Starting in Junos OS Release 18.2R1, the version 9 MPLS template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS template is supported for the QFX10002-60C switch. [Table 100 on page 514](#) shows the fields that are included in the template.

Table 100: Version 9 MPLS Template Fields for PTX Series and the QFX10002-60C Switch

Field	Element ID
Input Interface	10
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22

Table 100: Version 9 MPLS Template Fields for PTX Series and the QFX10002-60C Switch (Continued)

Field	Element ID
Time the flow ended with respect to system up time (FPC up time)	21
MPLS Label 1	70
MPLS Label 2	71
MPLS Label 3	72

Fields Included in the Version 9 MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Within an IP Network Having an IPv4 Payload

Starting in Junos OS Release 18.1R1, the version 9 MPLS-over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

[Table 101 on page 515](#) shows the fields that are available in the Version 9 template for MPLS-over-UDP flows that are within an IP network and have an IPv4 payload. The fields are shown in the order in which they appear in the template.

Table 101: Version 9 MPLS-over-UDP Carried on IP Network Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch

Field	Element ID
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12
UDP source port for tunnel endpoint	7

Table 101: Version 9 MPLS-over-UDP Carried on IP Network Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch (Continued)

Field	Element ID
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15
BGP IPv4 next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14
MPLS label 1—VPN bottom of stack label	70
IP protocol version of IP payload on MPLS VPN	60
IPv4 source address of tunnel payload	8
IPv4 destination address of tunnel payload	12
IP protocol of tunnel payload	4
IP TOS	5

Table 101: Version 9 MPLS-over-UDP Carried on IP Network Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch (Continued)

Field	Element ID
Source transport port	7
Destination transport port	11
ICMP type	32
TCP flags	6
Number of flow bytes	1
Number of flow packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Encapsulated in an RSVP-TE LSP and Having an IPv4 Payload

Starting in Junos OS Release 18.1R1, the version 9 MPLS-over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

[Table 102 on page 518](#) shows the fields that are available in the Version 9 template for MPLS-over-UDP flows that are encapsulated in an RSVP-TE LSP in the inner MPLS network and have an IPv4 payload. The fields are shown in the order in which they appear in the template.

Table 102: Version 9 MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch

Field	Element ID
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12
UDP source port for tunnel endpoint	7
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15
BGP IPv4 next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14
MPLS label 1—VPN bottom of stack label	70

Table 102: Version 9 MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv4 Payload) for PTX Series and the QFX10002-60C Switch (Continued)

Field	Element ID
IP protocol version of IP payload on MPLS VPN	60
IPv4 source address of tunnel payload	8
IPv4 destination address of tunnel payload	12
IP protocol of tunnel payload	4
IP TOS	5
Source transport port	7
Destination transport port	11
ICMP type	32
TCP flags	6
Number of flow bytes	1
Number of flow packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Within an IP Network Having an IPv6 Payload

Starting in Junos OS Release 18.1R1, the version 9 MPLS-over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

[Table 103 on page 520](#) shows the fields that are available in the Version 9 template for MPLS-over-UDP flows that are within an IP network and have an IPv6 payload. The fields are shown in the order in which they appear in the template.

Table 103: Version 9 MPLS-over-UDP Carried on IP Network Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch

Field	Element ID
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12
UDP source port for tunnel endpoint	7
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15

Table 103: Version 9 MPLS-over-UDP Carried on IP Network Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch *(Continued)*

Field	Element ID
BGP next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14
MPLS label 1—VPN bottom of stack label	70
IP protocol version of IP payload on MPLS VPN	60
IPv6 source address of tunnel payload	27
IPv6 destination address of tunnel payload	28
IP protocol of tunnel payload	4
IP TOS	5
Source transport port	7
Destination transport port	11
ICMP type V6	32
TCP flags	6
Number of flow bytes	1
Number of flow packets	2

Table 103: Version 9 MPLS-over-UDP Carried on IP Network Template Fields (IPv6 Payload) for PTX Series and the QFX10002-60C Switch *(Continued)*

Field	Element ID
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Fields Included in the Version 9 MPLS-over-UDP Template for PTX Series and the QFX10002-60C Switch for Flows Encapsulated in an RSVP-TE LSP and Having an IPv6 Payload

Starting in Junos OS Release 18.1R1, the version 9 MPLS-over-UDP template is supported for the PTX Series. Starting in Junos OS Release 19.4R1, the IPFIX MPLS-over-UDP template is supported for the PTX10002-60C router. Starting in Junos OS Release 21.2R1, the IPFIX MPLS-over-UDP template is supported for the QFX10002-60C switch.

Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

[Table 104 on page 522](#) shows the fields that are available in the Version 9 template for MPLS-over-UDP flows that are encapsulated in an RSVP-TE LSP in the inner MPLS network and have an IPv6 payload. The fields are shown in the order in which they appear in the template.

Table 104: Version 9 MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv6 Payload) for PTX Series and QFX10002-60C Switch

Field	Element ID
RSVP label (top MPLS label stack entry) for MPLS tunnel	70
RSVP label pushed before top label	71
IPv4 source address for tunnel endpoint	8
IPv4 destination address for tunnel endpoint	12

Table 104: Version 9 MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv6 Payload) for PTX Series and QFX10002-60C Switch (Continued)

Field	Element ID
UDP source port for tunnel endpoint	7
Tunnel endpoint destination transport port	11
IPv4 source mask for tunnel source IP address	9
IPv4 destination mask for tunnel destination IP address	13
Source AS for tunnel	16
Destination AS for tunnel	17
IPv4 next hop address—gateway for tunnel destination IP address	15
BGP next hop address—tunnel destination IP BGP peer	18
Input SNMP index	10
Output SNMP index	14
MPLS label 1—VPN bottom of stack label	70
IP protocol version of IP payload on MPLS VPN	60
IPv6 source address of tunnel payload	27
IPv6 destination address of tunnel payload	28
IP protocol of tunnel payload	4

Table 104: Version 9 MPLS-over-UDP Encapsulated in RSVP-TE LSP Template Fields (IPv6 Payload) for PTX Series and QFX10002-60C Switch (Continued)

Field	Element ID
IP TOS	5
Source transport port	7
Destination transport port	11
ICMP type	32
TCP flags	6
Number of flow bytes	1
Number of flow packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21

Release History Table

Release	Description
22.4R1	Starting in Junos OS Release 22.4R1 for the MX240, MX480, MX960, MX10004, MX10008, MX2008, MX2010, and MX2020 routers, we support ingress and egress sampling of IPv4, IPv6, and MPLS traffic on abstracted fabric (af) interfaces between guest network functions (GNFs) in a node slicing scenario, for both the IPFIX and version 9 export formats.
22.2R1-EVO	Starting in Junos OS Evolved Release 22.2R1 for the PTX10003 router, Information Element 63, IPv6 BGP NextHop Address, is available in both the IPv6 template and the MPLS-IPv6 template for the IPFIX and version 9 export formats.

21.4R1-EVO	Starting in Junos OS Evolved Release 21.4R1 for PTX Series, you can export BGP community and AS path information using IP Flow Information Export (IPFIX) information elements 483 through 491, 16, and 17, per RFCs 8549 and 6313. Content providers can use this information to identify a transit service provider degrading the quality of the service. You configure these elements with the statement data-record-fields at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level.
21.3R1-EVO	Starting in Junos OS Evolved Release 21.3R1 for the PTX10001-36MR, PTX10004, and PTX10008 routers, Information Element 63, IPv6 BGP NextHop Address, is available in both the IPv6 template and the MPLS-IPv6 template for the IPFIX and version 9 export formats.
21.3R1	Starting with Junos OS Release 21.3R1 for PTX Series routers, no flows are maintained. Every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, the number of records sent to the collector is higher than before.
21.2R1-Evo	Starting with Junos OS Evolved 21.2R1, no flows are maintained. Every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, the number of records sent to the collector is higher than before.
21.2R1	Starting with Junos OS Release 21.2R1 on the QFX10002-60C switch, you can perform inline active flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel. MPLS IPv4 and IPv6 payloads and IPFIX and version 9 templates are supported. Only ingress sampling is supported.
21.1R1-EVO	Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring.
20.4R1-EVO	Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
20.3R1-EVO	Starting with Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
19.4R1	Starting with Junos OS Release 19.4R1 on the PTX10002-60C router, you can perform flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel. MPLS IPv4 and IPv6 payloads and IPFIX and version 9 templates are supported. Only ingress sampling is supported.

19.3R2	Inline active flow monitoring is not supported when you enable Next Gen Services on an MX Series router.
19.2R1	Starting in Junos OS Release 19.2R1 for MX and PTX Series routers, Information Element 63, IPv6 BGP NextHop Address, is available in both the IPv6 template and the MPLS-IPv6 template for the IPFIX and version 9 export formats.
18.4R1	Starting in Junos OS Release 18.4R1, you can configure inline active flow monitoring for MPLS-IPv6 traffic for MX Series routers.
18.4R1	Starting in Junos OS Release 18.4R1, the IPFIX and version 9 MPLS-IPv6 templates are supported for the MX Series.
18.2R1	Starting in Junos OS Release 18.2R1, you can configure inline active flow monitoring for MPLS, MPLS-IPv4, and MPLS-IPv6 traffic for PTX3000 and PTX5000 Series routers.
18.2R1	Starting in Junos OS Release 18.2R1, you can configure inline active flow monitoring for bridge traffic for MX Series routers.
18.2R1	Starting in Junos OS Release 18.2R1, the IPFIX and version 9 MPLS templates are supported for the PTX Series.
18.2R1	Starting in Junos OS Release 18.2R1, the IPFIX and version 9 MPLS-IPv4 templates are supported for the PTX Series.
18.2R1	Starting in Junos OS Release 18.2R1, the IPFIX and version 9 MPLS-IPv6 templates are supported for the PTX Series.
18.1R1	Starting in Junos OS Release 18.1R1, you can also apply version 9 flow templates to IPv6 traffic.
18.1R1	Starting in Junos OS Release 18.1R1, you can configure inline active flow monitoring for MPLS-over-UDP traffic for PTX3000 and PTX5000 Series routers.
18.1R1	Starting in Junos OS Release 18.1R1, the IPFIX and version 9 MPLS-over-UDP templates are supported for the PTX Series.
17.4R1	Starting with Junos OS Release 17.4R1, you can apply IPFIX flow templates to unicast IPv4 and IPv6 traffic on QFX10008 and QFX10016 switches.
17.2R1	Starting with Junos OS Release 17.2R1, you can apply IPFIX flow templates to unicast IPv4 and IPv6 traffic on QFX10002 switches.

16.1R4	In Junos OS Release 16.2 and in Junos OS Release 16.1R3 and earlier, you can configure only one collector under a family for inline active flow monitoring. Starting with Junos OS Release 16.1R4 and 17.2R1, you can configure up to four collectors under a family for inline active flow monitoring.
16.1R1	Starting with Junos OS Release 16.1R1 and 15.1F2, the IPv4 table is assigned a default value of 1024.
16.1R1	Starting in Junos OS Release 16.1R1, you can also apply IPFIX and version 9 flow templates to MPLS and MPLS-IPv4 traffic.
16.1R1	Flow Direction (Starting in Junos OS Release 16.1R1)
16.1R1	Starting in Junos OS Release 16.1R1, the IPFIX VPLS template is supported.

RELATED DOCUMENTATION

[Example: Configuring Inline Active Flow Monitoring on MX Series and T4000 Routers | 567](#)

[Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers | 537](#)

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

Inline active flow monitoring is implemented on the Packet Forwarding Engine. The Packet Forwarding Engine performs functions such as creating and updating flows, and updating flow records. The flow records are sent out in industry-standard IPFIX or version 9 format.

Support for active flow monitoring with IPFIX templates on QFX10002 switches was added in Junos OS Release 17.2R1. Starting in Junos OS Release 20.3R1 on QFX10002-60C switches, you can configure inline active flow monitoring for IPv4 and IPv6 traffic. Both IPFIX and version 9 templates are supported.

On routers with MS-PICs or MS-DPCs, IPv4 and IPv6 fragments are processed accurately. The flow monitoring application creates two flows for every fragmented flow. The first fragment that has the complete Layer 4 information forms the first flow with 5-tuple data and subsequently, all the fragmented packets related to this flow form another flow with the Layer 4 fields set to zero.

The following considerations apply to the inline flow-monitoring instance configuration:

- Sampling run-length and clip-size are not supported.

- For inline configurations, collectors are not reachable via `fxp0`.
- Inline flow monitoring does not support `cflowd`. Therefore, inline flow monitoring does not support the local dump option, which is available only with `cflowd`.
- Inline active flow monitoring is not supported when you enable Next Gen Services on an MX Series router.
- The number of collectors that are supported depends on the device:
 - In Junos OS Release 16.2 and in Junos OS Release 16.1R3 and earlier, you can configure only one collector under a family for inline active flow monitoring. Starting with Junos OS Release 16.1R4 and 17.2R1, you can configure up to four collectors under a family for inline active flow monitoring. Starting with Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring. The Packet Forwarding Engine (PFE) can export the flow record, flow record template, option data, and option data template packet to all configured collectors. To configure a collector under a family for inline active flow monitoring, configure the `flow-server` statement at the `edit forwarding-options sampling-instance instance-name family (inet | inet6) output` hierarchy level. To specify up to four collectors, include up to four `flow-server` statements.
 - For inline configurations on all other devices, each family can support only one collector.

Inline active flow monitoring is available in four hierarchy levels:

- `[edit chassis]`—At this level, you associate the sampling instance with the FPC on which the media interface is present (except on the MX80 and MX104—see ["Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers" on page 537](#)). If you are configuring sampling of IPv4 flows, IPv6 flows or VPLS flows (Junos OS only), you can configure the flow hash table size for each family, as described below.
- `[edit firewall]`—At this level, you configure a firewall filter for the family of traffic to be sampled. You must attach this filter to the interface on which you want to sample the traffic.
- `[edit forwarding-options]`—At this level, you configure a sampling instance and associate the template with the sampling instance. At this level, you also configure the flow-server IP address and port number as well as the flow export rate.
- `[edit services flow-monitoring]`—At this level, you configure the template properties for inline flow monitoring.

Before you configure inline active flow monitoring, you should ensure that you have adequately-sized hash tables for IPv4, IPv6, MPLS, and VPLS flow sampling. (VPLS flow sampling is Junos OS only). These tables can use one to fifteen 256K areas. Starting with Junos OS Release 16.1R1 and 15.1F2, the IPv4 table is assigned a default value of 1024. Prior to Junos OS Release 16.1 and 15.1F2, the IPv4 table is assigned a default value of fifteen 256K areas. The IPv6 table is assigned a default value of 1024, and the VPLS table is assigned a default value of 1024. When anticipated traffic volume requires larger tables, allocate larger tables.

To allocate flow hash tables:

1. Go to the [edit chassis fpc 0 inline-services flow-table-size] hierarchy level for inline services on the FPC that processes the monitored flows.

```
[edit]
user@host# edit chassis fpc 0 inline-services flow-table-size
```

2. Specify the required sizes for the sampling hash tables.

```
[edit chassis fpc 0 inline-services flow-table-size]
user@host# set bridge-flow-table-size units
user@host# set ipv4-flow-table-size units
user@host# set ipv6-flow-table-size units
user@host# set mpls-flow-table-size units
user@host# set vpls-flow-table-size units
```

NOTE: Starting in Junos OS Release 18.2R1, the bridge-flow-table-size option is available and the vpls-flow-table-size option is deprecated; use the bridge-flow-table-size option instead. The bridge-flow-table-size option supports both VPLS and bridge records.

NOTE: The total number of units used for IPv4, IPv6, MPLS, and VPLS cannot exceed 15. Also, starting in Junos OS Release 16.1R1 and 15.1F2, changing the flow hash table size does *not* automatically reboot the FPC (for earlier releases changing the flow hash table size triggers the FPC to reboot).

To configure inline active flow monitoring on MX Series routers (except for MX80 and MX104 routers), EX Series switches, and T4000 routers with Type 5 FPC:

1. Enable inline active flow monitoring and specify the source address for the traffic.

```
[edit forwarding-options sampling instance instance-name family (bridge | inet | inet6 |
mpls | vpls ) output]
user@host# set inline-jflow source address address
```

2. Specify the template to use with the sampling instance.

```
[edit forwarding-options sampling instance instance-name family (bridge | inet | inet6 |
mpls | vpls ) output flow-server hostname]
user@host# set (version9 | version-ipfix) template template-name
```

3. Configure a template to specify output properties.

```
[edit services flow-monitoring]
user@host# set (version-ipfix | version9) template template-name
```

4. (Optional) Configure the interval after which an active flow is exported.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-active-timeout seconds
```

5. (Optional) Configure the interval of activity that marks a flow as inactive.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-inactive-timeout seconds
```

6. (Optional) Configure the template refresh rate in either number of packets or number of seconds.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set template-refresh-rate (packets packets | seconds seconds)
```

7. (Optional) Configure the refresh rate in either number of packets or number of seconds.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set option-refresh-rate (packets packets | seconds seconds)
```

8. Specify the type of record that the template is used for.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set (bridge-template | ipv4-template | ipv6-template | mpls-ipv4-template | mpls-
template | peer-as-billing-template | vpls-template)
```

The vpls-template option is only for IPFIX templates.

Starting in Junos OS Release 18.2R1, the bridge-template option is available and the vpls-template option is deprecated; use the bridge-template option instead. The bridge-template option (Junos OS only) supports both VPLS and bridge records and is for both IPFIX and version9 templates.

Starting in Junos OS Release 18.4R1, the mpls-ipv4-template option is deprecated for inline flow monitoring. To configure MPLS records starting in Junos OS Release 18.4R1, use the mpls-template option and the tunnel-observation option. This is described in step "9" on page 531.

9. Starting in Junos OS Release 18.4R1 for the MX Series, if you are configuring any type of MPLS flow records, perform the following:

- a. Specify the MPLS template.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set mpls-template
```

- b. Configure the type of MPLS flow records to create.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set tunnel-observation [ipv4 | ipv6]
```

The tunnel-observation values enable the creation of the following types of flow records:

- ipv4—MPLS-IPv4 flows
- ipv6—MPLS-IPv6 flows

You can configure multiple values for tunnel-observation.

For an MPLS traffic type that does *not* match any of the tunnel-observation values, plain MPLS flow records are created. For example, if you only configure `ipv4`, then MPLS-IPv6 traffic results in plain MPLS flow records.

If you do not configure `tunnel-observation`, plain MPLS flow records are created.

- c. If you are running inline flow monitoring on a Lookup (LU) card, enable sideband mode to create MPLS-IPv6 flow records.

```
[edit chassis fpc slot-number inline-services]
user@host# set use-extended-flow-memory
```

If you are running inline flow monitoring on an LU card and do not enable sideband mode, then MPLS-IPv6 traffic results in plain MPLS flow records.

10. (Optional) Include the flow direction value in the template.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-key flow-direction
```

The reported data field contains 0x00 (ingress) or 0x01 (egress). If you do not include the `flow-key flow-direction` statement, the flow direction data field contains the invalid value 0xFF.

11. (Optional) Include VLAN IDs in both the ingress and egress directions in the flow key.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-key vlan-id
```

This statement is not required for ingress and egress VLAN ID reporting on interfaces.

12. Associate the sampling instance with the FPC on which you want to implement inline active flow monitoring.

For MX240, MX480, MX960, MX2010, MX2020, use the following command:

```
[edit ]
user@host# set chassis fpc fpc-number sampling-instance instance-name
```

- a. Confirm the configuration by running the following show command:

```
user@host# show chassis
```

```
fpc 0 {  
    sampling-instance sample-ins1;  
}
```

For MX5, MX10, MX40, and MX80, use the following command:

```
[edit ]  
user@host# set chassis tfeb slot 0 sampling-instance instance-name
```

- a. Confirm the configuration by running the following show command:

```
user@host# show chassis
```

```
tfeb {  
    slot 0 {  
        sampling-instance sample-ins1;  
    }  
}
```

For MX104, use the following command:

```
[edit ]  
user@host# set chassis afeb slot 0 sampling-instance instance-name
```


- a. Confirm the configuration by running the following show command:

```
user@host# show chassis
```

```
afeb {
  slot 0 {
    sampling-instance sample-ins1;
  }
}
```

This example shows the sampling configuration for an instance that supports inline active flow monitoring on family inet:

```
[edit]
user@host> show forwarding-options
sampling {
  instance {
    sample-ins1 {
      input {
        rate 1;
      }
      family inet {
        output {
          flow-server 192.0.2.2 {
            port 2055;
            version-ipfix {
              template {
                ipv4;
              }
            }
          }
        }
        inline-jflow {
          source-address 10.11.12.13;
        }
      }
    }
  }
}
```

Here is the output format configuration:

```
[edit]
user@host> show services flow-monitoring
services {
    flow-monitoring {
        version-ipfix {
            template ipv4 {
                flow-active-timeout 60;
                flow-inactive-timeout 60;
                ipv4-template;
                template-refresh-rate {
                    packets 1000;
                    seconds 10;
                }
                option-refresh-rate {
                    packets 1000;
                    seconds 10;
                }
            }
        }
    }
}
```

The following example shows the output format configuration for chassis fpc0:

```
[edit]
user@host> show services flow-monitoring
sampling-instance instance-1; {
    inline-services {
        flow-table-size {
            ipv4-flow-table-size 8;
            ipv6-flow-table-size 7;
        }
    }
}
```

Release History Table

Release	Description
21.1R1-Evo	Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring.
20.4R1-Evo	Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
20.3R1-Evo	Starting with Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
19.3R2	Inline active flow monitoring is not supported when you enable Next Gen Services on an MX Series router.
18.4R1	Starting in Junos OS Release 18.4R1, the <code>mpls-ipv4-template</code> option is deprecated for inline flow monitoring. To configure MPLS records starting in Junos OS Release 18.4R1, use the <code>mpls-template</code> option and the <code>tunnel-observation</code> option.
18.2R1	Starting in Junos OS Release 18.2R1, the <code>bridge-flow-table-size</code> option is available and the <code>vpls-flow-table-size</code> option is deprecated; use the <code>bridge-flow-table-size</code> option instead.
18.2R1	Starting in Junos OS Release 18.2R1, the <code>bridge-template</code> option is available and the <code>vpls-template</code> option is deprecated; use the <code>bridge-template</code> option instead.
16.1R4	In Junos OS Release 16.2 and in Junos OS Release 16.1R3 and earlier, you can configure only one collector under a family for inline active flow monitoring. Starting with Junos OS Release 16.1R4 and 17.2R1, you can configure up to four collectors under a family for inline active flow monitoring.
16.1R1	Also, starting in Junos OS Release 16.1R1 and 15.1F2, changing the flow hash table size does <i>not</i> automatically reboot the FPC.
16.1R1	Starting with Junos OS Release 16.1R1 and 15.1F2, the IPv4 table is assigned a default value of 1024.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers](#) | 537

Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers

To configure inline active flow monitoring on MX80 and MX104 routers:

1. Associate a sampling instance with the Forwarding Engine Processor.

For the MX80:

```
[edit]
user@host# set chassis tfeb slot 0 sampling-instance sampling-instance
```

The Forwarding Engine Processor slot is always 0 because MX80 and MX104 routers have only one Packet Forwarding Engine. In this MX80 configuration, the sampling instance is `sample-ins1`.

```
[edit]
user@host# set chassis tfeb slot 0 sampling-instance sample-ins1
```

For the MX104:

```
[edit]
user@host# set chassis afeb slot 0 sampling-instance sampling-instance
```

NOTE: MX80 and MX104 routers support only one sampling instance.

2. Under forwarding-options, configure a sampling instance for the flow server and inline jflow instances (these be configured in the following steps):

```
[edit forwarding-options sampling]
user@host# edit instance inline_sample
```

3. Configure the rate at the [edit forwarding-options sampling instance instance-name input] hierarchy level to apply specific values for the sampling instance sample-ins1.

```
[edit forwarding-options sampling instance sample-ins1 input]
user@host# set rate number
```

In this configuration, the rate is 1000.

```
[edit forwarding-options sampling instance sample-ins1 input]
user@host# set rate 1000
```

4. Navigate to the output hierarchy and from there, enable a flow server and then specify the output address and port:

```
[edit] forwarding-options sampling instance inline_sample family inet output]
user@host# edit flow-server hostname
```

```
[edit forwarding-options sampling instance inline_sample family inet output flow-server
hostname]
user@host# set port number
```

5. Return to the output hierarchy and specify the source address for inline jflow:

```
[edit forwarding-options sampling instance sample-ins1 family inet output]
user@host# set inline-jflow source-address address
```

In this configuration, the source address is 10.11.12.13.

```
[edit forwarding-options sampling instance sample-ins1 family inet output]
user@host# set inline-jflow source-address 10.11.12.13
```

6. Specify the output properties.

```
[edit services flow-monitoring]
user@host# set version-ipfix
```

The output format properties are common to other output formats and are described in [“Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices” on page 603](#).

The following is an example of the sampling configuration for an instance that supports inline active flow monitoring on MX80 routers:

```
[edit forwarding-options]
user@host# show
sampling {
  instance {
    inline_sample {
      input {
        rate 1000;
      }
      family inet{
        output {
          flow-server 192.168.64.143 {
            port 80;
          }
          inline-jflow {
            source-address 10.10.11.12;
          }
        }
      }
    }
  }
}
```

NOTE: You need not configure a Flexible PIC Concentrator (FPC) slot because MX80 routers have only one Packet Forwarding Engine.

The following considerations apply to the inline flow-monitoring instance configuration:

- This configuration does not support MPLS-IPv6.
- Clip-size is not supported.

RELATED DOCUMENTATION

Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices 603
Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 74
inline-jflow 1162

Configuring Inline Active Flow Monitoring on PTX Series Routers

IN THIS SECTION

- [Platform and Feature Support | 540](#)
- [How to Configure Inline Active Flow Monitoring on PTX Series Routers | 543](#)

This topic describes how to configure inline active flow monitoring on PTX Series routers for IPv4 and IPv6 traffic.

Platform and Feature Support

[Table 105 on page 540](#) lists the PTX Series platform support for various types of traffic for inline active flow monitoring.

Table 105: PTX Series Platform Support for Inline Active Flow Monitoring

Platform	Support
PTX3000 Series	Junos OS 18.1R1—IPv4 and IPv6 traffic (both IPFIX and version 9) Junos OS 18.2R1—MPLS, MPLS-IPv4, and MPLS-IPv6 traffic.
PTX5000 Series	Junos OS 18.1R1—IPv4 and IPv6 traffic (both IPFIX and version 9) Junos OS 18.2R1, MPLS, MPLS-IPv4, and MPLS-IPv6 traffic.

Table 105: PTX Series Platform Support for Inline Active Flow Monitoring (Continued)

Platform	Support
PTX1000	Junos OS 17.3R1—IPv4 and IPv6 traffic (version 9 only).
PTX10001-36MR	Junos OS Evolved 20.3R1—IPv4, IPv6, MPLS, MPLS-IPv4, and MPLS-IPv6 traffic.
PTX10002-60C	Junos OS 18.4R1—IPv4 and IPv6 traffic (both IPFIX and version 9). Junos OS 19.4R1—MPLS, MPLS-IPv4, and MPLS-IPv6 traffic.
PTX10003	Junos OS Evolved 19.3R1—IPv4 and IPv6 traffic (IPFIX and version 9). Junos OS Evolved 20.1R1—MPLS, MPLS-IPv4, and MPLS-IPv6 traffic.
PTX10004	Junos OS Evolved 20.4R1—IPv4, IPv6, MPLS, MPLS-IPv4, and MPLS-IPv6 traffic (IPFIX and version 9).
PTX10008 (with the JNP10008-SF3 and the JNP10K-LC1201 line card)	Junos OS Evolved 19.3R1—IPv4 and IPv6 traffic (IPFIX and version 9). Junos OS Evolved 20.1R1—MPLS, MPLS-IPv4, and MPLS-IPv6 traffic.
PTX10008 (with the JNP10008-SF3 and the JNP10K-LC1202 line card)	Junos OS Evolved 20.3R1—IPv4, IPv6, MPLS, MPLS-IPv4, and MPLS-IPv6 traffic (IPFIX and version 9).
PTX10008 (without the JNP10008-SF3) and PTX10016	Junos OS 18.1R1—IPv4 and IPv6 traffic (both IPFIX and version 9) Junos OS 18.2R1—MPLS, MPLS-IPv4, and MPLS-IPv6 traffic.

To configure inline flow monitoring for MPLS-over UDP traffic on PTX Series Routers, see ["Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers" on page 550](#). Inline active flow monitoring for MPLS-over-UDP traffic is not supported on the PTX10001-36MR, PTX10003, PTX10004, and the PTX10008 (with the JNP10008-SF3) routers.

Starting in Junos OS Release 18.2R1, you can configure up to four collectors under a family for inline active flow monitoring. In previous releases of Junos OS, you could configure only one collector under a family for inline active flow monitoring. Starting in Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring. Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring. To configure a collector under a family for inline active flow monitoring, configure the `flow-server` statement at the `edit forwarding-options sampling-instance instance-name family (inet | inet6) output` hierarchy level. To specify up to four collectors, include up to four `flow-server` statements.

Inline active flow monitoring is implemented on the Logical CPU (LCPU). All the functions like flow creation, flow update, and flow records export are done by the LCPU. The flow records are sent out in either the IPFIX format or the version 9 format.

Starting with Junos OS Evolved Release 21.2R1 and Junos OS Release 21.3R1, no flows are maintained. Every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, the number of records sent to the collector is higher than before. The IPFIX and version 9 Options Template Data Record now contains 0 in the Flow Active Timeout (Element ID 36) and Flow Inactive Timeout (Element ID 37) fields. Therefore, the Options Template Data Record is not compliant with IPFIX RFC 7011. The `show services accounting flow inline-jflow fpc-slot slot operational mode` command now displays 0 for all of the Active Flows and Timed Out fields. The values of the various Total Flows fields are now equal to their respective Flow Packets field values. The values of the various Flows Inactive Timed Out fields are now equal to their respective Flow Packets field values. The effect of the `nexthop-learning` statement at the `[edit services flow-monitoring version version template template-name]` hierarchy level on this no-flow behavior varies depending upon the operating system. For Junos OS Evolved, we do not recommend that you configure the `nexthop-learning` statement, as it reduces the number of packets that can be processed. For Junos OS, you can configure the `nexthop-learning` statement to change this default no-flow behavior and once again create and maintain flows, then attach the template to all sampling instances associated with FPCs that require the previous behavior.

The following limitations and restrictions apply to the inline active flow monitoring feature in Junos OS and Junos OS Evolved:

- Egress MPLS filters are not supported on the PTX10001-36MR, PTX10003, PTX10004, and the PTX10008 (with the JNP10008-SF3) routers.
- The PTX10001-36MR router does not support multiple FPC sampling collection because it has only 1 Routing Engine.
- True outgoing interface (OIF) reporting is not supported for egress sampling. In Junos OS Evolved, true outgoing interface (OIF) reporting is not supported for GRE de-encapsulated packets.

- The interface type field for the true incoming interface is not part of the version 9 template because this element is not present in the version 9 export version.
- For GRE tunnel traffic on PTX10003 routers, the physical interface is reported in the layer 2 header and is considered as one of the keys during flow creation. Therefore, when physical interfaces are moved in or out of the aggregated Ethernet bundle, a new flow is created and the old flows are timed out after a period of inactivity. Physical interface, logical interface, or the aggregated logical interface (based on the configuration) is reported as the incoming interface in export records based on the configuration.

For GRE tunnel traffic on PTX10008 (with the JNP10008-SF3) routers, an FTI interface is configured to terminate a GRE tunnel. This interface is used during flow creation as one of the keys instead of the physical interface. Hence when a physical interface is moved in or out of an aggregated Ethernet bundle, no new flow is created as the key remains unchanged. Physical interface, logical interface, or the aggregated logical interface (based on the configuration) is reported as the incoming interface in exported records.

How to Configure Inline Active Flow Monitoring on PTX Series Routers

SUMMARY

In this example, we configure a version-ipfix template for recording IPv4 and IPv6 traffic flows.

IN THIS SECTION

- [Configure a Template to Specify Output Properties | 543](#)
- [Configure a Sampling Instance to Specify Input Properties | 545](#)
- [Assign the Sampling Instance to an FPC | 546](#)
- [Configure a Firewall Filter to Accept and Sample Flows | 546](#)
- [Assign the Firewall Filter to an Interface | 546](#)
- [Results from a Sample Configuration | 546](#)

Configure a Template to Specify Output Properties

1. Define the template and configure the type of flow the template should record.

```
[edit services flow-monitoring]
user@host# set version-ipfix template template-name ipv4-template
```

```

user@host# set version-ipfix template template-name ipv6-template
user@host# set version-ipfix template template-name mpls-template

```

2. (Optional) Configure additional output properties for the template, such as flow timeout interval and template/option refresh rates, to control the flow records.

You can use the `template-refresh-rate` option to configure the frequency at which the flow generator sends updates about template definitions to the flow collector either using number of packets or seconds.

```

[edit services flow-monitoring]
user@host# set version-ipfix template template-name flow-active-timeout seconds
user@host# set version-ipfix template template-name flow-inactive-timeout seconds
user@host# set version-ipfix template template-name template-refresh-rate (packets packets |
seconds seconds)
user@host# set version-ipfix template template-name option-refresh-rate (packets packets |
seconds seconds)

```

3. (Optional)

If you are monitoring MPLS flows, that is, if the template in use is configured for the MPLS protocol family, use the `tunnel-observation` option to identify the types of MPLS flows.

```

[edit services flow-monitoring]
user@host# set version-ipfix template template-name tunnel-observation (ipv4 | ipv6 | mpls-over-
udp)

```

4. (Optional) Enable the learning of next-hop addresses so that the true outgoing interface is reported.

NOTE: Starting in Junos OS Evolved 21.2R1, we do not recommend that you enable learning of next-hop addresses, as it reduces the number of packets that can be processed. However, starting in Junos OS Release 21.3R1, you can configure the `nexthop-learning` statement to change the default no-flow behavior and once again create and maintain flows, then attach the template to all sampling instances associated with FPCs that require the previous behavior.

```

[edit services flow-monitoring]
user@host# set version-ipfix template template-name nexthop-learning enable

```

Configure a Sampling Instance to Specify Input Properties

1. Define the sampling instance and configure the ratio of number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

```
[edit forwarding-options sampling]
user@host# set instance instance-name input rate number
```

BEST PRACTICE: We recommend that you use a value of 1000 or higher for MPLS flows.

2. Configure the protocol family for the sampling instance and specify a flow collector to send the traffic aggregates.

```
[edit forwarding-options sampling]
user@host# set instance instance-name family (inet | inet6 | mpls) flow-server hostname
```

3. (Optional) Specify the UDP port for the flow collector and the template to use with the sampling instance.

```
[edit forwarding-options sampling]
user@host# set instance instance-name family (inet | inet6 | mpls) flow-server hostname port port-number
user@host# set instance instance-name family (inet | inet6 | mpls) flow-server hostname version-ipfix template template-name
```

4. Configure inline processing of the sampled packets.

```
[edit forwarding-options sampling]
user@host# set instance instance-name family (inet | inet6 | mpls) output inline-flow source-address address
```

Assign the Sampling Instance to an FPC

1. Assign the sampling instance to the FPC on which you want to implement flow monitoring.

```
[edit chassis]
user@host# set fpc slot-number sampling-instance instance-name
```

Configure a Firewall Filter to Accept and Sample Flows

1. Configure the firewall filter for the protocol family and enable sampling of traffic flows.

```
[edit firewall]
user@host# set family (inet | inet6 | mpls) filter filter-name
user@host# set family (inet | inet6 | mpls) filter filter-name term term-name then accept
user@host# set family (inet | inet6 | mpls) filter filter-name term term-name then sample
```

Assign the Firewall Filter to an Interface

1. Assign the input firewall filter to the interface you want to monitor.

```
[edit interfaces]
user@host# set interface-name unit unit-number family (inet | inet6 | mpls) filter input
filter-name
```

Results from a Sample Configuration

The following is an example of the sampling configuration for an instance that supports inline flow monitoring on family inet and on family inet6:

```
[edit chassis]
fpc 0 {
    sampling-instance sample-1;
}
```

```
[edit services]
```

```

flow-monitoring {
    version-ipfix {
        template test-template {
            flow-active-timeout 30;
            flow-inactive-timeout 60;
            nexthop-learning {
                enable;
            }
            template-refresh-rate {
                seconds 10;
            }
            ipv4-template;
        }
        template v6 {
            ipv6-template;
        }
    }
}

```

```

[edit interfaces]
et-1/0/0 {
    unit 0 {
        family inet {
            filter {
                input ipv4-filter;
            }
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.100.1/32;
        }
    }
}

```

```

[edit forwarding-options]

```

```
sampling {
  instance {
    ipv4 {
      input {
        rate 10;
      }
      family inet {
        output {
          flow-server 10.208.174.127 {
            port 2055;
            version-ipfix {
              template {
                test-template;
              }
            }
          }
          inline-jflow {
            source-address 192.168.100.1;
          }
        }
      }
    }
    family inet6 {
      output {
        flow-server 10.208.174.127 {
          port 2055;
          version-ipfix {
            template {
              v6;
            }
          }
          inline-jflow {
            source-address 192.168.100.1;
          }
        }
      }
    }
  }
}
```

```

    }
}

```

```

[edit services]
flow-monitoring {
  version-ipfix {
    template test-template {
      flow-active-timeout 30;
      flow-inactive-timeout 60;
      nexthop-learning {
        enable;
      }
      template-refresh-rate {
        seconds 10;
      }
      ipv4-template;
    }
    template v6 {
      ipv6-template;
    }
  }
}

```

You can use the ["show services accounting flow" on page 1645](#) command to verify active flow statistics.

Release History Table

Release	Description
21.3R1	
21.3R1	For the PTX Series, starting with Junos OS Release 21.3R1 , no flows are maintained. Every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, the number of records sent to the collector is higher than before.
21.2R1-Evo	For the PTX Series, starting with Junos OS Evolved Release 21.2R1, no flows are maintained. Every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, the number of records sent to the collector is higher than before.

21.1R1-Evo	Starting with Junos OS Evolved 21.1R1, for the PTX10004 router, you can configure up to four collectors for inline active flow monitoring.
20.4R1-Evo	Starting with Junos OS Evolved 20.4R1, for the PTX10001-36MR and the PTX10008 (with the JNP10K-LC1202 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
20.3R1-Evo	Starting in Junos OS Evolved 20.3R1, for the PTX10003 and PTX10008 (with the JNP10K-LC1201 line card and the JNP10008-SF3) routers, you can configure up to four collectors for inline active flow monitoring.
18.2R1	Starting in Junos OS Release 18.2R1, you can configure up to four collectors under a family for inline active flow monitoring. In previous releases of Junos OS, you could configure only one collector under a family for inline active flow monitoring.

Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers

IN THIS SECTION

- [MPLS-over-UDP Flow Monitoring Overview | 550](#)
- [Configuring Inline Active Flow Monitoring of MPLS-over-UDP Flows | 553](#)

You can enable inline active flow monitoring that reports the inner payload of MPLS-over-UDP flows on PTX Series routers and QFX10002-60C switches.

MPLS-over-UDP Flow Monitoring Overview

IN THIS SECTION

- [Benefits of Using MPLS-Over-UDP Flow Monitoring | 551](#)
- [Flow Monitoring Scenarios for MPLS-over-UDP | 551](#)

Starting with Junos OS Release 18.1R1 on PTX Series routers with an FPC3, PTX10K-LC1101, PTX10K-LC1102, or PTX1000 card, you can perform inline active flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel.

Starting with Junos OS Release 19.4R1, on the PTX10002-60C, you can perform inline active flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel.

Starting with Junos OS Release 19.4R1, the PTX10002-60C supports inline active flow monitoring for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP traffic. Both IPFIX and version 9 templates are supported.

Starting with Junos OS Release 21.2R1, the QFX10002-60C supports inline active flow monitoring for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP traffic. Both IPFIX and version 9 templates are supported.

For a description of the fields included in the templates, see ["Understand Inline Active Flow Monitoring" on page 437](#). Only ingress sampling is supported.

MPLS-over-UDP is not supported on the PTX10001-36MR, PTX10003, PTX10004, and PTX10008 (with the JNP10008-SF3) routers.

Benefits of Using MPLS-Over-UDP Flow Monitoring

- Gather and export detailed information on even the original IPv4 or IPv6 payload of the MPLS-over-UDP flow.

Flow Monitoring Scenarios for MPLS-over-UDP

Monitoring for MPLS-over-UDP tunnels includes the following scenarios:

- The MPLS-over-UDP flow is carried through a full IP network, using IPv4 endpoints on PTX Series routers (see [Figure 53 on page 552](#)). The inner payload may be IPv4 or IPv6. [Figure 54 on page 552](#) shows the encapsulated packet. Flow monitoring reports the inner IP header and payload, in addition to the tunnel and MPLS fields.

You can enable ingress monitoring for the MPLS-over-UDP tunnel at its transit and egress nodes. For example, in [Figure 53 on page 552](#), you can enable ingress monitoring on routers R4, R5, R6, and R7.

Figure 53: MPLS-over-UDP in Full IP Network




Figure 54: Encapsulated Packet for MPLS-over-UDP in Full IP Network



- The MPLS-over-UDP flow is carried through an IP-MPLS-IP network, using IPv4 endpoints on PTX Series routers (see [Figure 55 on page 553](#)). The inner payload may be IPv4 or IPv6. In the inner MPLS network, the MPLS-over-UDP flow is encapsulated in an RSVP-TE label-switched path (LSP). [Figure 56 on page 553](#) shows the encapsulated packet. Flow monitoring reports the inner IP header and payload, in addition to the RSVP label, tunnel, and MPLS fields.

You can enable ingress monitoring for the MPLS-over-UDP tunnel at its transit and egress nodes. For example, in [Figure 55 on page 553](#), you can enable ingress monitoring on routers R4, R5, R6, R7, R8, and R9.

Figure 55: MPLS-over-UDP Over IP-MPLS-IP Network




Figure 56: MPLS-over-UDP in RSVP-TE LSP Packet



Configuring Inline Active Flow Monitoring of MPLS-over-UDP Flows

IN THIS SECTION

- [Configuring the Template to Specify Output Properties | 554](#)
- [Configuring the Sampling Instance | 555](#)
- [Assigning the Sampling Instance to an FPC | 557](#)
- [Configuring a Firewall Filter | 557](#)
- [Assigning the Firewall Filter to the Monitored Interface | 557](#)

(Junos OS only) Configuring inline active monitoring of MPLS-over-UDP flows includes the following tasks:

Configuring the Template to Specify Output Properties

Configure a template to specify the output properties for the flow records:

1. Configure the template name.

```
[edit services flow-monitoring]
user@host# set (version-ipfix | version9) template template-name
```

2. (Optional) Configure the interval after which an active flow is exported.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-active-timeout seconds
```

3. (Optional) Configure the interval of activity that marks a flow as inactive.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-inactive-timeout seconds
```

4. (Optional) Configure the frequency at which the flow generator sends updates about template definitions to the flow collector. Specify either number of packets or number of seconds.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set template-refresh-rate packets packets seconds seconds
```

5. (Optional) Configure the refresh rate in either number of packets or number of seconds.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set option-refresh-rate packets packets seconds seconds
```

6. Enable flow monitoring of MPLS-over-UDP flows.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set tunnel-observation mpls-over-udp
```

7. Specify the template type.

- If you are monitoring an MPLS-over-UDP flow that is carried through a full IP network (see [Figure 53 on page 552](#)), use the `ipv4-template`:

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set ipv4-template
```

- If you are monitoring an MPLS-over-UDP flow that is carried through an IP-MPLS-IP network (see [Figure 55 on page 553](#)):

For the IP network transit and egress nodes (for example, R4, R5, R8, and R9 in [Figure 55 on page 553](#)), use the `ipv4-template` type.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set ipv4-template
```

For the transit and egress nodes where the MPLS-over-UDP flow is encapsulated in an RSVP-TE LSP (for example R6 and R7 in [Figure 55 on page 553](#)), use one of the following templates:

- Starting in Junos OS Release 18.2R1:

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set mpls-template
```

- In Junos OS Release 18.1:

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set mpls-ipvx-template
```

8. Enable the learning of next-hop addresses so that the true outgoing interface (OIF) is reported.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set nexthop-learning
```

Configuring the Sampling Instance

Configure a sampling instance:

1. Configure the sampling instance name.

```
[edit forwarding-options sampling]
user@host# set instance instance-name
```

2. Configure the MPLS protocol family for the sampling instance.

```
[edit forwarding-options sampling instance instance-name]
user@host# set family mpls
```

3. Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

```
[edit forwarding-options sampling instance instance-name input]
user@host# set rate number
```

4. Specify the source address for the traffic to be sampled.

```
[edit forwarding-options sampling instance instance-name family mpls output]
user@host# set inline-jflow source-address address
```

5. Specify the flow export rate of monitored packets in kpps.

```
[edit forwarding-options sampling instance instance-name family mpls output]
user@host# set inline-jflow flow-export-rate rate
```

6. Specify the output address and port for a flow server.

```
[edit forwarding-options sampling instance instance-name family mpls output]
user@host# set flow-server hostname port port-number
```

7. Specify the template to use with the sampling instance.

```
[edit forwarding-options sampling instance instance-name family mpls output flow-server
hostname]
user@host# set (version9 | version-ipfix) template template-name
```

Assigning the Sampling Instance to an FPC

- Assign the sampling instance to the FPC on which you want to implement flow monitoring.

```
[edit chassis]
user@host# set fpc slot-number sampling-instance instance-name
```

Configuring a Firewall Filter

Configure a firewall filter to accept and sample MPLS traffic.

1. Configure the MPLS firewall filter name.

```
[edit firewall]
user@host# edit family mpls filter filter-name
```

2. Configure a term to sample and accept MPLS packets.

```
[edit firewall family mpls filter filter-name]
user@host# set term term-name then accept
user@host# set term term-name then sample
```

Assigning the Firewall Filter to the Monitored Interface

- Assign the input firewall filter to the interface you want to monitor.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family mpls filter input filter-name
```

Release History Table

Release	Description
21.2R1	Starting with Junos OS Release 21.2R1, the QFX10002-60C supports inline active flow monitoring for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP traffic. Both IPFIX and version 9 templates are supported.

19.4R1	Starting with Junos OS Release 19.4R1, on the PTX10002-60C, you can perform inline flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel.
19.4R1	Starting with Junos OS Release 19.4R1, the PTX10002-60C supports inline active flow monitoring for MPLS, MPLS-IPv4, MPLS-IPv6, and MPLS-over-UDP traffic. Both IPFIX and version 9 templates are supported.
18.1R1	Starting with Junos OS Release 18.1R1 on PTX Series routers with an FPC3, PTX10K-LC1101, PTX10K-LC1102, or PTX1000 card, you can perform inline flow monitoring for MPLS-over-UDP flows to look past the tunnel header to sample and report on the inner payload at both the transit and egress nodes of the tunnel.

Inline Active Flow Monitoring on IRB Interfaces

IN THIS SECTION

- [Overview | 558](#)
- [Understand Inline Active Flow Monitoring on IRB interfaces | 559](#)
- [Configure Inline Active Flow Monitoring on IRB Interfaces on PTX Series Routers | 561](#)

You can perform inline active flow monitoring for IPv4 and IPv6 traffic on the integrated routing and bridging (IRB) interfaces on PTX Series routers.

Overview

On PTX Series routers, you can perform inline active flow monitoring for IPv4 and IPv6 traffic on the integrated routing and bridging (IRB) interfaces. Both IPFIX and version 9 templates for the flow monitoring are supported. For a description of the fields included in the templates, see "[Understand Inline Active Flow Monitoring](#)" on page 437.

Understand Inline Active Flow Monitoring on IRB interfaces

IN THIS SECTION


- Sampling on an IRB Interface with Traffic Routed to a Tunnelled Core | 559
- Layer 2 bridging and Layer 3 IP routing on an IRB interface | 560

You can enable inline active flow monitoring by configuring the IPFIX or V9 templates on IRB interfaces.

Sampling on an IRB Interface with Traffic Routed to a Tunnelled Core

Figure 57 on page 559 illustrates sampling on an IRB interface where the traffic is routed to a tunnelled core, primarily an MPLS tunnel. The packets are entering irb.10 on which you can enable ingress sampling. The packets can be forwarded to a next hop which is not a part of any user-defined VLAN.


Figure 57: Sampling on an IRB Interface Routing Traffic to a Tunnelled Core



Layer 2 bridging and Layer 3 IP routing on an IRB interface

Figure 58 on page 560 illustrates the topology where Layer 2 bridging and Layer 3 IP routing are supported on the same interface.

Figure 58: Layer 2 Bridging and Layer 3 IP Routing on the Same IRB Interface




PC1 and PC2 are in VLAN RED (ID 10) and PC3 is in VLAN BLUE (ID 20).

For traffic moving from PC1 to PC3 or from PC2 to PC3, an IRB interface must be configured with a logical unit with an address in the subnet for VLAN RED and a logical unit with an address in the subnet for VLAN BLUE. The switch automatically directs routes to these subnets and uses these routes to forward traffic between VLANs. If traffic is flowing from VLAN RED to VLAN BLUE, you can configure ingress sampling on irb.10 and egress sampling on irb.20.

Figure 59 on page 560 illustrates sampling in a topology where Layer 2 bridging and Layer 3 IP routing are supported on the same interface. The interfaces, et-0/0/36.0 and irb.10 belong to VLAN ID 2. The interfaces, et-0/0/48 and irb.20 belong to VLAN ID 3. Packets are entering irb.10 and exiting on irb.20. Hence, you can configure ingress sampling on irb.10 and egress sampling on irb.20.

Figure 59: Sampling on an IRB Interface Supporting Bridging and Routing



Configure Inline Active Flow Monitoring on IRB Interfaces on PTX Series Routers

IN THIS SECTION

- [Configure the Template to Specify Output Properties | 561](#)
- [Configure the Sampling Instance | 562](#)
- [Assign the Sampling Instance to an FPC | 564](#)
- [Configure a Firewall Filter | 564](#)
- [Associate a Layer 3 Interface with the VLAN to Route Traffic | 565](#)
- [Assign the Firewall Filter to the Monitored Interface | 566](#)

Configure the Template to Specify Output Properties

Configure a template to specify the output properties for the flow records:

1. Configure the template name.

```
[edit services flow-monitoring]
user@host# set (version-ipfix | version9) template template-name
```

For example:

```
[edit services flow-monitoring]
user@host# set version-ipfix template t1
```

2. (Optional) Configure the interval after which an active flow is exported.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-active-timeout seconds
```

For example:

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-active-timeout 10
```

3. (Optional) Configure the interval of activity that marks a flow as inactive.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-inactive-timeout seconds
```

For example:

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set flow-inactive-timeout 10
```

4. Specify the template type.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set template-name
```

For example:

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
user@host# set ipv4-template
```

Configure the Sampling Instance

Configure a sampling instance:

1. Configure the sampling instance name.

```
[edit forwarding-options sampling]
user@host# set instance instance-name
```

For example:

```
[edit forwarding-options sampling]
user@host# set instance s1
```

2. Configure the protocol family for the sampling instance.

```
[edit forwarding-options sampling instance instance-name]  
user@host# set family (inet | inet6 | mpls)
```

For example:

```
[edit forwarding-options sampling instance instance-name]  
user@host# set family inet
```

3. Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

```
[edit forwarding-options sampling instance instance-name input]  
user@host# set rate number
```

For example:

```
[edit forwarding-options sampling instance instance-name input]  
user@host# set rate 10
```

4. Specify the source address for the traffic to be sampled.

```
[edit forwarding-options sampling instance instance-name family inet output]  
user@host# set inline-jflow source-address address
```

For example:

```
[edit forwarding-options sampling instance instance-name family inet output]  
user@host# set inline-jflow source-address 10.10.0.1
```

5. Specify the output address and port for a flow server.

```
[edit forwarding-options sampling instance instance-name family inet output]  
user@host# set flow-server hostname port port-number
```

For example:

```
[edit forwarding-options sampling instance instance-name family inet output]
user@host# set flow-server 10.10.10.2 port 2055
```

6. Specify the template to use with the sampling instance.

```
[edit forwarding-options sampling instance instance-name family inet output flow-server
hostname]
user@host# set (version9 | version-ipfix) template template-name
```

For example:

```
[edit forwarding-options sampling instance instance-name family inet output]
user@host# set version-ipfix template t1
```

Assign the Sampling Instance to an FPC

Assign the sampling instance to the FPC on which you want to implement flow monitoring.

```
[edit chassis]
user@host# set fpc slot-number sampling-instance instance-name
```

For example:

```
[edit chassis]
user@host# set fpc 0 sampling-instance s1
```

Configure a Firewall Filter

Configure a firewall filter to specify the family of traffic to accept and sample.

1. Configure the firewall filter name and specify the family of traffic.

```
[edit firewall]
user@host# set family (inet | inet6 | mpls) filter filter-name
```

For example:

```
[edit firewall]
user@host# set family inet filter f2
```

2. Configure a term to sample and accept packets.

```
[edit firewall family mpls filter filter-name]
user@host# set term term-name then accept
user@host# set term term-name then sample
```

For example:

```
[edit firewall family mpls filter filter-name]
user@host# set term t1 then count c2
user@host# set term t1 then accept
user@host# set term t1 then sample
```

Associate a Layer 3 Interface with the VLAN to Route Traffic

Assign the IRB Interface to the VLAN.

```
[edit vlans vlan-name]
user@host# set vlan-name vlan-id vlan-id-number
user@host# set vlan-name l3-interface l3-interface-name .logical-interface-number
```

For example:

```
[edit vlans vlan-name]
user@host# set vlan2 vlan-id 2
user@host# set vlan2 l3-interface irb.10
```

For example, if you are configuring inline flow monitoring using IRB while supporting layer 2 bridging and layer 3 IP routing on the same interface (See [Figure 59 on page 560](#)):

```
[edit vlans vlan-name]
user@host# set vlan-2 vlan-id 2
user@host# set vlan-2 l3-interface irb.10
```



```
user@host# set vlan-3 vlan-id 3
user@host# set vlan-3 l3-interface irb.20
```

Assign the Firewall Filter to the Monitored Interface

Assign the input firewall filter to the interface you want to monitor. Also, configure the VLANs for which the interface can carry traffic.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family (inet | inet6 | mpls) filter input filter-name address
```

For example, if you are configuring inline flow monitoring using IRB while supporting layer 2 bridging and layer 3 IP routing on the same interface (See [Figure 59 on page 560](#)):

```
[edit interfaces]
user@host# set et-0/0/36 unit 0 family ethernet-switching vlan members vlan2
user@host# set et-0/0/48 unit 0 family ethernet-switching vlan members vlan3
user@host# set et-0/0/60 unit 0 family inet address 10.10.10.1
user@host# set irb unit 1 family inet filter input f2
user@host# set irb unit 1 family inet address 10.1.1.1
user@host# set irb unit 2 family inet address 10.20.1.1
user@host# set irb unit 1 family inet address 10.1.1.1
user@host# set irb unit 2 family inet filter output f2
```

Release History Table

Release	Description
22.2R1-EVO	Starting in Junos OS Evolved Release 22.2R1 on the PTX10003 router, you can perform inline active flow monitoring for IPv4 and IPv6 traffic on integrated routing and bridging (IRB) interfaces.
21.3R1-EVO	Starting in Junos OS Evolved Release 21.3R1 on the PTX10001-36MR, PTX10004, and PTX10008 routers, you can perform inline active flow monitoring for IPv4 and IPv6 traffic on integrated routing and bridging (IRB) interfaces.
19.1R1	Starting in Junos OS Release 19.1R1, on PTX Series routers, you can perform inline active flow monitoring for IPv4 and IPv6 traffic on integrated routing and bridging (IRB) interfaces.

Example: Configuring Inline Active Flow Monitoring on MX Series and T4000 Routers

IN THIS SECTION

- [Software and Hardware Requirements | 574](#)
- [Overview | 575](#)

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.

Configuring Template Properties

```
set services flow-monitoring version9 template template1 flow-active-timeout 120
set services flow-monitoring version9 template template1 flow-inactive-timeout 60
set services flow-monitoring version9 template template1 template-refresh-rate packets 100
set services flow-monitoring version9 template template1 template-refresh-rate seconds 600
set services flow-monitoring version9 template template1 option-refresh-rate packets 100
set services flow-monitoring version9 template template1 option-refresh-rate seconds 600
set services flow-monitoring version9 template template1 ipv4-template
set services flow-monitoring version-ipfix template template-v61 flow-active-timeout 150
set services flow-monitoring version-ipfix template template-v61 flow-inactive-timeout 100
set services flow-monitoring version-ipfix template template-v61 template-refresh-rate seconds 30
set services flow-monitoring version-ipfix template template-v61 ipv6-template
```

Configuring a Sampling Instance

```
set forwarding-options sampling instance instance-1 input rate 1
set forwarding-options sampling instance instance-1 family inet output flow-server 10.50.1.2
port 2055
set forwarding-options sampling instance instance-1 family inet output flow-server 10.50.1.2
version9 template template1
```

```

set forwarding-options sampling instance instance-1 family inet output inline-jflow source-
address 10.50.1.100
set forwarding-options sampling instance instance-1 family inet output inline-jflow flow-export-
rate 10
set forwarding-options sampling instance instance-1 family inet6 output flow-server 10.50.1.2
port 2055
set forwarding-options sampling instance instance-1 family inet6 output flow-server 10.50.1.2
version-ipfix template template-v61
set forwarding-options sampling instance instance-1 family inet6 output inline-jflow source-
address 10.50.1.110
set forwarding-options sampling instance instance-1 family inet6 output inline-jflow flow-export-
rate 6

```

Configuring FPC Parameters

```

set chassis fpc 0 sampling-instance instance-1
set chassis fpc 0 inline-services flow-table-size ipv4-flow-table-size 8
set chassis fpc 0 inline-services flow-table-size ipv6-flow-table-size 7

```

Configuring Firewall Filters

```

set firewall family inet filter inet-sample term t1 then sample
set firewall family inet filter inet-sample term t1 then accept
set firewall family inet6 filter inet6-sample term t1 then sample
set firewall family inet6 filter inet6-sample term t1 then accept

```

Configuring Interface Properties

```

set interfaces ge-0/0/4 unit 0 family inet filter input inet-sample
set interfaces ge-0/0/4 unit 0 family inet address 10.150.1.1/24
set interfaces ge-0/1/6 unit 0 family inet6 filter input inet6-sample
set interfaces ge-0/1/6 unit 0 family inet6 address 2001:db8:0:2::1/64

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

1. Configure the template properties for inline active flow monitoring.

```
[edit services flow-monitoring]
user@router1# set version9 template template1 ipv4-template
user@router1# set version9 template template1 flow-active-timeout 120
user@router1# set version9 template template1 flow-inactive-timeout 60
user@router1# set version9 template template1 template-refresh-rate packets 100
user@router1# set version9 template template1 option-refresh-rate packets 100
user@router1# set version-ipfix template template-v61 ipv6-template
user@router1# set version-ipfix template template-v61 flow-active-timeout 150
user@router1# set version-ipfix template template-v61 flow-inactive-timeout 100
user@router1# set version-ipfix template template-v61 template-refresh-rate seconds 30
user@router1# set version-ipfix template template-v61 option-refresh-rate seconds 30
```

2. Configure the sampling instance for inline active flow monitoring.

```
[edit forwarding-options sampling]
user@router1# set instance instance-1 input rate 1
user@router1# set instance instance-1 family inet output flow-server 10.50.1.2 port 2055
user@router1# set instance instance-1 family inet output flow-server 10.50.1.2 version9
template template1
user@router1# set instance instance-1 family inet output inline-jflow source-address
10.50.1.100
user@router1# set instance instance-1 family inet output inline-jflow flow-export-rate 10
user@router1# set instance instance-1 family inet6 output flow-server 10.50.1.2 port 2055
user@router1# set instance instance-1 family inet6 output flow-server 10.50.1.2 version-ipfix
template template-v61
user@router1# set instance instance-1 family inet6 output inline-jflow source-address
10.50.1.110
user@router1# set instance instance-1 family inet6 output inline-jflow flow-export-rate 6
```

NOTE: Until you complete the next step for associating the sampling instance with an FPC, the instance remains inactive and is marked `inactive` in the configuration.

3. Associate the sampling instance with the FPC on which you want to implement inline active flow monitoring, and also configure the hash table sizes.

NOTE: In Junos OS releases earlier than Release 12.1, the following conditions are applicable for supporting backward compatibility when you configure the IPv4 and IPv6 flow table sizes for inline active flow monitoring:

- If you do not configure the `flow-table-size` statement at the `[edit chassis fpc slot-number inline-services]` hierarchy level, fifteen 256K entries are allocated by default for the IPv4 flow table and one 1K entry is allocated by default for the IPv6 flow table on the Packet Forwarding Engine.
- If you configure the `ipv4-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level and do not configure the `ipv6-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level, the number of units of 256K entries that you configure for the IPv4 flow table is allocated. For the IPv6 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you do not configure the `ipv4-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level and if you configure the `ipv6-flow-table-size size` statement at the `[edit chassis fpc slot-number inline-services flow-table-size]` hierarchy level, the number of units of 256K entries that you configure for the IPv6 flow table is allocated. For the IPv4 flow table, a default size of one 1K entry is allocated on the Packet Forwarding Engine.
- If you configure the sizes of both the IPv4 and IPv6 flow tables, the flow tables are created on the Packet Forwarding Engine based on the size that you specified.

NOTE: When you configure inline active flow monitoring for VPLS flows, include the `vpls-flow-table-size` statement.

```
[edit chassis]
user@router1# set fpc 0 sampling-instance instance-1
user@router1# set fpc 0 inline-services flow-table-size ipv4-flow-table-size 8
user@router1# set fpc 0 inline-services flow-table-size ipv6-flow-table-size 7
```

4. Configure firewall filters.

```
[edit firewall]
user@router1# set family inet filter inet-sample term t1 then sample
user@router1# set family inet filter inet-sample term t1 then accept
```

```

user@router1# set family inet6 filter inet6-sample term t1 then sample
user@router1# set family inet6 filter inet6-sample term t1 then accept

```

5. Associate the firewall filters configured in the previous step with the interfaces on which you want to set up inline active flow monitoring.

```

[edit interfaces]
user@router1# set ge-0/0/4 unit 0 family inet filter input inet-sample
user@router1# set ge-0/0/4 unit 0 family inet address 10.150.1.1/24
user@router1# set ge-0/1/6 unit 0 family inet6 filter input inet6-sample
user@router1# set ge-0/1/6 unit 0 family inet6 address 2001:db8:0:2::1/64

```

6. Commit the configuration.

```

[edit]
user@router1# commit

```

Results

From the configuration mode, confirm your configuration by entering `show services flow-monitoring`, `show forwarding-options sampling`, `show chassis fpc 0`, `show firewall`, and `show interfaces` commands. If the output does not display the intended configuration, repeat the instructions in the example to correct the configuration.

- `show services flow-monitoring`

```

version9 {
  template template1 {
    flow-active-timeout 120;
    flow-inactive-timeout 60;
    template-refresh-rate {
      packets 100;
      seconds 600;
    }
    option-refresh-rate {
      packets 100;
      seconds 600;
    }
    ipv4-template;
  }
}

```

```

}
  version-ipfix {
    template template-v61 {
      flow-active-timeout 150;
      flow-inactive-timeout 100;
      template-refresh-rate {
        seconds 30;
      }
      ipv6-template;
    }
  }
}

```

- show forwarding-options sampling

```

instance {
  instance-1 {
    input {
      rate 1;
    }
    family inet {
      output {
        flow-server 10.50.1.2 {
          port 2055;
          version9 {
            template {
              template1;
            }
          }
        }
        inline-jflow {
          source-address 10.50.1.100;
          flow-export-rate 10;
        }
      }
    }
  }
  family inet6 {
    output {
      flow-server 10.50.1.2 {
        port 2055;
        version-ipfix {
          template {
            template-v61;
          }
        }
      }
    }
  }
}

```



```

    }
  }
}

```

- show interfaces

```

...
ge-0/1/6 {
  vlan-tagging;
  unit 0 {
    family inet6 {
      filter {
        input inet6-sample;
      }
      address 2001:db8:0:2::1/64;
    }
  }
}

ge-0/0/4 {
  vlan-tagging;
  unit 0 {
    family inet {
      filter {
        input inet-sample;
      }
      address 10.150.1.1/24;
    }
  }
}
...

```

Software and Hardware Requirements

- An MX Series router other than MX80
- Junos OS Release 13.2 or later.

NOTE:

- Junos OS Releases earlier than 13.2 also support inline active flow monitoring. However, some of the features discussed in this example are not supported on previous releases.
- You need Junos OS Release 14.2 or later for configuring inline active flow monitoring on T4000 routers with Type 5 FPC.

Overview

Inline active flow monitoring enables you to configure active sampling without making use of a services DPC. This topic explains the basic configuration for enabling inline active flow monitoring for IPv4 and IPv6 flows. You can also configure inline active flow monitoring for VPLS flows. To configure inline active flow monitoring for VPLS flows, you must specify the family as `vpls` and include `vpls-template` at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level.

RELATED DOCUMENTATION

[Understand Inline Active Flow Monitoring | 437](#)

[Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers | 537](#)

Sampling Data Using Flow Aggregation

IN THIS CHAPTER

- Understanding Flow Aggregation | 576
- Enabling Flow Aggregation | 577
- Configuring Flow Aggregation on MX, M and T Series Routers and NFX250 to Use Version 5 or Version 8 cflowd | 578
- Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583
- Configuring Flow Aggregation on PTX Series Routers to Use Version 9 Flow Templates | 596
- Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603
- Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers | 615
- Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows | 620
- Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows | 625
- Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers | 633
- Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers | 637
- Logging cflowd Flows on M and T Series Routers Before Export | 640
- Configuring Next-Hop Address Learning on MX Series and PTX Series Routers for Destinations Accessible Over Multiple Paths | 641

Understanding Flow Aggregation

You can collect an aggregate of sampled flows and send the aggregate to a specified host that runs either the cflowd application available from CAIDA (<http://www.caida.org>) or the newer version 9 format defined in RFC 3954, *Cisco Systems NetFlow Services Export Version 9*. Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process.

By using flow aggregation, you can obtain various types of byte and packet counts of flows through a router. The application collects the sampled flows over a period of 1 minute. At the end of the minute, the number of samples to be exported are divided over the period of another minute and are exported over the course of the same minute.

You configure flow aggregation in different ways, depending on whether you want to export flow records in cflowd version 5 or 8 format, or the separate version 9 format. The latter allows you to sample MPLS, IPv4, IPv6, and peer AS billing traffic. You can also combine configuration statements between the MPLS and IPv4 formats.

NOTE: When PIC-based sampling is enabled, collection of flow statistics for sampled packets on flows in virtual private networks (VPNs) is also supported. No additional CLI configuration is required.

RELATED DOCUMENTATION

[Enabling Flow Aggregation | 577](#)

[Configuring Flow Aggregation on MX, M and T Series Routers and NFX250 to Use Version 5 or Version 8 cflowd | 578](#)

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers | 637](#)

[Logging cflowd Flows on M and T Series Routers Before Export | 640](#)

Enabling Flow Aggregation

Before you can perform flow aggregation, the routing protocol process must export the autonomous system (AS) path and routing information to the sampling process. To enable the export of AS path and the routing information to the sampling process, one or more of the following needs to be configured:

- At the [edit forwarding-options] hierarchy level (for routing instances, at the [edit routing-instance *routing-instance-name* forwarding-options] hierarchy level), configure sampling family or sampling output or sampling instance or monitoring or accounting.
- At the [edit routing-options] hierarchy level (for routing instances, at the [edit routing-instance *routing-instance-name* routing-options] hierarchy level), configure route record.

- At the [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider] hierarchy level, configure forwarding-db-size.

RELATED DOCUMENTATION

[Understanding Flow Aggregation | 576](#)

[Configuring Flow Aggregation on MX, M and T Series Routers and NFX250 to Use Version 5 or Version 8 cflowd | 578](#)

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers | 637](#)

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

[Logging cflowd Flows on M and T Series Routers Before Export | 640](#)

Configuring Flow Aggregation on MX, M and T Series Routers and NFX250 to Use Version 5 or Version 8 cflowd

To enable the collection of cflowd version 5 or version 8 flow formats, include the `flow-server` statement:

```
flow-server hostname {
  aggregation {
    autonomous-system;
    destination-prefix;
    protocol-port;
    source-destination-prefix {
      caida-compliant;
    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  (local-dump | no-local-dump);
  port port-number;
  version format;
}
```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options sampling family (inet | inet6 | mpls) output]

- [edit forwarding-options sampling instance *instance-name* output]
- [edit forwarding-options accounting *name* output cflowd *hostname*]

You must configure the `family inet` statement on logical interface unit 0 on the monitoring interface, as in the following example:

```
[edit interfaces]
sp-3/0/0 {
  unit 0 {
    family inet {
      ...
    }
  }
}
```

NOTE: Boot images for monitoring services interfaces are specified at the [edit chassis images pic] hierarchy level. You must enable the NTP client to make the cflowd feature operable, by including the following configuration:

```
[edit system]
ntp {
  boot-server ntp.example.com;
  server 172.17.28.5;
}
processes {
  ntp enable;
}
```

You can also configure cflowd version 5 for flow-monitoring applications by including the `cflowd` statement at the [edit forwarding-options monitoring *name* family inet output] hierarchy level:

```
cflowd hostname {
  port port-number;
}
```

The following restrictions apply to cflowd flow formats:

- You can configure up to one version 5 and one version 8 flow format at the [edit forwarding-options accounting *name* output] hierarchy level.

- You can configure up to eight version 5 or one version 8 flow format at the [edit forwarding-options sampling family (inet | inet6 | mpls) output] hierarchy level for Routing Engine-based sampling by including the flow-server statement. In contrast, PIC-based sampling allows you to specify one cflowd version 5 server and one version 8 server simultaneously. However, the two cflowd servers must have different IP addresses.
- You can configure up to eight version 5 flow formats at the [edit forwarding-options monitoring *name* output] hierarchy level. Version 8 flow formats and aggregation are not supported for flow-monitoring applications.
- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC are dropped until the synchronization process is complete.
- The configuration includes a proprietary v5 extension template for supporting 4-byte AS information in flow records. Its template version is set to 500, indicating it to be proprietary. All other fields remain the same; the source AS and destination AS are each 4 bytes long, rather than 2 bytes as in the traditional v5 template. This option is available at the [edit forwarding-options sampling family inet output flow-server *server-name* version] hierarchy level.

In the cflowd statement, specify the name or identifier of the host that collects the flow aggregates. You must also include the User Datagram Protocol (UDP) port number on the host and the version, which gives the format of the exported cflowd aggregates. To collect cflowd records in a log file before exporting, include the local-dump statement.

NOTE: You can specify both host (cflowd) sampling and port mirroring in the same configuration; however, only one action takes effect at any one time. Port mirroring takes precedence. For more information, see [Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#).

For cflowd version 8 only, you can specify aggregation of specific types of traffic by including the aggregation statement. This conserves memory and bandwidth by enabling cflowd to export targeted flows rather than all aggregated traffic. To specify a flow type, include the aggregation statement:

```
aggregation {
  autonomous-system;
  destination-prefix;
  protocol-port;
```

```

    source-destination-prefix {
        caida-compliant;
    }
    source-prefix;
}

```

You can include this statement at the following hierarchy levels:

- [edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server *hostname*]
- [edit forwarding-options accounting *name* output cflowd *hostname*]

The `autonomous-system` statement configures aggregation by the AS number; this statement might require setting the separate `cflowd autonomous-system-type` statement to include either origin or peer AS numbers. The `origin` option specifies to use the origin AS of the packet source address in the Source Autonomous System `cflowd` field. The `peer` option specifies to use the peer AS through which the packet passed in the Source Autonomous System `cflowd` field. By default, `cflowd` exports the origin AS number.

The `destination-prefix` statement configures aggregation by the destination prefix only.

The `protocol-port` statement configures aggregation by the protocol and port number; requires setting the separate `cflowd port` statement.

The `source-destination-prefix` statement configures aggregation by the source and destination prefix. Version 2.1b1 of CAIDA's `cflowd` application does not record source and destination mask length values in compliance with CAIDA's *cflowd Configuration Guide*, dated August 30, 1999. If you configure the `caida-compliant` statement, the Junos OS complies with Version 2.1b1 of `cflowd`. If you do not include the `caida-compliant` statement in the configuration, the Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*.

The `source-prefix` statement configures aggregation by the source prefix only.

Collection of sampled packets in a local ASCII file is not affected by the `cflowd` statement.

The following commands enable Routing Engine- and PIC-based sampling at the `set forwarding options sampling` hierarchy level:

- `set input rate rate`
- `set input run-length length`
- `set family inet output flow-server flowcollector port udp port`
- `set family inet output flow-server flowcollector no-local-dump`
- `set family inet output flow-server flowcollector version <5/8>`

The following commands enable Routing Engine- and PIC-based sampling at the set interfaces hierarchy level:

- *interface to be sampled* unit *unit* family inet filter *input/output filtername*

The following commands enable Routing Engine- and PIC-based sampling at the set firewall family hierarchy level:

- set inet filter *filtername* term 1 then count *filtername*ing
- set inet filter *filtername* term 1 then sample
- set inet filter *filtername* term 1 then accept

The following command enables PIC-based sampling at the set forwarding options sampling hierarchy level:

- set family inet output interface *sp-*/*/** source address *source address*

The following example shows a PIC-based flow aggregation configuration using version 5:

```
family inet {
  output {
    flow-inactive-timeout 15;
    flow-active-timeout 60;
    flow-server 203.0.113.165 {
      port 9996;
      version 5;
    }
    interface sp-2/2/0 {
      engine-id 4;
      source-address 203.0.113.126;
    }
  }
}
```

The following example shows an Routing Engine-based flow aggregation configuration using version 5:

```
family inet {
  output {
    flow-inactive-timeout 15;
    flow-active-timeout 60;
    flow-server 203.0.113.165 {
      port 9996;
      source-address 203.0.113.126;
    }
  }
}
```

```

        version 5;
    }
}

```

RELATED DOCUMENTATION

[Understanding Flow Aggregation | 576](#)

[Enabling Flow Aggregation | 577](#)

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

IN THIS SECTION

- [Configuring the Traffic to Be Sampled | 584](#)
- [Configuring the Version 9 Template Properties | 585](#)
- [Customizing Template ID, Observation Domain ID, and Source ID for Version 9 Flow Templates | 586](#)
- [Restrictions | 587](#)
- [Fields Included in Each Template Type | 588](#)
- [MPLS Sampling Behavior | 590](#)
- [Verification | 590](#)
- [Examples: Configuring Version 9 Flow Templates | 590](#)

Use of version 9 flow template enables you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector does not affect the router configuration.

NOTE: Version 9 requires that you install a services PIC, such as the Adaptive Services PIC or MS-PIC in the router. On MX Series routers, the MS-DPC fulfills this requirement. For more information on determining which services PIC is suitable for your router, see *Enabling Service Packages* or the appropriate hardware documentation.

NOTE: If multiple protocol families are configured for a particular flow collector, the export packets originates from multiple Source IDs, with each Source ID corresponding to a particular protocol. The multiple Source IDs do not indicate that the export packets are originating from multiple Service PICs.

Configuring the Traffic to Be Sampled

To specify sampling of IPv4, IPv6, MPLS, or peer AS billing traffic, include the appropriate configuration of the family statement at the [edit forwarding-options sampling] hierarchy level:

```
[edit forwarding-options]
sampling {
  family (inet | inet6 | mpls);
}
```

You can include family inet ,family inet6, or family mpls.

NOTE: If you specify sampling for peer AS billing traffic, the family statement supports only IPv4 and IPv6 traffic (inet or inet6). Peer AS billing traffic is enabled only at the global instance hierarchy level and is not available for per Packet Forwarding Engine instances.

After you specify the family of traffic to be sampled, configure the sampling parameters such as:

- Maximum packet length (beyond which the packets are truncated).
- Maximum packets to be sampled per second (beyond which the packets are dropped).
- Rate (for example, if you specify 10, every 10th packet is sampled).

- Run length (which specifies the number of packets to be sampled after the trigger; that is, if the rate is set to 10 and run-length to 5, five packets starting at the 10th packet are sampled).

```
[edit forwarding-options sampling]
input {
  maximum-packet-length bytes
  max-packets-per-second number;
  rate number;
  run-length number;
}
```

Configuring the Version 9 Template Properties

To define the Version 9 templates, include the following statements at the [edit services flow-monitoring version9] hierarchy level:

```
[edit services flow-monitoring version9]
template template-name {
  options-template-id
  template-id
  source-id
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  option-refresh-rate packets packets seconds seconds;
  template-refresh-rate packets packets seconds seconds;
  (ipv4-template | ipv6-template | mpls-ipv4-template | mpls-template | peer-as-billing-template) {
    label-position [ positions ];
  }
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the `template template-name` statement.
- You then specify each template for the appropriate type of traffic by including the `ipv4-template`, `ipv6-template`, `mpls-ipv4-template`, or `mpls-template`.
- If the template is used for MPLS traffic, you can also specify up to three label positions for the MPLS header label data by including the `label-position` statement; the default values are [1 2 3].
- Within the template definition, you can optionally include values for the `flow-active-timeout` and `flow-inactive-timeout` statements. These statements have specific default and range values when they are

used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds. Values you specify in template definitions override the global timeout values configured at the [edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server] hierarchy level.

- You can also include settings for the option-refresh-rate and template-refresh-rate statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the seconds option, the default value is 60 and the range is from 10 through 600. For the packets option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPv6 traffic on a media interface, the following configuration is supported:

```
interfaces interface-name {
  unit 0 {
    family inet6 {
      sampling {
        input;
        output;
      }
    }
  }
}
```

Customizing Template ID, Observation Domain ID, and Source ID for Version 9 Flow Templates

Starting in Junos OS Release 14.1, you can define a Version 9 flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector does not affect the router configuration. You can specify the unique identifier for the version 9 and IPFIX templates. The identifier of a template is locally unique within a combination of a transport session and an observation domain. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, you can reassign template IDs.

This functionality to configure template ID, options template ID, observation domain ID, and source ID is supported on all routers with MPCs.

NOTE: The template IDs that include MPLS and MPLS-IPv4 template ID are applicable for IPFIX only. The V9 format carries a different template ID.

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

For more information about specifying the source ID, observation domain ID, template ID, and options template ID for version 9 and IPFIX flows, see ["Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows" on page 620](#) and ["Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows" on page 625](#).

Restrictions

The following restrictions apply to version 9 templates:

- You cannot apply the two different types of flow aggregation configuration at the same time.
- Flow export based on an `mpls-ipv4` template assumes that the IPv4 header follows the MPLS header. In the case of Layer 2 VPNs, the packet on the provider router (P router) looks like this:

MPLS | Layer 2 Header | IPv4

In this case, `mpls-ipv4` flows are not created on the PIC, because the IPv4 header does not directly follow the MPLS header. Packets are dropped on the PIC and are accounted as parser errors.

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created on the monitoring PIC only after the route record resynchronization operation is complete, which is 60 seconds after the PIC comes up. Any packets sent to the PIC are dropped until the synchronization process is complete.

NOTE: Because the forwarding of a packet that arrives with MPLS labels is performed based on the MPLS label and not based on the IP address contained in the packet, the packet is sampled at the output interface with the MPLS label that was popped not being available at the time of sampling. In such a case, depending on the incoming interface (IIF), the VRF index is identified

and the route for the sampled packet is determined in the VRF table. Because a specific route is not available in the VRF that is different from the VRF on which the packet is received, the Output Interface Index, Source Mask, and Destination Mask fields are incorrectly populated. This behavior occurs when an IPv4 template is applied as a firewall filter on an egress interface with sample as the action.

Fields Included in Each Template Type

The following fields are common to all template types:

- Input interface
- Output interface
- Number of bytes
- Number of packets
- Flow start time
- Flow end time

The IPv4 template includes the following specific fields:

- IPv4 Source Address
- IPv4 Destination Address
- L4 Source Port
- L4 Destination Port
- IPv4 ToS
- IPv4 Protocol
- ICMP type and code
- TCP Flags
- IPv4 Next Hop Address
- Source autonomous system (AS) number
- Destination AS number

The IPv6 template includes the following specific fields:

- IPv6 Source Address and Mask
- IPv6 Destination Address and Mask
- L4 Source Port
- L4 Destination Port
- IPv6 ToS
- IPv6 Protocol
- TCP Flags
- IP Protocol Version
- IPv6 Next Hop Address
- Egress Interface Information
- Source Autonomous System (AS) number
- Destination AS number

The MPLS template includes the following specific fields:

- MPLS Label #1
- MPLS Label #2
- MPLS Label #3
- MPLS EXP Information
- FEC IP Address

The MPLS-IPv4 template includes all the fields found in the IPv4 and MPLS templates.

The peer AS billing template includes the following specific fields:

- IPv4 Class of Service (ToS)
- Ingress Interface
- BGP IPv4 Next Hop Address
- BGP Peer Destination AS Number

MPLS Sampling Behavior

This section describes the behavior when MPLS sampling is used on egress interfaces in various scenarios (label pop or swap) on provider routers (P routers). For more information on configuration and background specific to MPLS applications, see the [MPLS Applications User Guide](#).

- You configure MPLS sampling on an egress interface on the P router and configure an MPLS flow aggregation template. The route action is label *pop* because penultimate hop popping (PHP) is enabled.

With the current capability of applying MPLS templates, MPLS flows are created.

- As in the first case, you configure MPLS sampling on an egress interface on the P router and configure an MPLS flow aggregation template. The route action is label swap and the swapped label is 0 (explicit null).

The resulting behavior is that MPLS packets are sent to the PIC. The flow being sampled corresponds to the label before the swap.

- You configure a Layer 3 VPN network, in which a customer edge router (CE-1) sends traffic to a provider edge router (PE-A), through the P router, to a similar provider edge router (PE-B) and customer edge router (CE-2) on the remote end.

The resulting behavior is that you cannot sample MPLS packets on the PE-A to P router link.

Verification

To verify the configuration properties, you can use the `show services accounting aggregation template-name name operational` mode command.

All other `show services accounting` commands also support version 9 templates, except for `show services accounting flow-detail` and `show services accounting aggregation aggregation-type`. For more information about operational mode commands, see the [CLI Explorer](#).

Examples: Configuring Version 9 Flow Templates

The following example shows a version 9 template configuration:

```
services {
  flow-monitoring {
    version9 {
      template ip-template {
        flow-active-timeout 20;
        flow-inactive-timeout 120;
      }
    }
  }
}
```

```

        ipv4-template;
    }
    template mpls-template-1 {
        mpls-template {
            label-position [1 3 4];
        }
    }
    template mpls-ipv4-template-1 {
        mpls-ipv4-template {
            label-position [1 5 7];
        }
    }
    template vpls-template-1 {
        vpls-template;
    }
}
}
}
}

```

The following example shows a firewall filter configuration for MPLS traffic:

```

firewall {
    family mpls {
        filter mpls_sample {
            term default {
                then {
                    accept;
                    sample;
                }
            }
        }
    }
}

```

The following example applies the MPLS sampling filter on a networking interface and configures the AS PIC to accept both IPv4 and MPLS traffic:

```

interfaces {
    at-0/1/1 {
        unit 0 {

```

```

        family mpls {
            filter {
                input mpls_sample;
            }
        }
    }
}
sp-7/0/0 {
    unit 0 {
        family inet;
        family mpls;
    }
}
}

```

The following example applies the MPLS version 9 template to the sampling output and sends it to the AS PIC:

```

forwarding-options {
    sampling {
        input {
            family mpls {
                rate 1;
            }
        }
        family mpls {
            output {
                flow-active-timeout 60;
                flow-inactive-timeout 30;
                flow-server 192.0.2.4 {
                    port 2055;
                    version9 {
                        template mpls-ipv4-template-1;
                    }
                }
            }
        }
        interface sp-7/0/0 {
            source-address 198.51.100.1;
        }
    }
}

```

```

    }
}

```

The following example shows a firewall filter configuration for the peer AS billing traffic:

```

firewall {
  family inet {
    filter peer-as-filter {
      term 0 {
        from {
          destination-class dcu-1;
          interface ge-2/1/0;
          forwarding-class class-1;
        }
        then count count_team_0;
      }
    }
    term 1 {
      from {
        destination-class dcu-2;
        interface ge-2/1/0;
        forwarding-class class-1;
      }
      then count count_team_1;
    }
    term 2 {
      from {
        destination-class dcu-3;
        interface ge-2/1/0;
        forwarding-class class-1;
      }
      then count count_team_2;
    }
  }
}
}

```

The following example applies the peer AS firewall filter as a filter attribute under the forwarding-options hierarchy for CoS-level data traffic usage information collection:

```
forwarding-options {
  family inet {
    filter output peer-as-filter;
  }
}
```

The following example applies the peer AS DCU policy options to collect usage statistics for the traffic stream for as-path ingressing at a specific input interface with the firewall configuration hierarchy applied as Forwarding Table Filters (FTFs). The configuration functionality with CoS capability can be achieved through FTFs for destination-class usage with forwarding-class for specific input interfaces:

```
policy-options {
  policy-statement P1 {
    from {
      protocol bgp;
      neighbor 10.2.25.5; #BGP router configuration;
      as-path AS-1; #AS path configuration;
    }
    then destination-class dcu-1; #Destination class configuration;
  }
  policy-statement P2 {
    from {
      neighbor 203.0.113.5;
      as-path AS-2;
    }
    then destination-class dcu2;
  }
  policy-statement P3 {
    from {
      protocol bgp;
      neighbor 192.0.2.129;
      as-path AS-3;
    }
    then destination-class dcu3;
  }
  as-path AS-1 3131:1111:1123;
  as-path AS-2 100000;
```

```
as-path AS-3 192:29283:2;
}
```

The following example applies the vpls version 9 template to enable sampling of traffic for billing purposes:

```
forwarding-options {
  sampling {
  }
  input {
    rate 1;
  }
  family inet {
    output {
      flow-server 10.209.15.58 {
        port 300;
        version9 {
          template {
            peer-as;
          }
        }
      }
      interface sp-5/2/0 {
        source-address 203.0.113.133;
      }
    }
  }
}
family inet {
  filter {
    output peer-as-filter;
  }
}
```

Release History Table

Release	Description
14.1	Starting in Junos OS Release 14.1, you can define a Version 9 flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic.

RELATED DOCUMENTATION

[Understanding Flow Aggregation | 576](#)

[Enabling Flow Aggregation | 577](#)

[Configuring Flow Aggregation on MX, M and T Series Routers and NFX250 to Use Version 5 or Version 8 cflowd | 578](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

Configuring Flow Aggregation on PTX Series Routers to Use Version 9 Flow Templates

IN THIS SECTION

- [Configuring the Version 9 Template Properties | 597](#)
- [Restrictions | 598](#)
- [Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates | 598](#)
- [Fields Included in the IPv4 Templates for PTX Series Routers | 598](#)
- [Fields Included in the IPv6 Templates for PTX Series Routers | 600](#)
- [Verification | 601](#)
- [Example: Configuring an version 9 Flow Templates and Flow Sampling | 602](#)

You can define a flow record template suitable for IPv4 traffic or IPv6 traffic using a version 9 flow template. Templates and the fields included in the template are transmitted to the collector periodically. The collector does not affect the router configuration. You can define template refresh rate, flow active timeout and inactive timeout.

If flow records are being sent for multiple protocol families (for example, for IPv4 and IPv6), each protocol family flow will have a unique Observation Domain ID.

Configuring the Version 9 Template Properties


To define the version 9 templates, include the following statements at the [edit services flow-monitoring version9] hierarchy level:

```
[edit services flow-monitoring version9]
template name {
    options-template-id
    template-id
    observation-domain-id
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    option-refresh-rate packets packets seconds seconds;
    template-refresh-rate packets packets seconds seconds;
    (ipv4-template | ipv6-template);
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the `template name` statement.
- You specify each template for the appropriate type of traffic by including the `ipv4-template` or `ipv6-template`.
- Within the template definition, you can optionally include values for the `flow-active-timeout` and `flow-inactive-timeout` statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds.
- You can also include settings for the `option-refresh-rate` and `template-refresh-rate` statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the `seconds` option, the default value is 600 and the range is from 10 through 600. For the `packets` option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPv6 traffic on a media interface, the following configuration is supported:

```
interfaces interface-name {
    unit 0 {
        family inet6 {
            sampling {
                input;
                output;
```

Restrictions

The following restrictions apply to version 9 templates:

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created only after the route record resynchronization operation is complete, which takes 120 seconds.

Customizing Template ID, Observation Domain ID, and Source ID for Version 9 flow Templates

NOTE: For PTX Series routers with third generation FPCs installed, the FPC’s slot number is used for the observation domain ID.

Use of version 9 flow templates allow you to define a flow record template suitable for IPv4 traffic or IPv6 traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector does not need to be aware of the router configuration. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, template IDs can be reassigned.

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

Fields Included in the IPv4 Templates for PTX Series Routers

Table 106 on page 599 shows the fields that are available in the IPv4 templates.

Table 106: IPv4 Template Fields

Field	Element ID
IPv4 Source Address	8
IPv4 Destination Address	12
IPv4 TOS	5
IPv4 Protocol	4
L4 Source Port	7
L4 Destination Port	11
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
BGP Next Hop Address	18
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22

Table 106: IPv4 Template Fields *(Continued)*

Field	Element ID
Time the flow ended with respect to system up time (FPC up time)	21
IPv4 Next Hop Address	15
IPv4 Source Mask	9
IPv4 Destination Mask	13
TCP Flags	6
IP Protocol Version	60

Fields Included in the IPv6 Templates for PTX Series Routers

[Table 107 on page 600](#) shows the fields that are available in the IPv6 templates.

Table 107: IPv6 Template Fields

Field	Element ID
IPv6 Source Address	27
IPv6 Destination Address	28
IPv6 TOS	5
IPv6 Protocol	4
L4 Source Port	7
L4 Destination Port	11

Table 107: IPv6 Template Fields *(Continued)*

Field	Element ID
ICMP Type and Code	32
Input Interface	10
Source AS	16
Destination AS	17
Output Interface	14
Number of Flow Bytes	1
Number of Flow Packets	2
Time the flow started with respect to system up time (FPC up time)	22
Time the flow ended with respect to system up time (FPC up time)	21
IPv6 Next Hop Address	62
IPv6 Source Mask	29
IPv6 Destination Mask	30
TCP Flags	6
IP Protocol Version	60

Verification

The following show commands are supported for version 9:

- show services accounting flow inline-jflow fpc-slot *fpc-slot*
- show services accounting errors inline-jflow fpc-slot *fpc-slot*
- show services accounting status inline-jflow fpc-slot *fpc-slot*

Example: Configuring an version 9 Flow Templates and Flow Sampling

The following is a sample version 9 template configuration:

```
services {
  flow-monitoring {
    version9 {
      template ipv4 {
        flow-active-timeout 60;
        flow-inactive-timeout 70;
        template-refresh-rate seconds 30;
        option-refresh-rate seconds 30;
        ipv4-template;
      }
    }
  }
}
```

```
chassis;
  fpc 0 {
    sampling-instance s1;
  }
```

The following example applies the version 9 template to enable sampling of traffic for billing:

```
forwarding-options {
  sampling {
    instance {
      s1 {
        input {
          rate 10;
        }
        family inet {
          output {
```

Configuring Inline Active Flow Monitoring on PTX Series Routers | 540

version9 (Flow Monitoring) | 1537

ipv4-template | 1191

ipv6-template | 1195

IN THIS SECTION

- Configuring the IPFIX Template Properties | 604
- Restrictions | 605
- Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates | 606

- [IPFIX Templates | 606](#)
- [Verification | 606](#)
- [Example: Configuring IPFIX Flow Templates and Flow Sampling | 607](#)
- [Example: Configuring Inline Active Flow Monitoring Version 9 Flow Templates and Flow Sampling | 608](#)
- [Example: Configuring IPFIX Flow Templates and Flow Sampling | 612](#)

Use of IPFIX allows you to define a flow record template suitable for IPv4 traffic or IPv6 traffic. Templates are transmitted to the collector periodically, and the collector does not affect the router configuration. You can define template refresh rate, flow active timeout and inactive timeout.

If flow records are being sent for multiple protocol families (for example, for IPv4 and IPv6), each protocol family flow has a unique Observation Domain ID. The following sections contain additional information:

Starting with Junos OS Release 17.3R1, IPFIX flow templates are supported on QFX10002 switches.

Starting with Junos OS Release 17.4R1, IPFIX flow templates are supported on QFX10008 and QFX10016 switches.

Starting with Junos OS Release 19.4R1, IPFIX flow templates are supported on SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vSRX, and vSRX3.0 devices.

Starting with Junos OS Release 20.1R1, IPFIX flow templates are supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.

Starting with Junos OS Release 20.4R1, IPFIX flow templates are supported on NFX150, NFX250 NextGen, and NFX350 devices.

Configuring the IPFIX Template Properties

To define the IPFIX templates, include the following statements at the [edit services flow-monitoring version-ipfix] hierarchy level:

```
[edit services flow-monitoring version-ipfix]
template template-name {
  options-template-id
  template-id
  observation-domain-id
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
```

```

option-refresh-rate packets packets seconds seconds;
template-refresh-rate packets packets seconds seconds;
(ipv4-template | ipv6-template);
}

```

The following details apply to the configuration statements:

- You assign each template a unique name by including the `template template-name` statement.
- You then specify each template for the appropriate type of traffic by including the `ipv4-template` or `ipv6-template`.
- Within the template definition, you can optionally include values for the `flow-active-timeout` and `flow-inactive-timeout` statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds.
- You can also include settings for the `option-refresh-rate` and `template-refresh-rate` statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the `seconds` option, the default value is 600 and the range is from 10 through 600. For the `packets` option, the default value is 4800 and the range is from 1 through 480,000.
- To filter IPv6 traffic on a media interface, the following configuration is supported:

```

interfaces interface-name {
    unit 0 {
        family inet6 {
            sampling {
                input;
                output;
            }
        }
    }
}

```

Restrictions

The following restrictions apply to IPFIX templates:

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works

correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.

- Flows are created only after the route record resynchronization operation is complete, which takes 120 seconds.
- The VLAN ID field is updated when a new flow record is created and so, any change in VLAN ID after the record has been created might not be updated in the record.

Customizing Template ID, Observation Domain ID, and Source ID for IPFIX flow Templates

Starting in Junos OS Release 14.1, you can define an IPFIX flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector need not be aware of the router configuration. You can specify the unique identifier for the version 9 and IPFIX templates. The identifier of a template is locally unique within a combination of a transport session and an observation domain. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, you can reassign template IDs.

This functionality to configure template ID, options template ID, observation domain ID, and source ID is supported on all routers with MPCs.

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

For more information about specifying the source ID, observation domain ID, template ID, and options template ID for version 9 and IPFIX flows, see ["Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows" on page 620](#) and ["Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows" on page 625](#).

IPFIX Templates

For information about the definitions of the fields included in IPFIX IPv4 and IPv6 templates, see ["IPFIX and Version 9 Templates" on page 442](#).

Verification

The following show commands are supported for IPFIX:

- show services accounting flow inline-jflow fpc-slot *fpc-slot*
- show services accounting errors inline-jflow fpc-slot *fpc-slot*
- show services accounting status inline-jflow fpc-slot *fpc-slot*

Example: Configuring IPFIX Flow Templates and Flow Sampling

The following example shows an IPFIX template configuration:

```
services {
  flow-monitoring {
    version-ipfix {
      template ipv4 {
        flow-active-timeout 60;
        flow-inactive-timeout 70;
        template-refresh-rate seconds 30;
        option-refresh-rate seconds 30;
        ipv4-template;
      }
    }
  }
}
```

```
chassis;
  fpc 0 {
    sampling-instance s1;
  }
```

The following example applies the IPFIX template to enable sampling of traffic for billing:

```
forwarding-options {
  sampling {
    instance {
      s1 {
        input {
          rate 10;
        }
        family inet {
          output {
```

```

        flow-server 192.0.2.2 {
            port 2055;
            version-ipfix {
                template {
                    ipv4;
                }
            }
            inline-jflow {
                source-address 198.51.100.1;
            }
        }
    }
}

```

Example: Configuring Inline Active Flow Monitoring Version 9 Flow Templates and Flow Sampling

The following example shows inline Active Flow Monitoring version 9 IPv4 template configuration:

```

services {
    flow-monitoring {
        version9 {
            template ipv4-v9 {
                flow-active-timeout 60;
                flow-inactive-timeout 15;
                template-refresh-rate {
                    packets 1000;
                }
                option-refresh-rate {
                    seconds 100;
                }
                ipv4-template;
            }
        }
    }
}

```

The following example shows inline Active Flow Monitoring version 9 IPv6 template configuration:

```
services {
  flow-monitoring {
    version9 {
      template ipv6-v9 {
        flow-active-timeout 60;
        flow-inactive-timeout 15;
        template-refresh-rate {
          packets 1000;
        }
        option-refresh-rate {
          seconds 100;
        }
        Ipv6-template;
      }
    }
  }
}
```

The following example shows inline Active Flow Monitoring version 9 IPv4 sampling traffic and export configuration:

```
forwarding-options {
  sampling {
    traceoptions {
      file testsample size 1g world-readable;
      flag all;
    }
    instance {
      sample-ins1 {
        input {
          rate 1;
          run-length 0;
        }
        family inet {
          output {
            flow-server 10.207.18.113 {
              port 2055;
              version9 {
                template {
```


Example: Configuring IPFIX Flow Templates and Flow Sampling

The following example shows IPFIX IPv4 template configuration:

```
flow-monitoring {
  version-ipfix {
    template ipv4-ipfix {
      flow-active-timeout 60;
      flow-inactive-timeout 60;
      template-refresh-rate {
        packets 1000;
        seconds 30;
      }
      option-refresh-rate {
        packets 500;
        seconds 60;
      }
      ipv4-template;
    }
  }
}
```

The following example shows IPFIX IPv6 template configuration:

```
flow-monitoring {
  version-ipfix {
    template ipv6-ipfix {
      flow-active-timeout 60;
      flow-inactive-timeout 60;
      template-refresh-rate {
        packets 1000;
        seconds 30;
      }
      option-refresh-rate {
        packets 500;
        seconds 60;
      }
      Ipv6-template;
    }
  }
}
```

The following example shows IPFIX IPv4 sampling traffic and export configuration:

```
forwarding-options {
  sampling {
    traceoptions {
      file testsample size 1g world-readable;
      flag all;
    }
    instance {
      sample-ins1 {
        input {
          rate 1;
          run-length 0;
        }
        family inet {
          output {
            flow-server 10.207.18.113 {
              port 4739;
              version-ipfix {
                template {
                  ipv4-ipfix;
                }
              }
            }
            inline-jflow {
              source-address 10.207.18.232;
              flow-export-rate 2;
            }
          }
        }
      }
    }
  }
}
```

The following example shows IPFIX IPv6 sampling traffic and export configuration:

```
forwarding-options {
  sampling {
    traceoptions {
      file testsample size 1g world-readable;
```



```

    flag all;
  }
  instance {
    sample-ins1 {
      input {
        rate 1;
        run-length 0;
      }
      family inet {
        output {
          flow-server 2001::2 {
            port 4739;
            version9 {
              template {
                ipv6-ipfix;
              }
            }
          }
          inline-jflow {
            source-address 2001::1;
            flow-export-rate 2;
          }
        }
      }
    }
  }
}
}
}
}
}
}
```

Release History Table

Release	Description
20.4R1	Starting with Junos OS Release 20.4R1, IPFIX flow templates are supported on NFX150, NFX250 NextGen, and NFX350 devices.
20.1R1	Starting with Junos OS Release 20.1R1, IPFIX flow templates are supported on SRX300, SRX320, SRX340, SRX345, and SRX550HM devices.
19.4R1	Starting with Junos OS Release 19.4R1, IPFIX flow templates are supported on SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vSRX, and vSRX3.0 devices.

17.4R1	Starting with Junos OS Release 17.4R1, IPFIX flow templates are supported on QFX10008 and QFX10016 switches.
17.2R1	Starting with Junos OS Release 17.3R1, IPFIX flow templates are supported on QFX10002 switches.
14.2	Fragment Identification (Starting in Junos OS Release 14.2)
14.2	IPv6 Extension Headers (Starting in Junos OS Release 14.2)
14.1	Starting in Junos OS Release 14.1, you can define an IPFIX flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic.

RELATED DOCUMENTATION

[Understanding Flow Aggregation | 576](#)

[Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers | 633](#)

[Enabling Flow Aggregation | 577](#)

[Configuring Flow Aggregation on MX, M and T Series Routers and NFX250 to Use Version 5 or Version 8 cflowd | 578](#)

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers

IN THIS SECTION

- [Configuring the IPFIX Template Properties | 616](#)
- [Restrictions | 617](#)
- [Customizing Template ID, Observation Domain ID, and Source ID for IPFIX Flow Templates | 617](#)
- [Verification | 618](#)
- [Example: Configuring an IPFIX Flow Template and Flow Sampling | 618](#)

Use of IPFIX allows you to define a flow record template suitable for IPv4 traffic or IPv6 traffic. Templates are transmitted to the collector periodically, and the collector is not aware of the router configuration. You can define template refresh rate, flow active timeout and inactive timeout.

If flow records are being sent for multiple protocol families (for example, for IPv4 and IPv6), each protocol family flow will have a unique Observation Domain ID.

To learn about the fields included in the templates, see ["Understanding Inline Active Flow Monitoring" on page 437](#).

Configuring the IPFIX Template Properties

To define the IPFIX templates, include the following statements at the [edit services flow-monitoring version-ipfix] hierarchy level:

```
[edit services flow-monitoring version-ipfix]
template name {
    options-template-id
    template-id
    observation-domain-id
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    option-refresh-rate packets packets seconds seconds;
    template-refresh-rate packets packets seconds seconds;
    (ipv4-template | ipv6-template);
}
```

The following details apply to the configuration statements:

- You assign each template a unique name by including the `template name` statement.
- You then specify each template for the appropriate type of traffic by including the `ipv4-template` or `ipv6-template`.
- Within the template definition, you can optionally include values for the `flow-active-timeout` and `flow-inactive-timeout` statements. These statements have specific default and range values when they are used in template definitions; the default is 60 seconds and the range is from 10 through 600 seconds.
- You can also include settings for the `option-refresh-rate` and `template-refresh-rate` statements within a template definition. For both of these properties, you can include a timer value (in seconds) or a packet count (in number of packets). For the `seconds` option, the default value is 600 and the range is from 10 through 600. For the `packets` option, the default value is 4800 and the range is from 1 through 480,000.

- To filter IPv6 traffic on a media interface, the following configuration is supported:

```

interfaces interface-name {
    unit 0 {
        family inet6 {
            sampling {
                input;
                output;
            }
        }
    }
}

```

Restrictions

The following restrictions apply to IPFIX templates:

- Outbound Routing Engine traffic is not sampled. A firewall filter is applied as output on the egress interface, which samples packets and exports the data. For transit traffic, egress sampling works correctly. For internal traffic, the next hop is installed in the Packet Forwarding Engine but sampled packets are not exported.
- Flows are created only after the route record resynchronization operation is complete, which takes 120 seconds.

Customizing Template ID, Observation Domain ID, and Source ID for IPFIX Flow Templates

NOTE: For PTX Series routers with third generation FPCs installed, the FPC's slot number is used for the observation domain ID.

Use of IPFIX flow templates allow you to define a flow record template suitable for IPv4 traffic, IPv6 traffic, MPLS traffic, a combination of IPv4 and MPLS traffic, or peer AS billing traffic. Templates and the fields included in the template are transmitted to the collector periodically, and the collector does not need to be aware of the router configuration. Template IDs 0 through 255 are reserved for template sets, options template sets, and other sets for future use. Template IDs of data sets are numbered from 256 through 65535. Typically, this information element or field in the template is used to define the characteristics or properties of other information elements in a template. After a restart of the export process of templates is performed, template IDs can be reassigned.

The corresponding data sets and option data sets contain the value of the template IDs and options template IDs respectively in the set ID field. This method enables the collector to match a data record with a template record.

Verification

The following show commands are supported for IPFIX:

- `show services accounting flow inline-jflow fpc-slot fpc-slot`
- `show services accounting errors inline-jflow fpc-slot fpc-slot`
- `show services accounting status inline-jflow fpc-slot fpc-slot`

Example: Configuring an IPFIX Flow Template and Flow Sampling

The following is a sample IPFIX template configuration:

```
services {
  flow-monitoring {
    version-ipfix {
      template ipv4 {
        flow-active-timeout 60;
        flow-inactive-timeout 70;
        template-refresh-rate seconds 30;
        option-refresh-rate seconds 30;
        ipv4-template;
      }
    }
  }
}
```

```
chassis;
  fpc 0 {
    sampling-instance s1;
  }
```

The following example applies the IPFIX template to enable sampling of traffic for billing:

```
forwarding-options {
  sampling {
    instance {
      s1 {
        input {
          rate 10;
        }
        family inet {
          output {
            flow-server 11.11.4.2 {
              port 2055;
              version-ipfix {
                template {
                  ipv4;
                }
              }
            }
          }
          inline-jflow {
            source-address 11.11.2.1;
          }
        }
      }
    }
  }
}
```

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring on PTX Series Routers | 540](#)

[version-ipfix \(Services\) | 1540](#)

[ipv4-template | 1191](#)

[ipv6-template | 1195](#)

Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows

IN THIS SECTION

- [Considerations for MX and QFX Series | 621](#)
- [Considerations for PTX Series | 624](#)

For IPFIX flows, an identifier of an observation domain is locally unique to an exporting process of the templates. The export process uses the observation domain ID to uniquely identify to the collection process in which the flows were metered. We recommend that you configure this ID to be unique for each IPFIX flow. A value of 0 indicates that no specific observation domain is identified by this information element. Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an IPFIX device.

For version 9 flows, a 32-bit value that identifies the Exporter Observation Domain is called the source ID. NetFlow collectors use the combination of the source IP address and the source ID field to separate different export streams originating from the same exporter.

To specify the observation domain ID for IPFIX flows, include the `observation-domain-id domain-id` statement at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level.

```
[edit services flow-monitoring version-ipfix]
template template-name {
    observation-domain-id domain-id;
}
```

To specify the source ID for version 9 flows, include the `source-id source-id` statement at the `[edit services flow-monitoring version9 template template-name]` hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
    source-id source-id;
}
```

Considerations for MX and QFX Series

If you configure the same Observation Domain ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base observation domain ID is transmitted in the flow. The actual observation domain ID is derived from the value you configure and also in conjunction with other parameters such as the slot number, lookup chip (LU) instance, Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two IPFIX devices.

Until Junos OS Release 13.3, the observation domain ID is predefined and is set to a fixed value, which is derived from the combination of FPC slot, sampling protocol, PFE Instance and LU Instance fields. This derivation creates a unique observation domain per LU per family. Starting with Junos OS Release 14.1, you can configure the observation domain ID, which causes the first 8 bits of the field to be configured.

The following modifications have been made:

- FPC slots are expanded to 8 bits to enable more slots to be configured in an MX Series Virtual Chassis configuration.
- 8 bits of the configured observation domain ID are used.
- You can configure a value for the observation domain ID in the range of 0 through 255.
- The Protocol field is increased to 3 bits to provide support for additional protocols in inline flow monitoring.
- You can associate the observation domain ID with templates by using the `observation-domain-id domain-id` statement at the `[edit services flow- monitoring version-ipfix template template-name]` hierarchy level.

[Table 108 on page 622](#) describes observation domain ID values for different combinations of the configured domain ID, protocol family, FPC slot, and the Packet Forwarding Engine and lookup chip instances.

Table 108: MX Series: Example of Observation Domain ID

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	IPv4 (0)	1	1	0	0000 0000 0000 1000 0000 0001 0000 0001 0x00080101
None	IPv6 (1)	1	1	0	0000 0000 0000 1001 0000 0001 0000 0001 0x00090101
None	VPLS (2)	1	1	0	0000 0000 0000 1010 0000 0001 0000 0001 0x000A0101
None	MPLS (3)	1	1	0	0000 0000 0000 1011 0000 0001 0000 0001 0x000B0101
4	IPv4 (0)	1	1	0	0000 0100 0000 1000 0000 0001 0000 0001 0x04080101

Table 108: MX Series: Example of Observation Domain ID (Continued)

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
190	IPv4 (0)	1	1	0	1101 1110 0000 1000 0000 0001 0000 0001 0xBE080101
4	IPv4 (0)	2	1	1	0000 0100 0000 1000 0000 0010 0001 0001 0x04080211
4	IPv6 (1)	1	1	0	0000 0100 0000 1001 0000 0001 0001 0000 0x04090110
190	IPv6 (1)	1	1	0	1101 1110 0000 1001 0000 0001 0001 0000 0xBE090110
4	VPLS (2)	2	2	0	0000 0100 0000 1010 0000 0010 0010 0000 0x040A0220

Table 108: MX Series: Example of Observation Domain ID (Continued)

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
10	IPv4 (0)	28	2	1	0000 1010 0000 1000 0001 1100 0010 0001 0x0A081C21

Considerations for PTX Series

When you configure the observation domain ID, the software attaches the ID to a particular template type.

If you configure the same observation domain ID for two different template types, such as for IPv4 and IPv6, this does not impact flow monitoring, because the configured ID is not what is being sent. The value sent in the packets is derived from that configured value and the FPC slot value. This method ensures two IPFIX devices can never have the same value of observation domain ID. As you can see in [Table 109 on page 624](#):

- The configurable observation domain ID value is 8 bits. Therefore, the value range is 0 to 255.
- One bit is always set to 1, ensuring that the observation domain ID value sent in the packet is never 0.

Table 109: PTX Series: Format of the Observation Domain ID Value Sent in the Packet

Configured observation domain ID value (8 bits)	(15 bits set to zero)	1 bit (set to 1)	FPC slot (8 bits)
---	-----------------------	------------------	-------------------

Release History Table

Release	Description
22.4R1-EVO	Starting in Junos OS Evolved Release 22.4R1, you can configure the observation domain ID and source ID for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 routers.
22.4R1	Starting in Junos OS Release 22.4R1, you can configure the observation domain ID and source ID for the PTX1000, PTX10008, and PTX10016 routers.
17.4R1	Starting with Junos OS Release 17.4R1, IPFIX flows are supported on QFX10008 and QFX10016 switches.
17.2R1	Starting with Junos OS Release 17.2R1, IPFIX flows are supported on QFX10002 switches.

RELATED DOCUMENTATION

[Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows](#) | 625

Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows

IN THIS SECTION

- [Considerations for MX and QFX Series](#) | 627
- [Considerations for PTX Series](#) | 632

Starting with Junos OS Release 14.1, you can define the template ID for version 9 and IPFIX templates for inline flow monitoring. To specify the template ID for version 9 flows, include the `template-id id` statement at the `[edit services flow-monitoring version9 template template-name]` hierarchy level.

NOTE: Template ID is implemented differently on SRX Series devices. You cannot configure the template ID, instead you should assign the template ID dynamically.

```
[edit services flow-monitoring version9]
template template-name {
    template-id id;
}
```

To specify the template ID for version IPFIX flows, include the `template-id` statement at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level.

```
[edit services flow-monitoring version-ipfix]
template template-name {
    template-id id;
}
```

To specify the options template ID for version 9 flows, include the `options-template-id` statement at the `[edit services flow-monitoring version9 template template-name]` hierarchy level.

```
[edit services flow-monitoring version9]
template template-name {
    options-template-id id;
}
```

To specify the options template ID for IPFIX flows, include the `options-template-id` statement at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level. The template ID and options template ID can be a value in the range of 1024 through 65535. For PTX Series routers, if you do not configure the template ID or options template ID, the software assigns an ID in the default range of 256-511, and the ID is different for each template.

```
[edit services flow-monitoring version-ipfix]
template template-name {
    options-template-id id;
}
```

Considerations for MX and QFX Series

The template ID and options template ID can be a value in the range of 1024 through 65535. If you do not configure values for the template ID and options template ID, default values are assumed for these IDs, which are different for the various address families. If you configure the same template ID or options template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same template ID value for both IPv4 and IPv6, the collector validates the export data based on the template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4.

Up to Junos OS Release 13.3R1, the default values of template IDs for IPFIX flows for the different protocols or address families are:

- IPv4 IPFIX flow template ID—256
- IPv6 IPFIX flow template ID—257
- VPLS IPFIX flow template ID—258
- MPLS IPFIX flow template ID—259

Starting with Junos OS Release 14.1R1, the default values of template IDs for version 9 flows for the different protocols or address families are:

- IPv4 version 9 flow template ID—320
- IPv6 version 9 flow template ID—321
- VPLS version 9 flow template ID—322
- MPLS version 9 flow template ID—323

Up to Junos OS Release 13.3R1, the default values of options template IDs for IPFIX flows for the different protocols or address families are:

- IPv4 IPFIX flow options template ID—512
- IPv6 IPFIX flow options template ID—513
- VPLS IPFIX flow options template ID—514
- MPLS IPFIX flow options template ID—515

Starting with Junos OS Release 14.1R1, the default values of options template IDs for version 9 flows for the different protocols or address families are:

- IPv4 version 9 flow options template ID—576

- IPv6 version 9 flow options template ID—577
- VPLS version 9 flow options template ID—578
- MPLS version 9 flow options template ID—579

Table 110 on page 628 describes for the MX Series the values of data template and option template IDs for different protocols with default and configured values for IPFIX flows.

Table 110: MX Series: Values of Template and Option Template IDs for IPFIX Flows

Family	Configured Value	Data Template	Option Template
IPv4	None	256	576
IPv4	1024-65535	1024-65535	1024-65535
IPv6	None	257	577
IPv6	1024-65535	1024-65535	1024-65535
VPLS	None	258	578
VPLS	1024-65535	1024-65535	1024-65535
MPLS	None	259	579
MPLS	1024-65535	1024-65535	1024-65535

Table 111 on page 628 describes for the MX Series the values of data template and option template IDs for different protocols with default and configured values for version 9 flows.

Table 111: MX Series: Values of Template and Option Template IDs for Version 9 Flows

Family	Configured Value	Data Template	Option Template
IPv4	None	320	576

Table 111: MX Series: Values of Template and Option Template IDs for Version 9 Flows (Continued)

Family	Configured Value	Data Template	Option Template
IPv4	1024-65535	1024-65535	1024-65535
IPv6	None	321	577
IPv6	1024-65535	1024-65535	1024-65535
VPLS	None	322	578
VPLS	1024-65535	1024-65535	1024-65535
MPLS	None	323	579
MPLS	1024-65535	1024-65535	1024-65535

[Table 112 on page 629](#) describes for the MX Series the values of data template and option template IDs for different protocols with default and configured values for IPFIX flows.

Table 112: MX Series: Values of Template and Option Template IDs for IPFIX Flows

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	IPv4 (0)	1	1	0	0000 0000 0000 1000 0000 0001 0000 0001 0x00080101

Table 112: MX Series: Values of Template and Option Template IDs for IPFIX Flows *(Continued)*

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
None	IPv6 (1)	1	1	0	0000 0000 0000 1001 0000 0001 0000 0001 0x00090101
None	VPLS (2)	1	1	0	0000 0000 0000 1010 0000 0001 0000 0001 0x000A0101
None	MPLS (3)	1	1	0	0000 0000 0000 1011 0000 0001 0000 0001 0x000B0101
4	IPv4 (0)	1	1	0	0000 0100 0000 1000 0000 0001 0000 0001 0x04080101
190	IPv4 (0)	1	1	0	1101 1110 0000 1000 0000 0001 0000 0001 0xBE080101

Table 112: MX Series: Values of Template and Option Template IDs for IPFIX Flows *(Continued)*

Configured Value	Protocol Family	FPC Slot	PFE Inst	LU Inst	Observation Domain Id Conf val rsvd 1proto slot LUInst PFEInst xxxx xxxx xxxx 1xxx xxxx xxxx xxxx xxxx
4	IPv4 (0)	2	1	1	0000 0100 0000 1000 0000 0010 0001 0001 0x04080211
4	IPv6 (1)	1	1	0	0000 0100 0000 1001 0000 0001 0001 0000 0x04090110
190	IPv6 (1)	1	1	0	1101 1110 0000 1001 0000 0001 0001 0000 0xBE090110
4	VPLS (2)	2	2	0	0000 0100 0000 1010 0000 0010 0010 0000 0x040A0220
10	IPv4 (0)	28	2	1	0000 1010 0000 1000 0001 1100 0010 0001 0x0A081C21

Considerations for PTX Series

If you choose to configure the template ID and options template ID, the range is 1024 to 65520. If you do not configure these IDs, the default values that are set are in the range 256-511 and are different for each template.

You can configure the `template-id` and `option-template-id` statements for family `inet`, `inet6`, and `mpls` only.

You must not configure the same IDs for different templates (option or data template).

NOTE: The operating system does not check to ensure that you do not configure the same ID value for different templates. If you configure the same ID value, such a setting is not processed properly and might cause unexpected behavior.

The template ID or options template ID range [configured `template-id` or `options-template-id` value + 20) is reserved and you must not configure any another ID in this range. The difference between configured template IDs or options template IDs across families should be at least 20; for example, if `template-id 1056` is configured for family `inet`, you should not configure template IDs in the range of 1056 to 1075 for any other families.

For Junos OS, if you change the template ID or options template ID, all flows are inactively timed out. New flows are learned afresh.

For Junos OS Evolved, if you change the template ID or options template ID, this change does not impact the flows.

Release History Table

Release	Description
22.4R1-EVO	Starting in Junos OS Evolved Release 22.4R1, you can configure the option template ID and template ID for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 routers.
22.4R1	Starting in Junos OS Release 22.4R1, you can configure the options template ID and the template ID for the PTX1000, PTX10008, and PTX10016 routers.
17.2R1	Starting with Junos OS Release 17.2R1, IPFIX templates are supported on QFX10002 switches.

RELATED DOCUMENTATION

[Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows](#) | 620

Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers

Starting with Junos OS Release 14.2, the following attributes can be contained in IPFIX flow templates that are sent to the flow collector:

- `fragmentIdentification` (element ID 54)
- `ipv6ExtensionHeaders` (element ID 64)

A flow can receive many fragments in a given interval. For a given set of fragments of a packet, there is a unique fragment Identification. Hence, multiple such values can be received in a given interval. RFC 5102 for `fragmentIdentification` 54 does not clearly indicate which fragment identification needs to be shipped in the flow record information (first fragment observed after sending the flow record information or the last observed before shipping the flow record information). However, the last observed fragment Identification for a given flow is also transmitted to the flow collector.

Unlike in IPv4, IPv6 routers never fragment IPv6 packets. Packets exceeding the size of the maximum transmission unit of the destination link are dropped and this condition is signaled by a Packet Too Big ICMPv6 type 2 message to the originating node, similarly to the IPv4 method when the Don't Fragment (DF) bit is set.

The `fragmentIdentification` element is supported for both IPv4 and IPv6 flow templates. The `fragmentIdentification` element is added in the record template. The `fragmentIdentification` attribute is 32 bits in size for both IPv4 and IPv6. For IPv6, this field is present in fragment Extension header and Fragment Identifier is updated as 0 if there is no Fragment extension header.

Ports are a part of the key used to identify a Flow and the subsequent packets after the first fragmented packet does not have the port information. For a fragmented packet that is destined to the router, the packets that are split assume different flows (the first and the subsequent packets). Also, because the port is denoted as zeroes for fragmented packets, all the traffic destined to a particular destination from a particular source might be reported as the same flow, although no association exists between them in terms of destination ports. Fragment ID is not part of the key. Although the fragment ID attribute is unique between each source and destination, they might end up as same flows in the intermediate router.

With ports being used in the key for the flow lookup, the fragmented packets of a stream are accounted in two different flows. The first fragmented packet, which contains the port information in its packet, is part of one flow. Subsequent packets after the first fragments, which do not contain the port information, are accounted under a different flow. Because the second flow does not contain the port information to identify itself, it consolidates all the other traffic streams with same source IP and destination IP address prefixes (also includes the non-first fragmented packets sent on different ports).

Destination nodes or endpoints in IPv6 are expected to perform path MTU discovery to determine the maximum size of packets to send, and the upper-layer protocol is expected to limit the payload size.

However, if the upper-layer protocol is unable to do so, the sending host can use the Fragment extension header in order to perform end-to-end fragmentation of IPv6 packets. Any data link layer conveying IPv6 data must be capable of delivering an IP packet containing 1280 bytes without the need to invoke end-to-end fragmentation at the IP layer.

The `ipv6ExtensionHeaders` information element is a set for 32 bit fields. Each bit in this set represents one IPv6 Extension header. An extension header bit is set if that particular extension header is observed for the flow. The bit is set to 1 if any observed packet of this Flow contains the corresponding IPv6 extension header. Otherwise, if no observed packet of this Flow contained the respective IPv6 extension header, the value of the corresponding bit is 0. The `ipv6ExtensionHeaders` element is added in the record template. The number of flows that are created depends on the number of IPv6 packets that include the IPv6 extender header attribute.

To enable the inclusion of element ID, 54, `fragmentIdentification` and element ID, 64, `ipv6ExtensionHeaders` in IPFIX flow templates that are exported to the flow collector, include the `ipv6-extended-attrib` statement at the `[edit chassis fpc slot-number inline- services flow-table-size]` hierarchy level. Collection of IP4 fragmentation IDs occurs automatically without having to configure this setting explicitly.

```
[edit chassis]
fpc slot-number {
  inline-services {
    flow-table-size {
      ipv6-extended-attrib;
    }
  }
}
```

Starting in Junos OS Releases 17.3R4, 17.4R3, 18.1R4, 18.2R2, 18.3R2, and 18.4R1, the values of the IPv6 options and their functions that are contained in IPv6 packets are described in [Table 113 on page 634](#).

Table 113: Values of IPv6 Options and Extension Headers in Packets

Bit Value	IPv6 Option	Next Header Code	Description
0	DST	60	Destination option header
1	HOP	0	Hop-by-hop option header

Table 113: Values of IPv6 Options and Extension Headers in Packets *(Continued)*

Bit Value	IPv6 Option	Next Header Code	Description
2	Res	Not applicable	Reserved
3	UNK	Not applicable	Unknown layer 4 header (compressed, encrypted, not supported)
4	FRA0	44	Fragment header – first fragment
5	RH	43	Routing header
6	FRA1	44	Fragmentation header – not first fragment
7	Res	Not applicable	Reserved
8 through 11	Res	Not applicable	Reserved
12	MOB	135	IPv6 mobility (RFC3775)
13	ESP	50	Encrypted security payload
14	AH	51	Authentication header
15	PAY	108	Payload compression header
16 through 31	Res	Not applicable	Reserved

For Junos OS Releases prior to 17.3R4, 17.4R3, 18.1R4, 18.2R2, and 18.3R2, the values of the IPv6 options and their functions that are contained in IPv6 packets are described in [Table 114 on page 636](#).

Table 114: Values of IPv6 Options and Extension Headers in Packets

Bit Value	IPv6 Option	Next Header Code	Description
0	Res	Not applicable	Reserved
1	FRA1	44	Fragmentation Header
2	RH	43	Routing Header
3	FRA0	44	Fragment Header – First Fragment
4	UNK	Not applicable	Unknown Layer 4 header (compressed, encrypted, not supported)
5	Res	Not applicable	Reserved
6	HOP	0	Hop-by-hop option header
7	DST	60	Destination option header
8	PAY	108	Payload compression header
9	AH	51	Authentication header
10	ESP	50	Encrypted security payload
11 through 31	Res	Not applicable	Reserved

Release History Table

Release	Description
17.3R4	Starting in Junos OS Releases 17.3R4, 17.4R3, 18.1R4, 18.2R2, 18.3R2, and 18.4R1, the values of the IPv6 options and their functions that are contained in IPv6 packets are described in Table 113 on page 634 .

RELATED DOCUMENTATION

- [Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)
- [ipv6-extended-attrib | 1194](#)

Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers

IN THIS SECTION

- [Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers | 637](#)
- [Directing Replicated Version 9 Flow Aggregates to Multiple Servers | 638](#)

You can configure replication of the sampled flow records for use by multiple flow servers. You can use either sampling based on the Routing Engine, using cflowd version 5 or version 8, or sampling based on the services PIC, using flow aggregation version 9.

Directing Replicated Routing Engine–Based Sampling Flows to Multiple Servers

Routing Engine–based sampling supports up to eight flow servers for both cflowd version 5 and version 8 configurations. The total number of servers is limited to eight regardless of how many are configured for cflowd v5 or v8.

When you configure cflowd-based sampling, the export packets are replicated to all flow servers configured to receive them. If two servers are configured to receive v5 records, both the servers receive records for a specified flow.

NOTE: With Routing Engine–based sampling, if multiple flow servers are configured with version 8 export format, all of them must use the same aggregation type. For example, all servers receiving version 8 export can be configured for source-destination aggregation type.

The following configuration example allows replication of export packets to two flow servers.

```
forwarding-options {
  sampling {
    instance inst1 {
      input {
        rate 1;
      }
      family inet;
      output {
        flow-server 10.10.3.2 {
          port 2055;
          version 5;
          source-address 192.168.164.119;
        }
        flow-server 172.17.20.62 {
          port 2055;
          version 5;
          source-address 192.168.164.119;
        }
      }
    }
  }
}
```

Directing Replicated Version 9 Flow Aggregates to Multiple Servers

The export packets generated for a template are replicated to all the flow servers that are configured to receive information for that template. The maximum number of servers supported is eight.

This also implies that periodic updates required by version 9 (RFC 3954) are sent to each configured collector. The following updates are sent periodically as part of this requirement:

- Options data

- Template definition

The refresh period for options data and template definition is configured on a per-template basis at the [edit services flow-monitoring] hierarchy level.

The following configuration example allows replication of version 9 export packets to two flow servers.

```
forwarding-options {
  sampling {
    instance inst1 {
      input {
        rate 1;
      }
      family inet;
      output {
        flow-server 10.10.3.2 {
          port 2055;
          version9 {
            template {
              ipv4;
            }
          }
        }
        flow-server 172.17.20.62 {
          port 2055;
          version9 {
            template {
              ipv4;
            }
          }
        }
      }
      flow-inactive-timeout 30;
      flow-active-timeout 60;
      interface sp-4/0/0 {
        source-address 10.10.3.4;
      }
    }
  }
}
```

RELATED DOCUMENTATION

[Active Flow Monitoring Overview | 54](#)

[Configuring Flow Monitoring | 5](#)

[Configuring Services Interface Redundancy with Flow Monitoring | 72](#)

[Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System | 58](#)

Logging cflowd Flows on M and T Series Routers Before Export

To collect the cflowd flows in a log file before they are exported, include the `local-dump` statement at the `[edit forwarding-options sampling output flow-server hostname]` hierarchy level:

```
[edit forwarding-options sampling output flow-server hostname]
local-dump;
```

By default, the flows are collected in `/var/log/sampled`; to change the filename, include the `filename` statement at the `[edit forwarding-options sampling traceoptions]` hierarchy level. For more information about changing the filename, see ["Configuring Traffic Sampling Output" on page 420](#).

NOTE: Because the `local-dump` statement adds extra overhead, you should use it only while debugging cflowd problems, not during normal operation.

The following is an example of the flow information. The AS number exported is the origin AS number. All flows that belong under a cflowd header are dumped, followed by the header itself:

```
Jun 27 18:35:43 v5 flow entry
Jun 27 18:35:43   Src addr: 192.0.2.1
Jun 27 18:35:43   Dst addr: 198.51.100.15
Jun 27 18:35:43   Nhop addr: 198.51.100.240
Jun 27 18:35:43   Input interface: 5
Jun 27 18:35:43   Output interface: 3
Jun 27 18:35:43   Pkts in flow: 15
Jun 27 18:35:43   Bytes in flow: 600
Jun 27 18:35:43   Start time of flow: 7230
Jun 27 18:35:43   End time of flow: 7271
Jun 27 18:35:43   Src port: 26629
Jun 27 18:35:43   Dst port: 179
```

```

Jun 27 18:35:43    TCP flags: 0x10
Jun 27 18:35:43    IP proto num: 6
Jun 27 18:35:43    TOS: 0xc0
Jun 27 18:35:43    Src AS: 7018
Jun 27 18:35:43    Dst AS: 11111
Jun 27 18:35:43    Src netmask len: 16
Jun 27 18:35:43    Dst netmask len: 0

```

[... 41 more version 5 flow entries; then the following header:]

```

Jun 27 18:35:43 cflowd header:
Jun 27 18:35:43    Num-records: 42
Jun 27 18:35:43    Version: 5
Jun 27 18:35:43    low seq num: 118
Jun 27 18:35:43    Engine id: 0
Jun 27 18:35:43    Engine type: 3

```

RELATED DOCUMENTATION

[Active Flow Monitoring Overview | 54](#)

[Configuring Flow Monitoring | 5](#)

[Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers | 637](#)

[Configuring Services Interface Redundancy with Flow Monitoring | 72](#)

[Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System | 58](#)

Configuring Next-Hop Address Learning on MX Series and PTX Series Routers for Destinations Accessible Over Multiple Paths

Starting in Junos OS Release 16.1, you can enable learning of next-hop addresses to correctly report the next hop address, output SNMP, destination IP address, and destination IP mask values in the flow records when a destination is reachable through multiple paths. By default, this behavior of learning the next-hop addresses is disabled for inline active flow monitoring.

Starting in Junos OS Release 20.3R1 for the PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016 routers, if you enable learning of next-hop addresses, the packet loss priority (PLP) and the first two characters of the configured forwarding class name are reported in the IPv4 and IPv6 IPFIX flow records. The collector uses this information to derive the DSCP bits that the packet would contain

when exiting the router. The first two letters of a configured forwarding class name must be unique. For tunnel termination, 0xFF is exported in the PLP field and NULL (0) is exported in the forwarding class name field. The mapping between the PLP exported in the record and the loss priority names is as follows:

- 0x00: Low
- 0x01: Medium-low
- 0x02: Medium-high
- 0x03: High
- 0xFF: Unknown

When learning next-hop addresses is disabled, data is reported as follows:

- If the destination address of the sampled IPv4 flow is reachable through multiple paths, the IPv4 next hop address and the output SNMP address are reported in the flow records as the same as the gateway address and SNMP index of the first path seen in the forwarding table.
- If the destination address of the sampled IPv6 flow is reachable through multiple paths, the IPv4 next hop address and the output SNMP address are reported as 0 in the flow records.
- If the Incoming Interface (IIF) and Outgoing Interface (OIF) are not in the same VRF, then the destination IP address, destination IP mask, IPv4 next hop address, and the output SNMP address are reported as 0 in the flow records.
- The packet loss priority and forwarding class information is not reported for the PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016 routers.

When learning of next-hop addresses is enabled, the output SNMP, destination IP address, destination IP mask values, packet loss priority, and the first two characters of the configured forwarding class name in the flow records are reported correctly when a destination is reachable through multiple paths. To enable next-hop learning, include the `nexthop-learning enable` statement at the `[edit services flow-monitoring (version-ipfix | version9) template template-name]` hierarchy level.

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
set nexthop-learning enable;
```

Release History Table

Release	Description
20.3R1	Starting in Junos OS Release 20.3R1 for the PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016 routers, if you enable learning of next-hop addresses, the packet loss priority (PLP) and the first two characters of the configured forwarding class name are reported in the IPv4 and IPv6 IPFIX flow records. The collector uses this information to derive the DSCP bits that the packet would contain when exiting the router. The first two letters of a configured forwarding class name must be unique.
16.1	Starting in Junos OS Release 16.1, you can enable learning of next-hop addresses to correctly report the next hop address, output SNMP, destination IP address, and destination IP mask values in the flow records when a destination is reachable through multiple paths.

RELATED DOCUMENTATION

| [nexthop-learning](#) | [1258](#)

5

PART

Real-Time Performance Monitoring and Video Monitoring Services

[Monitoring Traffic Using Real-Time Performance Monitoring | 645](#)

[Managing License Server for Throughput Data Export | 724](#)

[Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking | 728](#)

[Configuring RFC 2544-Based Benchmarking Tests on ACX Series | 855](#)

[Tracking Streaming Media Traffic Using Inline Video Monitoring | 927](#)

Monitoring Traffic Using Real-Time Performance Monitoring

IN THIS CHAPTER

- Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | **646**
- Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | **651**
- Configuring RPM Receiver Servers | **662**
- Limiting the Number of Concurrent RPM Probes on M, MX, T and PTX Routers and EX Series Switches | **663**
- Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches | **663**
- Analyzing Network Efficiency in IPv6 Networks on MX Series Routers Using RPM Probes | **668**
- Configuring BGP Neighbor Discovery Through RPM | **671**
- Examples: Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers With RPM | **674**
- Trace RPM Operations | **676**
- Examples: Configuring Real-Time Performance Monitoring on MX, M, T and PTX Series Routers | **680**
- Enabling RPM on MX, M and T Series Routers and SRX Firewalls for the Services SDK | **685**
- Understand Two-Way Active Measurement Protocol | **686**
- Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches | **694**
- Example: Configuring TWAMP Client and Server on MX Series Routers | **705**
- Understanding TWAMP Auto-Restart | **713**
- Configuring TWAMP Client and TWAMP Server to Reconnect Automatically After TWAMP Server Unavailability | **716**

Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches

Real-time performance monitoring (RPM) enables you to configure active probes to track and monitor traffic. Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets. RPM provides Management Information Base (MIB) support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

When RPM is configured on a Junos device, the device calculates network performance based on packet response time, jitter, and packet loss. The device gathers RPM statistics by sending out probes to a specified probe target, identified by an IP address. When the target receives a probe, it generates responses that are received by the device.

Starting in Junos OS Release 17.3R1, you can apply RPM to IPsec tunnels and GRE tunnels for PIC-based and Routing Engine-based RPM clients and servers if you are using MS-MPCs or MS-MICs. Packet Forwarding Engine-based RPM is not supported for IPsec tunnels. Support of RPM on IPsec tunnels enables service level agreement (SLA) monitoring for traffic transported in IPsec tunnels.

NOTE: RPM is not supported on logical systems.

Starting in Junos OS Evolved Release 20.1R1, you can configure RPM probes. For Junos OS Evolved, RPM is configured at the `[edit services monitoring rpm]` hierarchy level. The scope of support is limited to:

- Probe generation and reception (client) as well as reflection (server) for the following RPM probe types:
 - icmp-ping
 - icmp-timestamp
 - udp-ping
 - udp-timestamp
- Probe history management
- Reporting through syslog only

Starting in Junos OS Evolved Release 21.2R1, reporting through SNMP MIB objects is supported for RPM. For more information about SNMP MIBs that Juniper supports, see [SNMP MIB Explorer](#).

In Junos OS, you can also configure RPM services to determine automatically whether a path exists between a host device and its configured BGP neighbors. You can view the results of the discovery using an SNMP client. Results are stored in `pingResultsTable`, `jnxPingResultsTable`, `jnxPingProbeHistoryTable`, and `pingProbeHistoryTable`.

Starting in Junos OS Release 18.4R1 for MX Series routers, you can use RPM probes to detect link status, and change the preferred-route state on the basis of the probe results. RPM-tracked routes can be IPv4 or IPv6, and support a single IPv4 or IPv6 next hop. You configure this feature with the `rpm-tracking` statement at the `[edit routing-options]` or `[edit routing-instances routing-options]` hierarchy level. For example, RPM probes can be sent to an IP address to determine if the link is up, and if so, the software installs a static route in the route table. RPM-tracked static routes are installed with preference 1 and thus are preferred over any existing static routes for the same prefix. Starting in Junos OS Release 19.1R1, you can track up to 16 next hops for each IPv4 or IPv6 RPM-tracked static route, for MX Series routers. Starting in Junos OS Release 20.4R1, we've extended support to the PTX Series routers. In addition, for this feature, you can configure route preference and tag values for each IPv4 or IPv6 destination prefix. Starting in Junos OS Release 22.3R1, you can configure RPM-tracked static routes for the ACX710 and ACX5448 routers.

In Junos OS, probe configuration and probe results are supported by both the command-line interface (CLI) and SNMP. You set the probe options in the `test test-name` statement at the `[edit services rpm probe owner]` hierarchy level. You use the `show services rpm probe-results` command to view the results of the most recent RPM probes.

The following probe types are supported with DSCP marking:

- HTTP get (not available for BGP RPM services)
- ICMP echo
- ICMP timestamp
- TCP connection
- UDP echo
- UDP timestamp

NOTE: For ACX routers:

- The ACX710 and ACX5448 Series routers support the `hardware-timestamp` statement configuration, starting in Junos OS Release 22.3R1.
- The ACX500 Series, ACX1000 Series, ACX2000 Series, ACX4000 Series, ACX5048 router, and the ACX5096 router do not support the `hardware-timestamp` statement configuration.

With probes, you can monitor:

- Average round-trip time
- *Jitter* of the round-trip time—The difference between the minimum and maximum round-trip time
- Maximum round-trip time
- Minimum round-trip time
- Standard deviation of the round-trip time (Junos OS only)

One-way measurements for ICMP timestamp probes include:

- Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times
- Number of probe responses received
- Number of probes sent
- Percentage of lost probes

You can configure the following RPM thresholds:

- Ingress/egress delay
- Jitter
- Round-trip time
- Standard deviation (Junos OS only)
- Successive lost probes
- Total lost probes (per test)

You can also configure CoS classifiers and prioritization of RPM packets over regular data packets received on an input interface with the `dscp-code-points` configuration statement.

[Table 115 on page 649](#) provides information about RPM and related timestamp support on MPC, MS-MIC/MPC, and Routing Engine:

Table 115: RPM and related timestamp support for ICMP probes

Feature	Role	IP Version	Support (Y/N)	Timestamp on Routing Engine	Timestamp on MPC (hardware-timestamp)	Timestamp on MPC (si-interface)	Timestamp on MS-MIC/MPC (delegate-probes)
RPM	Client	IPv4	Y	Y (µsec) 2000 maximum probes	Y (µsec) 2000 maximum probes	N	Y (msec) 1 million maximum probes
		IPv6	Y	Y (µsec) 2000 maximum probes	N	N	Y (msec) 1 million maximum probes
	Server	IPv4	Y	Y (µsec) 2000 maximum probes	Y (µsec) 2000 maximum probes	N	Y (msec) 1 million maximum probes
		IPv6	Y	Y (µsec) 2000 maximum probes	N	N	Y (msec) 1 million maximum probes

Release History Table

Release	Description
22.4R1-EVO	Starting in Junos OS Evolved Release 22.4R1, you can configure RPM probes for the QFX5130-CD, QFX5220, and QFX5700. We've also added reporting through MIB objects for these devices. For Junos OS Evolved, RPM is configured at the [edit services monitoring rpm] hierarchy level.
22.3R1	Starting in Junos OS Release 22.3R1, you can configure RPM probes and enable timestamps on RPM probe messages in the Packet Forwarding Engine for the ACX710 and ACX5448 Series routers.

22.3R1	Starting in Junos OS Release 22.3R1, you can configure RPM-tracked static routes for the ACX710 and ACX5448 routers, including multiple next hops and the setting of preference and tag values for each IPv4 or IPv6 destination prefix.
21.4R1	Starting in Junos OS Release 21.4R1, you can configure RPM probes and enable timestamps on RPM probe messages in the Packet Forwarding Engine for the EX9200 Series switches.
21.3R1	Starting in Junos OS Release 21.3R1, you can configure RPM probes and enable timestamps on RPM probe messages in the Packet Forwarding Engine for the QFX10002, QFX10008, and QFX10016 switches.
21.2R1-EVO	Starting in Junos OS Evolved Release 21.2R1, reporting through SNMP MIB objects is supported for RPM.
21.2R1	Starting in Junos OS Release 21.2R1, you can enable timestamps on RPM probe messages in the Packet Forwarding Engine for the PTX5000 router.
20.4R1	Starting in Junos OS Release 20.4R1, we've extended support for the RPM-tracked static routes feature to the PTX Series routers. In addition, for this feature, you can configure route preference and tag values for each IPv4 or IPv6 destination prefix.
20.1R1-EVO	Starting in Junos OS Evolved Release 20.1R1, you can configure RPM probes. For Junos OS Evolved, RPM is configured at the [edit services monitoring rpm] hierarchy level.
19.3R2	RPM is not supported when you enable Next Gen Services on an MX Series router.
19.2R1	Starting in Junos OS Release 19.2R1, you can enable timestamps on RPM probe messages in the Packet Forwarding Engine host processor for the MPC10E-15C-MRATE line card on MX240, MX480, and MX960 routers, and on the MPC11E line card on the MX2008, MX2010, and MX2020 routers.
19.1R1	Starting in Junos OS Release 19.1R1, you can track up to 16 next hops for each IPv4 or IPv6 RPM-tracked static route, for MX Series routers.
19.1R1	Starting in Junos OS Release 19.1R1, PTX Series routers support timestamping of RPM probe messages on the Packet Forwarding Engine.

18.4R1	Starting in Junos OS Release 18.4R1 for MX Series routers, you can use RPM probes to detect link status, and change the preferred-route state on the basis of the probe results. RPM-tracked routes can be IPv4 or IPv6, and support a single IPv4 or IPv6 next hop. For example, RPM probes can be sent to an IP address to determine if the link is up, and if so, the software installs a static route in the route table. RPM-tracked static routes are installed with preference 1 and thus are preferred over any existing static routes for the same prefix.
17.3R1	Starting in Junos OS Release 17.3R1, you can apply RPM to IPsec tunnels and GRE tunnels for PIC-based and Routing Engine-based RPM clients and servers if you are using MS-MPCs or MS-MICs.
12.3X51-D10	Starting in Junos OS Release 12.3X51-D10, we extended support for RPM to ACX Series routers.

RELATED DOCUMENTATION

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

[Examples: Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers With RPM | 674](#)

Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches

The probe owner and test name of an RPM probe together represent a single RPM configuration instance. When you specify the test name, you also can configure the test parameters.

To configure the probe owner, test name, and test parameters, include the probe statement at the [edit services rpm] hierarchy level:

```
[edit services rpm]
probe owner {
  delegate-probes;
  test test-name {
    data-fill data;
    data-size size;
    destination-interface interface-name;
    destination-port (RPM) port;
    dscp-code-points (RPM) dscp-bits;
    hardware-timestamp;
    history-size size;
```

```

inet6-options;
moving-average-size number;
one-way-hardware-timestamp;
probe-count count;
probe-interval seconds;
probe-type type;
routing-instance (RPM) instance-name;
rpm-scale {
    destination {
        interface interface-name.logical-unit-number;
        subunit-cnt subunit-cnt;
    }
    source {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    source-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    target {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    target-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    tests-count tests-count;
}
source-address address;
target (url url | address address);
test-interval interval;
thresholds (Junos OS) thresholds;
traps traps;
ttl [hop-count]
}
}

```

Keep the following points in mind when you configure RPM clients and RPM servers:

- RPM is not supported on logical systems.
- You cannot configure an RPM client that is PIC-based and an RPM server that is based on either the Packet Forwarding Engine or Routing Engine to receive the RPM probes.
- You cannot configure an RPM client that is Packet Forwarding Engine-based and an RPM server that receives the RPM probes to be on the PIC or Routing Engine.
- The RPM client and RPM server must be located on the same type of module. For example, if the RPM client is PIC-based, the RPM server must also be PIC-based, and if the RPM server is Packet Forwarding Engine-based, the RPM client must also be Packet Forwarding Engine-based.
- Starting in Junos OS Release 17.3R1, PIC-based and Routing Engine-based RPM is supported for IPsec tunnels and GRE tunnels if you are using MS-MPCs or MS-MICs. Packet Forwarding Engine-based RPM is not supported for IPsec tunnels. Support of RPM on IPsec tunnels enables service level agreement (SLA) monitoring for traffic transported in IPsec tunnels.
- Starting in Junos OS Release 17.3R1, you can configure the generation of IPv4 `icmp-ping` and `icmp-ping-timestamp` RPM probes on an MS-MPC or MS-MIC, which increases the number of probes generated upto 1 million per second on every service-NPU compared to the number of probes that are generated on the Packet Forwarding Engine. Starting in Junos OS Release 18.1R1, you can configure the generation of `icmp6-ping` RPM probes on an MS-MPC or MS-MIC. To configure the generation of RPM probes on an MS-MPC or MS-MIC:
 - Include the destination-interface `interface-name.logical-unit-number` at the `[edit services rpm probe owner test test-name]` hierarchy level, and include the `delegate-probes` statement at the `[edit services rpm probe owner]` hierarchy level. The `interface-name.logical-unit-number` specifies a logical interface on an MS-MPC or MS-MIC slot, PIC, and port that has a valid IP address defined on it (for example, `ms-1/2/1.1`). The interface cannot be an aggregated multiservices interface (`ams-`).
 - Include the `rpm client-delegate-probes` and the family (`inet` | `inet6`) address `address` statements at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level. The `interface-name` and the `logical-unit-number` must match the `interface-name.logical-unit-number` that you used for the destination-interface.

For RPM probes configured on an MS-MPC or MS-MIC, you cannot configure the `routing-instance` statement at the `[edit services rpm probe owner test test-name]` hierarchy level, and you cannot configure both IPv4 and IPv6 probes within the same test.

Starting in Junos OS Release 18.1R1, you can use additional filters to limit the output of the ["show services rpm probe-results"](#) on page 1775 and ["show services rpm history-results"](#) on page 1769 commands for RPM probes generated on an MS-MPC or MS-MIC.

- Starting in Junos OS Release 17.4R1, you can optimize the CLI configuration for RPM tests for IPv4. Starting in Junos OS Release 18.2R1, you can also optimize the CLI configuration for RPM tests for IPv6. This optimization allows the use of minimal RPM configuration statements to generate multiple tests (up to 100K tests) with pre-defined, reserved RPM test names. This optimization can be configured for tests with probes that are generated by either the Packet Forwarding Engine or by an MS-MPC or MS-MIC. Tests are generated for multiple combinations of source and target addresses, which are incremented based on your configuration.

The maximum number of concurrent RPM probes supported for various Junos releases are as follows:

- Junos OS release older than 17.3R1—500
- Junos OS release 17.3R1 and later—2000 for ICMP and ICMP-Timestamp probe types. For probes of other types (UDP and TCP) the limit is 500.
- Junos OS Release 17.3R1 and later (with the implementation of ["delegate-probes" on page 1041](#))—1 Million per Service-NPU.

NOTE: One MS-MIC contains one service-NPU and one MS-MPC contains four service-NPUs.

With the implementation of ["delegate-probes" on page 1041](#), the RPM probes are compliant to RFC792 and RFC4443. Hence, they can be used to monitor any IP device compliant to either RFC and are able to respond to icmp-timestamp and/or icmp6-ping packets.

Tests are first generated for all the source addresses with the initial target address, then tests are generated for all the source addresses with the next available target address, and so on. You can also configure a group that contains global values for a particular probe owner, and apply the group to the probe owner.

To generate multiple RPM tests, configure the following:

```
[edit services rpm probe owner]
apply-groups group-name;
test test-name {
  rpm-scale {
    destination {
      interface interface-name.logical-unit-number;
      subunit-cnt subunit-cnt;
    }
    source {
```

```

        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    source-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    target {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    target-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    tests-count tests-count;
}

```

The options are:

<i>ipv4-address-base</i>	The IPv4 source or target address that is incremented to generate the addresses used in the RPM tests.
<i>ipv6-address-base</i>	The IPv6 source or target address that is incremented to generate the addresses used in the RPM tests.
<i>ipv4-step</i>	The amount to increment the IPv4 source or target address for each generated RPM test.
<i>ipv6-step</i>	The amount to increment the IPv6 source or target address for each generated RPM test.
<i>ipv4-count</i>	The maximum number of IPv4 source or target addresses to use for the generated RPM tests.
<i>ipv6-count</i>	The maximum number of IPv6 source or target addresses to use for the generated RPM tests.

<i>interface-name.logical-unit-number</i>	The services interface that is generating RPM probes and the logical unit number that is used for the first test that is generated.
<i>subunit-cnt</i>	The maximum number of logical units used by the services interface in the generated tests. The first generated test uses the logical unit specified in the <i>interface-name.logical-unit-number</i> option, and each successive test increments the logical unit number by one. Once the maximum number of logical units has been used, the next generated test cycles back to the logical unit that was used in the first test.
<i>tests-count</i>	The maximum number of RPM tests to generate. This number must be less than or equal to the number of generated source addresses multiplied by the number of generated target addresses.

To configure a group with global values for a particular probe owner:

```
[edit groups group-name]
services {
  rpm {
    probe <*> {
      test {
        data-fill data;
        data-size size;
        dscp-code-points (RPM) dscp-bits;
        history-size size;
        moving-average-size number;
        probe-count count;
        probe-type type;
        test-interval interval;
        thresholds (Junos OS) thresholds;
      }
    }
  }
}
```

- To specify a probe owner, include the probe statement at the [edit services rpm] hierarchy level. The probe owner identifier can be up to 32 characters in length.
- To specify a test name, include the test statement at the [edit services rpm probe owner] hierarchy level. The test name identifier can be up to 32 characters in length. A test represents the range of probes over which the standard deviation, average, and jitter are calculated.

- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the `data-fill` statement at the `[edit services rpm probe owner]` hierarchy level. The value can be a hexadecimal value. The `data-fill` statement is not valid with the `http-get` or `http-metadata-get` probe types.
- To specify the size of the data portion of ICMP probes, include the `data-size` statement at the `[edit services rpm probe owner]` hierarchy level. The size can be from 0 through 65400 and the default size is 0. The `data-size` statement is not valid with the `http-get` or `http-metadata-get` probe types.

NOTE: If you configure the hardware timestamp feature (see ["Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches" on page 663](#)):

- This is a deprecated element `data-size` default value is 32 bytes and this is a deprecated element 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 44 bytes.
 - The `data-size` must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface.
- On M Series and T Series routers, you configure the `destination-interface` statement to enable hardware timestamping of RPM probe packets. You specify an `sp-` interface to have the AS or Multiservices PIC add the hardware timestamps; for more information, see ["Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches" on page 663](#). You can also include the `one-way-hardware-timestamp` statement to enable one-way delay and jitter measurements.
 - To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the `destination-port` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. The `destination-port` statement is used only for the UDP and TCP probe types. The value can be 7 or from 49160 through 65535.

When you configure either `probe-type udp-ping` or `probe-type udp-ping-timestamp` along with hardware timestamping, the value for the `destination-port` can be only 7. A constraint check prevents you from configuring any other value for the destination port in this case. This constraint does not apply when you are using one-way hardware timestamping.

- To specify the value of the Differentiated Services (DiffServ) field within the IP header, include the `dscp-code-point` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. The DiffServ code point (DSCP) bits value can be set to a valid 6-bit pattern; for example, 001111. It also can be set using an alias configured at the `[edit class-of-service code-point-aliases dscp]` hierarchy level. The default is 000000.
- To specify the number of stored history entries, include the `history-size` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. Specify a value from 0 to 512. The default is 50.

- To specify a number of samples for making statistical calculations, include the `moving-average-size` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. Specify a value from 0 through 255.
- To specify the number of probes within a test, include the `probe-count` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. Specify a value from 1 through 15.
- To specify the time to wait between sending packets, include the `probe-interval` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. Specify a value from 1 through 255 seconds.
- To specify the packet and protocol contents of the probe, include the `probe-type` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. The following probe types are supported:
 - `http-get`—Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.
 - `http-metadata-get`—Sends an HTTP get request for metadata to a target URL.
 - `icmp-ping`—Sends ICMP echo requests to a target address.
 - `icmp-ping-timestamp`—Sends ICMP timestamp requests to a target address.
 - `tcp-ping`—Sends TCP packets to a target.
 - `udp-ping`—Sends UDP packets to a target.
 - `udp-ping-timestamp`—Sends UDP timestamp requests to a target address.

The following probe types support hardware timestamping of probe packets: `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, `udp-ping-timestamp`. Starting in Junos OS Release 17.3R3, the delegate probes are distributed evenly across the interval of 3 seconds to avoid the packet bursts in the network due to real-time performance monitoring (RPM). RPM syslogs are processed with the increase in the ramp up time of RPM delegates tests to 60 seconds. With RPM syslogs processed, the chances of multiple tests starting and ending at the same time are smaller, thus a potential restriction in event-processing.

NOTE: Some probe types require additional parameters to be configured. For example, when you specify the `tcp-ping` or `udp-ping` option, you must configure the destination port using the `destination-port` statement. The `udp-ping-timestamp` option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

When you configure either `probe-type udp-ping` or `probe-type udp-ping-timestamp` along with the `one-way-hardware-timestamp` command, the value for the `destination-port` can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.

- To specify the routing instance used by ICMP probes, include the `routing-instance` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. The default routing instance is Internet routing table `inet.0`.
- To specify the source IP address used for ICMP probes, include the `source-address` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. If the source IP address is not one of the router's assigned addresses, the packet uses the outgoing interface's address as its source.
- Starting in Junos OS Release 16.1R1, to specify the source IPv6 address to be used for RPM probes that are sent from the RPM client (the device that originates the RPM packets) to the RPM server (the device that receives the RPM probes), include the `inet6-options source-address ipv6-address` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. If the source IPv6 address is not one of the router's or switch's assigned addresses, the packet use the outgoing interface's address as its source.
- To specify the destination address used for the probes, include the `target` statement at the `[edit services rpm probe owner test test-name]` hierarchy level.
 - For HTTP probe types, specify a fully formed URL that includes `http://` in the URL address.
 - For all other probe types, specify an IP version 4 (IPv4) or IP version 6 (IPv6) (IPv6 support starts in Junos OS release 16.1R1) address for the target host.
- To specify the time to wait between tests, include the `test-interval` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. Specify a value from 0 through 86400 seconds. A value of 0 seconds causes the RPM test to stop after one iteration. The default value is 1.
- To specify thresholds used for the probes, include the `thresholds` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded. The following options are supported:
 - `egress-time`—Measures maximum source-to-destination time per probe.
 - `ingress-time`—Measures maximum destination-to-source time per probe.
 - `jitter-egress`—Measures maximum source-to-destination jitter per test.
 - `jitter-ingress`—Measures maximum destination-to-source jitter per test.
 - `jitter-rtt`—Measures maximum jitter per test, from 0 through 60000000 microseconds.
 - `rtt`—Measures maximum round-trip time per probe, in microseconds.
 - `std-dev-egress`—Measures maximum source-to-destination standard deviation per test.
 - `std-dev-ingress`—Measures maximum destination-to-source standard deviation per test.

- `std-dev-rtt`—Measures maximum standard deviation per test, in microseconds.
- `successive-loss`—Measures successive probe loss count, indicating probe failure.
- `total-loss`—Measures total probe loss count indicating test failure, from 0 through 15.
- Traps are sent if the configured threshold is met or exceeded. To set the trap bit to generate traps, include the traps statement at the `[edit services rpm probe owner test test-name]` hierarchy level. The following options are supported:
 - `egress-jitter-exceeded`—Generates traps when the jitter in egress time threshold is met or exceeded.
 - `egress-std-dev-exceeded`—Generates traps when the egress time standard deviation threshold is met or exceeded.
 - `egress-time-exceeded`—Generates traps when the maximum egress time threshold is met or exceeded.
 - `ingress-jitter-exceeded`—Generates traps when the jitter in ingress time threshold is met or exceeded.
 - `ingress-std-dev-exceeded`—Generates traps when the ingress time standard deviation threshold is met or exceeded.
 - `ingress-time-exceeded`—Generates traps when the maximum ingress time threshold is met or exceeded.
 - `jitter-exceeded`—Generates traps when the jitter in round-trip time threshold is met or exceeded.
 - `probe-failure`—Generates traps for successive probe loss thresholds crossed.
 - `rtt-exceeded`—Generates traps when the maximum round-trip time threshold is met or exceeded.
 - `std-dev-exceeded`—Generates traps when the round-trip time standard deviation threshold is met or exceeded.
 - `test-completion`—Generates traps when a test is completed.
 - `test-failure`—Generates traps when the total probe loss threshold is met or exceeded.

Release History Table

Release	Description
18.2R1	Starting in Junos OS Release 18.2R1, you can also optimize the CLI configuration for RPM tests for IPv6.
18.1R1	Starting in Junos OS Release 18.1R1, you can configure the generation of icmp6-ping RPM probes on an MS-MPC or MS-MIC.

18.1R1	Starting in Junos OS Release 18.1R1, you can use additional filters to limit the output of the "show services rpm probe-results" on page 1775 and "show services rpm history-results" on page 1769 commands for RPM probes generated on an MS-MPC or MS-MIC.
17.4R1	Starting in Junos OS Release 17.4R1, you can optimize the CLI configuration for RPM tests for IPv4.
17.3R3	Starting in Junos OS Release 17.3R3, the delegate probes are distributed evenly across the interval of 3 seconds to avoid the packet bursts in the network due to real-time performance monitoring (RPM). RPM syslogs are processed with the increase in the ramp up time of RPM delegates tests to 60 seconds. With RPM syslogs processed, the chances of multiple tests starting and ending at the same time are smaller, thus a potential restriction in event-processing.
17.3R1	Starting in Junos OS Release 17.3R1, PIC-based and Routing Engine-based RPM is supported for IPsec tunnels and GRE tunnels if you are using MS-MPCs or MS-MICs.
17.3R1	Starting in Junos OS Release 17.3R1, you can configure the generation of IPv4 icmp-ping and icmp-ping-timestamp RPM probes on an MS-MPC or MS-MIC, which increases the number of probes generated upto 1 million per second on every service-NPU compared to the number of probes that are generated on the Packet Forwarding Engine.
16.1	Starting in Junos OS Release 16.1R1, to specify the source IPv6 address to be used for RPM probes that are sent from the RPM client (the device that originates the RPM packets) to the RPM server (the device that receives the RPM probes), include the inet6-options source-address <i>ipv6-address</i> statement at the [edit services rpm probe <i>owner</i> test <i>test-name</i>] hierarchy level.
16.1	For all other probe types, specify an IP version 4 (IPv4) or IP version 6 (IPv6) (IPv6 support starts in Junos OS release 16.1R1) address for the target host.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Examples: Configuring Real-Time Performance Monitoring on MX, M, T and PTX Series Routers | 680](#)

Configuring RPM Receiver Servers

The RPM TCP and UDP probes are proprietary to Juniper Networks and require a receiver to receive the probes. To configure a server to receive the probes, include the `probe-server` statement at the `[edit services rpm]` hierarchy level:

```
[edit services rpm]
probe-server {
  tcp {
    destination-interface interface-name;
    port number;
  }
  udp {
    port number;
  }
}
```

The port number specified for the UDP and TCP server can be 7 or from 49160 through 65535.

NOTE: The `destination-interface` statement is not supported on PTX Series Packet Transport routers.

When you configure either `probe-type udp-ping` or `probe-type udp-ping-timestamp` along with the `one-way-hardware-timestamp` command, the value for the destination-port can be only 7. A constraint check prevents you for configuring any other value for the destination port in this case.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Examples: Configuring Real-Time Performance Monitoring on MX, M, T and PTX Series Routers | 680](#)

Limiting the Number of Concurrent RPM Probes on M, MX, T and PTX Routers and EX Series Switches

To configure the maximum number of concurrent probes allowed, include the `probe-limit` statement at the `[edit services rpm]` hierarchy level:

```
probe-limit limit;
```

Specify a limit from 1 through 500. The default maximum number is 100.

Starting in Junos OS Release 17.2R2 and 17.3R1 for MX Series routers only, the `probe-limit` is 1 through 2000.

Release History Table

Release	Description
17.2R2	Starting in Junos OS Release 17.2R2 and 17.3R1 for MX Series routers only, the <code>probe-limit</code> is 1 through 2000.

RELATED DOCUMENTATION

- Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646
- Examples: Configuring Real-Time Performance Monitoring on MX, M, T and PTX Series Routers | 680

Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You can timestamp the following RPM probe types: `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp`.

On M Series and T Series routers with an MS-PIC, on MX Series routers with an MS-DPC, MS-MIC, or MS-MPC linecard, on MX10000 Series routers, on PTX10008 and PTX10016 routers, and on EX Series switches, you can enable hardware timestamping of RPM probe messages. The timestamp is applied on both the RPM client device (the router or switch that originates the RPM probes) and the RPM probe server and applies only to IPv4 traffic. It is supported on the following:

- Layer 2 service package on MS-PICs, MS-DPCs, MS-MPCs, and MS-MICs.
- Layer 3 service package on MS-PICs, MS-DPCs, MS-MPCs, and MS-MICs.
- Extension-provider services package on M Series, MX Series, and T Series services PICs that support the Extension-Provider packages (In Junos OS releases earlier than Release 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.)
- Layer 2, Layer 3, SDK Services, and PFE RPM timestamping interoperate with each other. Here, the RPM client can be on the Layer 3 sp- interface and the RPM server can be on an SDK Services package.

Two-way timestamping is available on sp- and ms- interfaces. To configure two-way timestamping on M Series and T Series routers, include the `destination-interface` statement at the `[edit services rpm probe probe-owner test test-name]` hierarchy level:

```
destination-interface sp-fpc/pic/port.logical-unit
destination-interface ms-fpc/pic/port.logical-unit
```

Specify the RPM client router and the RPM server router on the services logical interface or the multiservices interface by including the `rpm` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level:

```
rpm (client | server);
```

The logical interface must be dedicated to the RPM task. It requires configuration of the `family inet` statement and a /32 address, as shown in the example. This configuration is also needed for other services such as NAT and stateful firewall. You cannot configure RPM service on unit 0 because RPM requires a dedicated logical interface; the same unit cannot support both RPM and other services. Because active flow monitoring requires unit 0, but RPM can function on any logical interface, a constraint check prevents you from committing an RPM configuration there.

On MX Series routers, on M320 Series routers using the Enhanced Queuing MPC, and on EX Series switches, you include the `hardware-timestamp` statement at the `[edit services rpm probe probe-name test test-name]` hierarchy level to specify that the probes are to be timestamped in the Packet Forwarding Engine host processor:

On MX Series routers, on MX10000 Series routers, on PTX5000, PTX10008, and PTX10016 routers, and on EX Series switches, you can include the `hardware-timestamp` statement at the `[edit services rpm probe probe-name test test-name]` hierarchy level to specify that the probes are to be timestamped in the Packet Forwarding Engine host processor. On MX Series routers, hardware timestamping is supported on the following line cards:

- DPC
- DPCE
- MPC1
- MPC2
- MPC3
- MPC4
- MPC5
- MPC6
- MPC7

```
hardware-timestamp;
```

On the client side, these probes are timestamped in the Packet Forwarding Engine host processor on the egress DPC on the MX Series or M320 Series router or EX Series switch originating the RPM probes (RPM client). On the responder side (RPM server), the RPM probes to be timestamped are handled by the Packet Forwarding Engine host processor, which generates the response instead of the RPM process. The RPM probes are timestamped only on the router that originates them (RPM client). As a result, only round-trip time is measured for these probes.

When using the `hardware-timestamp` statement, the `data-size` value for the probe must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface (see ["Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches" on page 651](#)). If hardware timestamping of RPM probe messages is enabled, the maximum data size that you can configure by using the `data-size` statement is limited to 1400.

NOTE: The Packet Forwarding Engine-based RPM feature does not support any stateful firewall configurations. If you need to combine RPM timestamping with a stateful firewall, use the interface-based RPM timestamping service described earlier in this section. MS-DPCs support stateful firewall processing as well as RPM timestamping.

To configure one-way timestamping, you must also include the `one-way-hardware-timestamp` statement at the `[edit services rpm probe probe-owner test test-name]` hierarchy level:

```
one-way-hardware-timestamp;
```

NOTE: If you configure RPM probes for a services interface (sp-), you need to announce local routes in a specific way for the following routing protocols:

- For OSPF, you can announce the local route by including the services interface in the OSPF area. To configure this setting, include the interface `sp-fpc/pic/port` statement at the `[edit protocols ospf area area-number]` hierarchy level.
- For BGP and IS-IS, you must export interface routes and create a policy that accepts the services interface local route. To export interface routes, include the `point-to-point` and `lan` statements at the `[edit routing-options interface-routes family inet export]` hierarchy level. To configure an export policy that accepts the services interface local route, include the `protocol local`, `rib inet.0`, and `route-filter sp-interface-ip-address/32 exact` statements at the `[edit policy-options policy-statement policy-name term term-name from]` hierarchy level and the `accept` action at the `[edit policy-options policy-statement policy-name term term-name then]` hierarchy level. For the export policy to take effect, apply the policy to BGP or IS-IS with the `export policy-name` statement at the `[edit protocols protocol-name]` hierarchy level.

For more information about these configurations, see the [Routing Policies, Firewall Filters, and Traffic Policers User Guide](#).

Routing the probe packets through the multiservices card also enables you to filter the probe packets to particular queues. The following example shows the RPM configuration and the filter that specifies queuing:

```
services rpm {
  probe p1 {
    test t1 {
      probe-type icmp-ping;
      target address 10.8.4.1;
      probe-count 10;
      probe-interval 10;
      test-interval 10;
      dscp-code-points af11;
      data-size 100;
      destination-interface sp-1/2/0.0;
    }
  }
}
firewall {
  filter f1 {
    term t1 {
```


Analyzing Network Efficiency in IPv6 Networks on MX Series Routers Using RPM Probes

IN THIS SECTION

- [Guidelines for Configuring RPM Probes for IPv6 Destinations | 669](#)

Real-time performance monitoring (RPM) is a mechanism that enables you to monitor network performance in real time and to assess and analyze network efficiency. Typically, network performance is assessed in real time based on the jitter, delay, and packet loss experienced on the network. RPM is a service available in Junos OS that enables a router to measure metrics such as round-trip delays and unanswered echo requests. To compute these parameters, RPM exchanges a set of probes with other IP hosts in the network for monitoring and network tracking purposes. These probes are sent from a source node to other destination devices in the network that require tracking. Data such as transit delay and jitter can be collected from these probes, and this data can be used to provide an approximation of the delay and jitter experienced by live traffic in the network. Different live traffic metrics such as round-trip time (RTT), positive egress jitter, negative egress jitter, positive ingress jitter, negative ingress jitter, positive round-trip jitter, and negative round-trip jitter can be obtained from the results of the RPM test. RPM calculates minimum, maximum, average, peak-to-peak, standard deviation, and sum calculations for each of these measurements. RPM probes can also be used to verify the path between BGP neighbors.

Starting with Junos OS release 16.1, the RPM client router (the router or switch that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address. To specify the destination IPv6 address used for the probes, include the `target (url ipv6-url | address ipv6-address)` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. The protocol family for IPv6 is named `inet6`.

```
[edit services rpm]
probe owner {
  test test-name {
    target (url ipv6-url | address ipv6-address);
  }
}
```

To specify the IPv6 protocol-related settings and the source IPv6 address of the client from which the RPM probes are sent, include the `inet6-options source-address ipv6-address` statement at the `[edit services rpm probe owner test test-name]` hierarchy level. A probe request is a standard packet with corresponding TCP, UDP, and ICMP headers over the IPv6 header. No RPM header is appended to the standard packet

for Routing Engine-based RPM implementation. A probe response is also a standard packet with corresponding TCP, UDP, and ICMP headers over the IPv6 header. No RPM header is appended to the standard packet for Routing Engine-based RPM implementation.

```
[edit services rpm]
probe owner {
  test test-name {
    inet6-options source-address ipv6-address;
  }
}
```

The output of the `show services rpm probe-results owner probe-name test test-name` and `show services rpm history-results owner owner test name` commands that display the results of the most recent RPM probes and results of historical RPM probes respectively have been enhanced to display the target address as IPv6 address and other IPv6 information for probes sent to IPv6 servers or destinations. The existing SNMP Get requests and traps for IPv6 are applicable for IPv6 probes. The target type field in the SNMP set operation contains IPv6 source and destination addresses.

Guidelines for Configuring RPM Probes for IPv6 Destinations

Keep the following points in mind when you configure IPv6 addresses for RPM destinations or servers:

- Only Routing Engine-based RPM is supported for IPv6 targets including VRF support, specification of the size of the data portion of ICMP probes, data pattern, and traffic class.
- You can configure probes with a combination of IPv4 and IPv6 tests. However, a test can be either IPv4 or IPv6-based at a point in time. The OS impacts the accuracy of the measurements because the variability factor introduced by the general OS that performs the system processing proved is significantly larger than the amount of time spent by the packet traversing on the wire. This condition causes round-trip time (RTT) spikes to be seen even with a single test.
- Routing Engine-based RPM does not support one-way hardware-based timestamping.
- One-way measurements are not supported here because timestamping is done only on the RPM client side.
- The maximum number of concurrent probes allowed (by including the `probe-limit` statement at the `[edit services rpm]` hierarchy level) is 1000. We recommend that the limit on concurrent probes be set as 10. Higher concurrent probes can result in higher spikes. The maximum number of tests you can configure is 1000. RPM cannot be configured on logical systems. SNMP set operation is permitted only on ICMP probes and it is not supported for other type of probes.
- The `hardware-timestamp` and `one-way-hardware-timestamp` statements at the `[edit services rpm probe owner test test-name]` hierarchy level are not supported for IPv6.

- You cannot specify the `icmp-ping` (which sends ICMP echo requests to a target address) and the `icmp-ping-timestamp` (which sends ICMP timestamp requests to a target address) options with the `probe-type` statement at the `[edit services rpm probe owner test test-name]` hierarchy level.
- Some of the RPM problems can be resolved by restarting the SNMP remote operations process (`rmopd`) on the Routing Engine by using the `restart remote-operations` command. If RPM needs to be disabled, the `rpm` statement at the `[edit services]` hierarchy level needs to be deleted or deactivated. PIC, Packet Forwarding Engine, and lookup chip (LU) based RPM implementation for IPv6 are not supported.
- The following table describes the IPv6 special address prefixes that are not supported.

IPv6 Address Type	IPv6 Address Prefix
Node-Scoped Unicast	<p>::1/128 is the loopback address</p> <p>::/128 is the unspecified address</p>
IPv4-Mapped Addresses	::FFFF:0:0/96
IPv4-Compatible Addresses	:<ipv4-address>/96
Link-Scoped Unicast	fe80::/10
Unique-Local	fc00::/7
Documentation Prefix	2001:db8::/32
6to4	2002::/16
6bone	5f00::/8
ORCHID	2001:10::/28
Teredo	2001::/32
Default Route	::/0

(Continued)

IPv6 Address Type	IPv6 Address Prefix
Multicast	ff00::/8

- The current scaling number for IPv4 probes is a maximum of 500 concurrent probes and the limit on the maximum number of configurable tests is 1000. These scaling parameters are applicable for IPv6 probes. The same scaling limits are applicable, even in cases where both IPv4-based tests and IPv6-based tests are run at the same time.
- The minimum rate of probes is 1 probe per second and the maximum interval between tests is 86400 seconds. These scaling and performance numbers vary based on whether the Two-Way Active Measurement Protocol (TWAMP) server and client are configured on the same router. This condition occurs because the TWAMP server/client has packet processing in RMOPD and it competes with RPM functionality in the same process. The RTT of IPv6-based RPM and ping utilities must be equivalent for data size. In Routing Engine-based RPM implementation, RTT spikes are seen owing to various queuing delays introduced in the system. This behavior can be noticed even with a single test.
- Some of the TCP and UDP ports might be opened to communicate between the RPM server and RPM client. Therefore, we recommend that you use firewalls and distributed denial-of-service (DDoS) attack filters to ensure that no security threats are possible by some third-party attackers or hackers.
- The different packet types that can be used within the probe include:
 - ICMP6 echo
 - UDP echo
 - UDP timestamp

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | 651

Configuring BGP Neighbor Discovery Through RPM

BGP neighbors can be configured at the following hierarchy levels:

- [edit protocols bgp group *group-name*]—Default logical system and default routing instance.

- [edit routing-instances *instance-name* protocols bgp group *group-name*]*—Default logical system with a specified routing instance.*
- [edit logical-systems *logical-system-name* protocols bgp group *group-name*]*—Configured logical system and default routing instance.*
- [edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name*]*—Configured logical system with a specified routing instance.*

When you configure BGP neighbor discovery through RPM, if you do not specify a logical system, the RPM probe applies to configured BGP neighbors for all logical systems. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. You can explicitly configure RPM probes to apply only to the default logical system, the default routing instance, or to a particular logical system or routing instance.

To configure BGP neighbor discovery through RPM, configure the probe properties at the [edit services rpm bgp] hierarchy:

```
data-fill data;
data-size size;
destination-port port;
history-size size;
logical-system logical-system-name [routing-instances routing-instance-name];
moving-average-size number;
probe-count count;
probe-interval seconds;
probe-type type;
routing-instances instance-name;
test-interval interval;
```

- To specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes, include the data-fill statement at the [edit services rpm bgp] hierarchy level. The value can be a hexadecimal value.
- To specify the size of the data portion of ICMP probes, include the data-size statement at the [edit services rpm bgp] hierarchy level. The size can be from 0 through 65400 and the default size is 0.
- To specify the User Datagram Protocol (UDP) port or Transmission Control Protocol (TCP) port to which the probe is sent, include the destination-port statement at the [edit services rpm bgp] hierarchy level. The destination-port statement is used only for the UDP and TCP probe types. The value can be 7 or from 49160 through 65535.
- To specify the number of stored history entries, include the history-size statement at the [edit services rpm bgp] hierarchy level. Specify a value from 0 to 512. The default is 50.

- To specify the logical system used by ICMP probes, include the `logical-system logical-system-name` statement at the `[edit services rpm bgp]` hierarchy level. If you do not specify a logical system, the RPM probe applies to configured BGP neighbors for all logical systems. To apply the probe to only the default logical system, you must set the value of *logical-system-name* to null.
- To specify a number of samples for making statistical calculations, include the `moving-average-size` statement at the `[edit services rpm bgp]` hierarchy level. Specify a value from 0 through 255.
- To specify the number of probes within a test, include the `probe-count` statement at the `[edit services rpm bgp]` hierarchy level. Specify a value from 1 through 15.
- To specify the time to wait between sending packets, include the `probe-interval` statement at the `[edit services rpm bgp]` hierarchy level. Specify a value from 1 through 255 seconds.
- To specify the packet and protocol contents of the probe, include the `probe-type` statement at the `[edit services rpm bgp]` hierarchy level. The following probe types are supported:
 - `icmp-ping`—Sends ICMP echo requests to a target address.
 - `icmp-ping-timestamp`—Sends ICMP timestamp requests to a target address.
 - `tcp-ping`—Sends TCP packets to a target.
 - `udp-ping`—Sends UDP packets to a target.
 - `udp-ping-timestamp`—Sends UDP timestamp requests to a target address.

NOTE: Some probe types require additional parameters to be configured. For example, when you specify the `tcp-ping` or `udp-ping` option, you must configure the destination port using the `destination-port port` statement. The `udp-ping-timestamp` option requires a minimum data size of 12; any smaller data size results in a commit error. The minimum data size for TCP probe packets is 1.

- To specify the routing instance used by ICMP probes, include the `routing-instances` statement at the `[edit services rpm bgp]` hierarchy level. The default routing instance is Internet routing table `inet.0`. If you do not specify a routing instance, the RPM probe applies to configured BGP neighbors in all routing instances. To apply the RPM probe to only the default routing instance, you must explicitly set the value of *instance-name* to default.
- To specify the time to wait between tests, include the `test-interval` statement at the `[edit services bgp probe]` hierarchy level. Specify a value from 0 through 86400 seconds. A value of 0 seconds causes the RPM test to stop after one iteration. The default value is 1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Examples: Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers With RPM | 674](#)

Examples: Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers With RPM

Configure BGP neighbor discovery with RPM for all logical systems and all routing instances:

```
[edit services rpm]
bgp {
    probe-type icmp-ping;
    probe-count 5;
    probe-interval 1;
    test-interval 60;
    history-size 10;
    data-size 255;
    data-fill 0123456789;
}
```

Configure BGP neighbor discovery with RPM for only the following logical systems and routing instances: LS1/RI1, LS1/RI2, LS2, and RI3:

```
[edit services rpm]
bgp {
    probe-type icmp-ping;
    probe-count 5;
    probe-interval 1;
    test-interval 60;
    history-size 10;
    data-size 255;
    data-fill 0123456789;
    logical-system {
        LS1 {
            routing-instances {
                RI1;
```

```

        RI2;
    }
}
LS2;
}
routing-instance {
    RI3;
}
}

```

NOTE: The `logical-system` statement is not supported on PTX Series Packet Transport routers.

Configure BGP neighbor discovery with RPM for only the default logical system and default routing instance:

```

[edit services rpm]
bgp {
    probe-type icmp-ping;
    probe-count 5;
    probe-interval 1;
    test-interval 60;
    history-size 10;
    data-size 255;
    data-fill 0123456789;
    logical-system {
        null {
            routing-instances {
                default;
            }
        }
    }
}
}

```

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

Trace RPM Operations

IN THIS SECTION

- [RPM Trace Operations Overview | 676](#)
- [Configure the Trace Operations | 677](#)
- [Configure the RPM Log File Name | 678](#)
- [Configure the Number and Size of RPM Log Files | 678](#)
- [Configure Access to the Log File | 679](#)
- [Configure a Regular Expression for Lines to Be Logged | 679](#)

RPM tracing operations track all RPM operations and record them in a log file. The logged error descriptions provide detailed information to help you solve problems faster.

RPM Trace Operations Overview

In Junos OS, you enable tracing operations by configuring the `traceoptions` statement at the specific hierarchy level you want to trace. Junos OS Evolved uses a different tracing architecture. All running applications create trace information, with multiple instances of the same application having their own trace information. Therefore, in Junos OS Evolved, trace messages are logged, viewed, and configured by application. As a result, Junos OS Evolved does not support the `traceoptions` statement at many of the hierarchy levels that Junos OS supports.

In Junos OS Evolved, you do not view trace files directly, and you should never add, edit, or remove trace files under the `/var/log/traces` directory because this can corrupt the traces. Instead, you use the `show trace application application-name node node-name` command to read and decode trace messages stored in the trace files. All running applications on Junos OS Evolved create trace information at the `info` level by default.

In Junos OS, by default, no events are traced. You can change this default behavior by using the `traceoptions` statement. If you include the `traceoptions` statement at the `[edit services rpm]` hierarchy level, the default tracing behavior is the following:

- Important events are logged in a file called **rmopd** located in the `/var/log` directory.
- When the log file reaches 128 kilobytes (KB), it is renamed **rmopd.0**, then **rmopd.1**, and so on, until there are three trace files. Then the oldest trace file (**rmopd.2**) is overwritten.
- Log files can be accessed only by the user who configures the tracing operation.

RPM is governed by the `rmopd` application. For Junos OS Evolved, to configure traces for a severity other than `info` for the `rmopd` application, include the application `rmopd` node *node-name* level *severity* statement at the `[edit system trace]` hierarchy level.

NOTE: For general monitoring and troubleshooting of devices running Junos OS or Junos OS Evolved, we recommend using standard tools such as CLI `show` commands, system log messages, SNMP, and telemetry data. You should avoid using trace messages for general debugging purposes and long-term solutions because they are subject to change without notice.

Configure the Trace Operations

By default, for Junos OS, if the `traceoptions` configuration is present, only important events are logged. You can configure the trace operations to be logged by including the following statements at the `[edit services rpm traceoptions]` hierarchy level:

```
flag {
  all;
  configuration;
  error;
  ipc;
  ppm;
  rpd;
  statistics
}
```

[Table 116 on page 677](#) describes the meaning of the RPM tracing flags.

Table 116: Junos OS RPM Tracing Flags

Flag	Description	Default Setting
all	Trace all operations.	Off
configuration	Trace configuration events.	Off
error	Trace events related to catastrophic errors in daemon.	Off
ipc	Trace IPC events.	Off

Table 116: Junos OS RPM Tracing Flags (Continued)

Flag	Description	Default Setting
ppm	Trace ppm events.	Off
rpd	Trace rpd events.	Off
statistics	Trace statistics.	Off

By default, for Junos OS Evolved, all running applications create trace information at the `info` level. To configure traces for a severity other than `info` for the `rmopd` application, include the application `rmopd` node `node-name` level `severity` statement at the `[edit system trace]` hierarchy level. For information about the various configurable severity levels for Junos OS Evolved, see *trace*.

SEE ALSO

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#) | 646

Configure the RPM Log File Name

(Junos OS only) By default, the name of the file that records RPM trace output is **rmopd**. To specify a different file name:

```
[edit services rpm traceoptions]
user@host set file filename
```

Configure the Number and Size of RPM Log Files

(Junos OS only) To configure the limits on the number and size of RPM trace files:

```
[edit services rpm traceoptions]
user@host set file filename files number size size
```

The number of files can be from 2 through 1000 files. The file size of each file can be from 10 KB through 1 gigabyte (GB).

For example, set the maximum file size to 2 MB, and the maximum number of files to 20 for a log file named **rpmtrace**:

```
[edit services rpm traceoptions]
user@host set file rpmtrace files 20 size 2MB
```

When the **rpmtrace** file reaches 2 MB, it is renamed **rpmtrace.0**, and a new file called **rpmtrace** is created. When the new **rpmtrace** reaches 2 MB, **rpmtrace.0** is renamed **rpmtrace.1** and **rpmtrace** is renamed **rpmtrace.0**. This process repeats until there are 20 trace files. Then the oldest file (**rpmtrace.19**) is overwritten by **rpmtrace.18**.

Configure Access to the Log File

(Junos OS only) By default, log files can be accessed only by the user who configures the tracing operation.

To specify that any user can read all log files:

```
[edit services rpm traceoptions]
user@host set file filename world-readable
```

To explicitly set the default behavior:

```
[edit services rpm traceoptions]
user@host set file filename no-world-readable
```

Configure a Regular Expression for Lines to Be Logged

(Junos OS only) By default, the trace operation output includes all lines relevant to the logged events.

To refine the output by specifying a regular expression (regex) to be matched:

```
[edit services rpm traceoptions]
user@host set file filename match regular-expression
```

Examples: Configuring Real-Time Performance Monitoring on MX, M, T and PTX Series Routers

Configure an RPM instance identified by the probe name `probe1` and the test name `test1`:

```
[edit services rpm]
probe probe1{
  test test1 {
    dscp-code-points 001111;
    probe-interval 1;
    probe-type icmp-ping;
    target address 172.17.20.182;
    test-interval 20;
    thresholds rtt 10;
    traps rtt-exceeded;
  }
}
probe-server {
  tcp {
    destination-interface lt-0/0/0.0
    port 50000;
  }
  udp {
    destination-interface lt-0/0/0.0
    port 50001;
  }
}
probe-limit 200;
```

Configure packet classification, using `lt-` interfaces to send the probe packets to a logical tunnel input interface. By sending the packet to the logical tunnel interface, you can configure regular and multifield classifiers, firewall filters, and header rewriting for the probe packets. To use the existing tunnel framework, the `dlci` and `encapsulation` statements must be configured.

```
[edit services rpm]
probe p1 {
  test t1 {
    probe-type icmp-ping;
    target address 10.8.4.1;
    probe-count 10;
```

```

        probe-interval 10;
        test-interval 10;
        source-address 10.8.4.2;
        dscp-code-points ef;
        data-size 100;
        destination-interface lt-0/0/0.0;
    }
}
[edit interfaces]
lt-0/0/0 {
    unit 0 {
        encapsulation frame-relay;
        dlci 10;
        peer-unit 1;
        family inet;
    }
    unit 1 {
        encapsulation frame-relay;
        dlci 10;
        peer-unit 0;
        family inet;
    }
}
[edit class-of-service]
interfaces {
    lt-0/0/0 {
        unit 1 {
            classifiers {
                dscp default;
            }
        }
    }
}
}

```

Configure an input filter on the interface on which the RPM probes are received. This filter enables prioritization of the received RPM packets, separating them from the regular data packets received on the same interface.

```

[edit firewall]
filter recos {
    term recos {
        from {

```

```

        source-address {
            10.8.4.1/32;
        }
        destination-address {
            10.8.4.2/32;
        }
    }
    then {
        loss-priority high;
        forwarding-class network-control;
    }
}
}
[edit interfaces]
fe-5/0/0 {
    unit 0 {
        family inet {
            filter {
                input recos;
            }
            address 10.8.4.2/24;
        }
    }
}

```

Configure an RPM instance and enable RPM for the extension-provider packages on the adaptive services interface:

```

[edit services rpm]
probe probe1{
    test test1 {
        data-size 1024;
        data-fill 0;
        destination-interface ms-1/2/0.10;
        dscp-code-points 001111;
        probe-count 10;
        probe-interval 1;
        probe-type icmp-ping;
        target address 172.17.20.182;
        test-interval 20;
        thresholds rtt 10;
        traps rtt-exceeded;
    }
}

```

```

    }
}
[edit interfaces]
ms-1/2/0 {
    unit 0 {
        family inet;
    }
    unit 10 {
        rpm client;
        family inet {
            address 192.0.2.1/32;
        }
    }
}
[edit chassis]
fpc 1 {
    pic 2 {
        adaptive-services {
            service-package {
                extension-provider {
                    control-cores 1;
                    data-cores 1;
                    object-cache-size 512;
                    policy-db-size 64;
                    package jservices-rpm;
                    syslog {
                        daemon any;
                    }
                }
            }
        }
    }
}
}
}
}

```

Configure the minimum statements necessary to enable TWAMP:

```

[edit services]
rpm {
    twamp {
        server {
            authentication-mode none;
            port 10000;                # Twamp server's listening port
        }
    }
}

```

```

        client-list LIST-1 {                # LIST-1 is the name of the client-list. Multiple
lists can be configured.
            address {
                198.51.100.2/30;            # IP address of the control client.
            }
        }
    }
}

[edit interfaces sp-5/0/0]
unit 0 {
    family inet;
}
unit 10 {
    rpm {
        twamp-server;                    # You must configure a separate logical interface on
the service PIC interface for the TWAMP server.
    }
    family inet {
        address 203.0.113.50/32;          # This address must be a host address with a 32-bit
mask.
    }
}

[edit chassis]
fpc 5 {
    pic 0 {
        adaptive-services {
            service-package layer-2;      # Configure the service PIC to run in Layer 2 mode.
        }
    }
}
}

```

Configure additional TWAMP settings:

```

[edit services]
rpm {
    twamp {
        server {
            maximum-sessions 5;
            maximum-sessions-per-connection 2;
            maximum-connections 3;
            maximum-connections-per-client 1;
            port 10000;
        }
    }
}

```

```

server-inactivity-timeout ;
client-list LIST-1 {
    address {
        198.51.100.2/30;
    }
}
}
}
}
}

```

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Examples: Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers With RPM | 674](#)

Enabling RPM on MX, M and T Series Routers and SRX Firewalls for the Services SDK

Real-time performance monitoring (RPM), which has been supported on the adaptive services interface, is now supported by the Services SDK. RPM is supported on all platforms and service PICs that support the Services SDK.

To enable RPM for the Junos OS extension-provider package on the adaptive services interface, configure the `object-cache-size`, `policy-db-size`, and `package` statements at the `[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]` hierarchy level. For the extension-provider package, `package-name` in the package `package-name` statement is `jservices-rpm`.

For more information about the extension-provider package, see the *SDK Applications Configuration Guide and Command Reference*.

The following example shows how to enable RPM for the extension-provider package on the adaptive services interface:

```

chassis fpc 1 {
    pic 2 {
        adaptive-services {
            service-package {

```



```

        extension-provider {
            control-cores 1;
            data-cores 1;
            object-cache-size 512;
            policy-db-size 64;
            package jservices-rpm;
            syslog daemon any;
        }
    }
}
}
}
}

```

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Examples: Configuring Real-Time Performance Monitoring on MX, M, T and PTX Series Routers | 680](#)

[destination-interface | 1046](#)

Understand Two-Way Active Measurement Protocol

IN THIS SECTION

- [TWAMP on MX Series Routers, EX9200 Series and QFX10000 Series Switches | 691](#)
- [TWAMP on PTX Series routers | 691](#)
- [TWAMP on QFX5000 Series switches | 692](#)
- [TWAMP on ACX Series routers | 693](#)

The Two-Way Active Management Protocol (TWAMP), described in RFC 5357, is an extension of the One-Way Active Management Protocol (OWAMP) that supplies two-way or round-trip measurements instead of unidirectional capabilities. Two-way measurements are helpful because round-trip delays do not require host clock synchronization and remote support might be a simple echo function. However, the Internet Control Message Protocol (ICMP) Echo Request/Reply (used by ping) for this purpose has

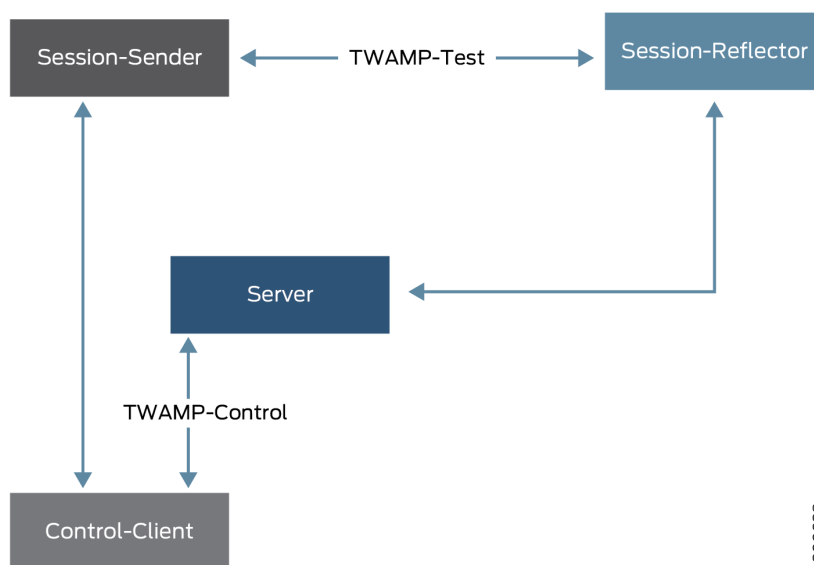
several shortcomings. TWAMP defines an open protocol for measuring two-way or round-trip metrics with greater accuracy than other methods by using time-stamps (processing delays can be factored as well).

Usually, TWAMP operates between interfaces on two devices playing specific roles. TWAMP is often used to check Service Level Agreement (SLA) compliance, and the TWAMP feature is often presented in that context. TWAMP uses two related protocols, running between several defined elements:

- TWAMP-Control—Initiates, starts, and ends test sessions. The TWAMP-Control protocol runs between a Control-Client element and a Server element.
- TWAMP-Test—Exchanges test packets between two TWAMP elements. The TWAMP-Test protocol runs between a Session-Sender element and a Session-Reflector element.

The four elements are shown in [Figure 60 on page 687](#):

Figure 60: Four Elements of TWAMP




Although four different TWAMP devices can perform the four logical roles of TWAMP Control-Client, Server, Session-Sender, and Session-Reflector, different devices can play different roles. A common implementation combines the roles of Control-Client and Session-Sender in one device (known as the TWAMP controller or TWAMP client) and the roles of Server and Session-Reflector in the other device (known as the TWAMP responder or TWAMP server). In this case, each device runs both the TWAMP-Control (between Control-Client and Server) and TWAMP-Test (between Session-Sender and Session-Reflector) protocols.

The TWAMP client-server architecture as implemented looks like this:

- TWAMP client
 - Control-Client sets up, starts and stops the TWAMP test sessions.
 - Session-Sender creates TWAMP test packets that are sent to the Session-Reflector in the TWAMP server.
- TWAMP server
 - Session-Reflector sends back a measurement packet when a test packet is received, but does not maintain a record of such information.
 - Server manages one or more sessions with the TWAMP client and listens for control messages on a TCP port.

The packaging of these elements into TWAMP client and TWAMP server processes is shown in [Figure 61 on page 688](#).

Figure 61: The Elements of TWAMP Implemented as Client (Left) and Server (Right).



[Table 117 on page 688](#) provides information about TWAMP and related timestamp support on MPC, MS-MIC/MPC, and inline:

Table 117: TWAMP and related timestamp support

Feature	Role	IP Version	Support (Y/N)	Timestamp Inline	Timestamp on MPC (hardware-timestamp)	Timestamp on MPC (si-interface)	Timestamp on MS-MIC/MPC (delegate-probes)
TWAMP	Client	IPv4	Y	N	Y (µsec) 500 maximum probes	Y (µsec) 500 maximum probes	N
		IPv6	N	N	N	N	N

Table 117: TWAMP and related timestamp support (Continued)

Feature	Role	IP Version	Support (Y/N)	Timestamp Inline	Timestamp on MPC (hardware-timestamp)	Timestamp on MPC (si-interface)	Timestamp on MS-MIC/MPC (delegate-probes)
	Server	IPv4	Y	N	Y (μsec) 500 maximum probes	Y (μsec) 500 maximum probes	N
		IPv6	N	N	N	N	N

[Table 118 on page 689](#) provides information about support for TWAMP Light, as defined in Appendix I of RFC 5357, which defines a light version of the TWAMP protocol, a stateless version of TWAMP where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away.

Support for IPv6 target addresses for TWAMP Light test sessions is introduced in Junos OS Release 21.3R1 and as mentioned in the table below.

Support for IPv6 link-local target addresses is introduced in Junos OS Release 21.4R1, for the MX Series and the PTX1000, PTX3000, and PTX5000 routers and in Junos OS Evolved Release 22.3R1, for the ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 routers.

Table 118: TWAMP Light Support

Device	Supported In
ACX710	Junos OS Release 22.3R1
ACX5448 Series	Junos OS Release 22.3R1
ACX7100 Series	Junos OS Evolved Release 21.2R1
ACX7509	Junos OS Evolved Release 22.3R1
MX Series, with LC480, LC2101, LC2103, and MPCs up to and including the MPC9E	Junos OS Release 21.1R1 (IPv4), Junos OS Release 21.3R1 (IPv6)

Table 118: TWAMP Light Support *(Continued)*

Device	Supported In
MX Series with the following line cards: LMIC16-BASE, LC9600, MPC10E, and MPC11E	<ul style="list-style-type: none"> IPv4 client: Junos OS Release 21.1R1 IPv4 server: Junos OS Release 22.2R1 IPv6 client and server: Junos OS Release 22.3R1
PTX Series running Junos OS, with MPCs up to and including the MPC9E	Junos OS Release 21.1R1 (IPv4), Junos OS Release 21.3R1 (IPv6)
PTX Series running Junos OS, with MPC10E and MPC11E line cards	<ul style="list-style-type: none"> client: Junos OS Release 21.1R1 (IPv4) server: Junos OS Release 22.2R1 (IPv4)
PTX10001-36MR	<ul style="list-style-type: none"> Junos OS Evolved Release 21.1R1 (IPv4) Junos OS Evolved Release 21.4R1 (IPv6)
PTX10003	<ul style="list-style-type: none"> Junos OS Evolved Release 20.3R1 (IPv4) Junos OS Evolved Release 21.4R1 (IPv6)
PTX10004	<ul style="list-style-type: none"> Junos OS Evolved Release 21.2R1 (IPv4) Junos OS Evolved Release 21.4R1 (IPv6)
PTX10008 and PTX10016 (with the JNP10008-SF3 and either the JNP10K-LC1201 or JNP10K-LC1202-36MR line card)	<ul style="list-style-type: none"> Junos OS Evolved Release 21.1R1 (IPv4) Junos OS Evolved Release 21.4R1 (IPv6)
QFX5130-32CD, QFX5220, and QFX5700	Junos OS Evolved 22.4R1 (IPv4 and IPv6)
QFX10002, QFX10008, and QFX10016	Junos OS Release 21.3R1 (IPv4)
EX9200	Junos OS Release 21.4R1

TWAMP on MX Series Routers, EX9200 Series and QFX10000 Series Switches

Both the control client and session sender (the TWAMP client) reside on the same Juniper Networks router. However, the TWAMP client does not require that the server and the session reflector to be on the same system. Therefore, the Juniper TWAMP client is capable of working with a third-party server implementation.

NOTE: TWAMP is not supported when you enable Next Gen Services on an MX Series router.

TWAMP on PTX Series routers

The TWAMP-Control protocol is used to set up performance measurement sessions between a TWAMP client and a TWAMP server, and the TWAMP-Test protocol is used to send and receive performance measurement probes. The destination interface `si-x/y/z` attribute, which is meant for enabling inline services, is not supported on PTX Series routers for TWAMP client configurations.

For Junos OS, TWAMP is configured at the `[edit services rpm twamp]` hierarchy level. For Junos OS Evolved, TWAMP is configured at the `[edit services monitoring twamp]` hierarchy level. [Table 119 on page 691](#) provides information about support for TWAMP.

Table 119: PTX Series TWAMP Support

Device	Supported In
PTX Series running Junos OS	Junos OS Release 19.2R1
PTX10001-36MR	<ul style="list-style-type: none"> Junos OS Evolved Release 21.1R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)
PTX10003	<ul style="list-style-type: none"> Junos OS Evolved Release 20.3R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)
PTX10004	<ul style="list-style-type: none"> Junos OS Evolved Release 21.2R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)

Table 119: PTX Series TWAMP Support (Continued)

Device	Supported In
PTX10008 (with the JNP10008-SF3 and either the JNP10K-LC1201 or JNP10K-LC1202-36MR line card)	<ul style="list-style-type: none"> Junos OS Evolved Release 21.1R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)
PTX10016 (with the JNP10008-SF3 and either the JNP10K-LC1201 or JNP10K-LC1202-36MR line card)	Junos OS Evolved Release 22.4R1 (IPv4 and IPv6)

The Junos OS Evolved support for TWAMP is limited to the following:

- IPv4 and IPv6 traffic only for control sessions and test sessions. Starting in Junos OS Evolved Release 21.4R1, IPv6 source and target addresses (except for link-local addresses) are supported for client lists, control connections, and test sessions.
- Probe statistics and history
- Control and test session status
- Test session probe generation and reception, as well as reflection
- Timestamps set by the Routing Engine or the Packet Forwarding Engine for IPv4 traffic. For IPv6 traffic, timestamps set by the Routing Engine only. For IPv6 traffic, starting in Junos OS Evolved 22.3R1, we support Packet Forwarding Engine timestamps. Prior to Junos OS Evolved Release 22.3R1, for IPv6 traffic, the `offload-type` statement at the `[edit services monitoring twamp client control-connection name test-session name]` hierarchy level should be configured as `none`.
- Error reporting through system log messages and SNMP traps only
- Unauthenticated mode only

TWAMP on QFX5000 Series switches

The TWAMP-Control protocol is used to set up performance measurement sessions between a TWAMP client and a TWAMP server, and the TWAMP-Test protocol is used to send and receive performance measurement probes. For Junos OS Evolved, TWAMP is configured at the `[edit services monitoring twamp]` hierarchy level.

Table 120: QFX5000 Series TWAMP Support

Device	Supported In
QFX5130-32CD	Junos OS Evolved Release 22.4R1
QFX5220	Junos OS Evolved Release 22.4R1
QFX5700	Junos OS Evolved Release 22.4R1

The Junos OS Evolved support for TWAMP is limited to the following:

- IPv4 and IPv6 source and target addresses (including link-local addresses) are supported for client lists, control connections, and test sessions.
- Probe statistics and history
- Control and test session status
- Test session probe generation and reception, as well as reflection
- Timestamps set by the Routing Engine or by the Packet Forwarding Engine for IPv4 and IPv6 traffic.
- Error reporting through system log messages and SNMP traps only
- Unauthenticated mode only

TWAMP on ACX Series routers

In Junos OS, TWAMP is supported for ACX routers. The ACX710 and ACX5448 Series routers support both reflection and generation. Other ACX Series routers running Junos OS support only reflection, not generation. For Junos OS, TWAMP is configured at the `[edit services rpm twamp]` hierarchy level.

In Junos OS Evolved, TWAMP is supported for ACX routers, for both reflection and generation. Starting in Junos OS Evolved 21.2R1, TWAMP (including TWAMP Light) is supported for the ACX7100 Series routers. For Junos OS Evolved, TWAMP is configured at the `[edit services monitoring twamp]` hierarchy level. The Junos OS Evolved support for TWAMP is limited to the following:

- IPv4 traffic only for control sessions and test sessions; IPv6 traffic support (except for link-local addresses) starting in Junos OS Evolved Release 21.4R1. Support for IPv6 link-local addresses for TWAMP Light test sessions only starting in Junos OS Evolved 22.3R1.
- Probe statistics and history
- Control and test session status

- Test session probe generation and reception, as well as reflection
- Timestamps set by the Routing Engine or the Packet Forwarding Engine for IPv4 traffic. For IPv6 traffic, timestamps set by the Routing Engine only. For IPv6 traffic, starting in Junos OS Evolved 22.3R1, we support Packet Forwarding Engine timestamps. Prior to Junos OS Evolved Release 22.3R1, for IPv6 traffic, the `offload-type` statement at the `[edit services monitoring twamp client control-connection name test-session name]` hierarchy level should be configured as `none`. Starting in Junos OS Evolved 22.4R1 for ACX routers, you can configure the `inline-timestamping` option of the `offload-type` statement to enable timestamps set inline by the hardware.
- Error reporting through system log messages only
- Unauthenticated mode only

RELATED DOCUMENTATION

[Example: Configuring TWAMP Client and Server on MX Series Routers | 705](#)

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches | 694](#)

Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches

IN THIS SECTION

- [Understand TWAMP Configuration | 695](#)
- [Configure a TWAMP Server | 699](#)
- [Configure a TWAMP Client | 702](#)

The Two-Way Active Measurement Protocol (TWAMP) defines a standard for measuring IP performance between two devices in a network. For more information on TWAMP, see RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)*. For more background information on TWAMP, see "[Understand Two-Way Active Measurement Protocol](#)" on page 686.

Understand TWAMP Configuration

Two-Way Active Measurement Protocol (TWAMP) support and configuration varies for hardware platform, physical interfaces, or virtual physical (services) interfaces. Support for RPM is not always an indicator of TWAMP support on a particular combination of platform and line card for Junos OS. The time stamps used in RPM and TWAMP are added in different places, depending on the hardware configuration. For example, different hardware components perform timestamping, either inline in the lookup (LU) chip, Routing Engine (Junos OS Evolved), the microkernel-based timestamping at the host Packet Forwarding Engine, or the line card such as a Multiservices Physical Interface Card (MS-PIC), Multiservices Modular Interface Card (MS-MIC), Multiservices Modular PIC Concentrator (MS-MPC), or Multiservices Dense Port Concentrator (MS-DPC).

The ACX710 and ACX5448 Series routers, which run Junos OS, support both reflection and generation. Other ACX Series routers running Junos OS support only reflection. ACX Series routers running Junos OS Evolved support both reflection and generation.

PTX Series routers running Junos OS do not support the destination interface `si-x/y/z` attribute, which is meant for enabling inline services, for TWAMP client configurations.

[Table 121 on page 695](#) provides information about support for TWAMP Light, as defined in Appendix I of RFC 5357, which defines a light version of the TWAMP protocol, a stateless version of TWAMP where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away.

Support for IPv6 target addresses for TWAMP Light test sessions is introduced in Junos OS Release 21.3R1 for MX Series and the PTX1000, PTX3000, and PTX5000 routers. For the Junos OS IPv6 TWAMP Light client, you must configure both the `target-address` and the `destination-port` statements at the `[edit services rpm twamp client control-connection control-client-name test-session test-session-name]` hierarchy level. Support for link-local target addresses for IPv6 TWAMP Light test sessions is introduced in Junos OS Release 21.4R1 for MX Series and the PTX1000, PTX3000, and PTX5000 routers and in Junos OS Evolved Release 22.3R1, for the ACX7100, ACX7509, PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 routers.

Table 121: TWAMP Light Support

Device	Supported In
ACX710	Junos OS Release 22.3R1
ACX5448 Series	Junos OS Release 22.3R1
ACX7100 Series	Junos OS Evolved Release 21.2R1
ACX7509	Junos OS Evolved Release 22.3R1

Table 121: TWAMP Light Support (Continued)

Device	Supported In
MX Series, with LC480, LC2101, LC2103, and MPCs up to and including the MPC9E	Junos OS Release 21.1R1 (IPv4), Junos OS Release 21.3R1 (IPv6)
MX Series with the following line cards: LMIC16-BASE, LC9600, MPC10E, and MPC11E	<ul style="list-style-type: none"> • IPv4 client: Junos OS Release 21.1R1 • IPv4 server: Junos OS Release 22.2R1 • IPv6 client and server: Junos OS Release 22.3R1
PTX Series running Junos OS, with MPCs up to and including the MPC9E	Junos OS Release 21.1R1 (IPv4), Junos OS Release 21.3R1 (IPv6)
PTX Series running Junos OS, with MPC10E and MPC11E line cards	<ul style="list-style-type: none"> • client: Junos OS Release 21.1R1 (IPv4) • server: Junos OS Release 22.2R1 (IPv4)
PTX10001-36MR	<ul style="list-style-type: none"> • Junos OS Evolved Release 21.1R1 (IPv4) • Junos OS Evolved Release 21.4R1 (IPv6)
PTX10003	<ul style="list-style-type: none"> • Junos OS Evolved Release 20.3R1 (IPv4) • Junos OS Evolved Release 21.4R1 (IPv6)
PTX10004	<ul style="list-style-type: none"> • Junos OS Evolved Release 21.2R1 (IPv4) • Junos OS Evolved Release 21.4R1 (IPv6)
PTX10008 and PTX10016 (with the JNP10008-SF3 and either the JNP10K-LC1201 or JNP10K-LC1202-36MR line card)	Junos OS Evolved Release 21.1R1
QFX5130-32CD, QFX5220, QFX5700	Junos OS Evolved Release 22.4R1
QFX10002, QFX10008, QFX10016	Junos OS Release 21.3R1 (IPv4)

Table 121: TWAMP Light Support (Continued)

Device	Supported In
EX9200	Junos OS Release 21.4R1

For Junos OS, TWAMP is configured at the [edit services rpm twamp] hierarchy level. For Junos OS Evolved, TWAMP is configured at the [edit services monitoring twamp] hierarchy level. [Table 122 on page 697](#) provides information about support for TWAMP.

Table 122: TWAMP Managed Support

Device	Supported In
ACX710	Junos OS Release 22.3R1 (IPv4)
ACX5448 Series	Junos OS Release 22.3R1 (IPv4)
ACX7100 Series	<ul style="list-style-type: none"> Junos OS Evolved Release 21.2R1 (IPv4) Junos OS Evolved Release 21.4R1 (IPv6)
ACX7509	Junos OS Evolved Release 22.3R1
MX Series	Junos OS Release 19.2R1
PTX Series running Junos OS	Junos OS Release 19.2R1
PTX10001-36MR	<ul style="list-style-type: none"> Junos OS Evolved Release 21.1R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)
PTX10003	<ul style="list-style-type: none"> Junos OS Evolved Release 20.3R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)
PTX10004	<ul style="list-style-type: none"> Junos OS Evolved Release 21.2R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)

Table 122: TWAMP Managed Support (*Continued*)

Device	Supported In
PTX10008 (with the JNP10008-SF3 and either the JNP10K-LC1201 or JNP10K-LC1202-36MR line card)	<ul style="list-style-type: none"> Junos OS Evolved Release 21.1R1 (IPv4) Junos OS Evolved Release 22.4R1 (IPv6)
PTX10016 (with the JNP10008-SF3 and either the JNP10K-LC1201 or JNP10K-LC1202-36MR line card)	Junos OS Evolved Release 22.4R1
QFX5130-32CD, QFX5220, QFX5700	Junos OS Evolved Release 22.4R1
QFX10002, QFX10008, QFX10016	Junos OS Release 21.3R1
EX9200	Junos OS Release 21.4R1

For Junos OS Evolved, TWAMP, including TWAMP Light, is supported, and is limited to the following:

- IPv4 and IPv6 traffic for control sessions and test sessions; IPv6 traffic support (except for link-local addresses) starting in Junos OS Evolved Release 21.4R1. Support for IPv6 link-local addresses for TWAMP Light test sessions only starting in Junos OS Evolved 22.3R1.
- Probe statistics and history
- Control and test session status
- Test session probe generation and reception, as well as reflection
- Timestamps set by the Routing Engine or the Packet Forwarding Engine for IPv4 traffic. For IPv6 traffic, timestamps set by the Routing Engine only. For IPv6 traffic, starting in Junos OS Evolved 22.3R1, we support Packet Forwarding Engine timestamps. Prior to Junos OS Evolved Release 22.3R1, for IPv6 traffic, the offload-type statement at the [edit services monitoring twamp client control-connection *name* test-session *name*] hierarchy level should be configured as none. Starting in Junos OS Evolved 22.4R1 for ACX routers, you can configure the inline-timestamping option of the offload-type statement to enable timestamps set inline by the hardware.
- Error reporting through system log messages and SNMP traps only
- Unauthenticated mode only

See ["TWAMP on ACX Series routers" on page 693](#) for information about IPv6 support for the ACX Series routers.

Table 123 on page 699 shows the relationship between RPM client and server support, TWAMP client (with the control component) and TWAMP server (with the responder component) support, and the hardware that performs timestamping.

Table 123: TWAMP Feature Support and Hardware for Junos OS, MX Series

TWAMP Feature Support	Routing Engine Timestamp	MS-PIC/MS-DPC Timestamp	MS-MIC/MS-MPC Timestamp	Packet Forwarding Engine (microkernel) Timestamp	Packet Forwarding Engine (LU) Timestamp (si-interface)
RPM Client	Yes	Yes	Yes	Yes	No
RPM Server	Yes	Yes	Yes	Yes	No
TWAMP Client	No	No	No	Yes	Yes
TWAMP Server	No	Yes	No	Yes (No responder configuration needed)	Yes

NOTE: Support for the services interfaces (sp-, ms-, and si- interfaces) are all slightly different.

Configure a TWAMP Server

With the exception of physical interfaces, TWAMP server configuration for Junos OS requires the following minimum configuration at the [edit services rpm twamp] hierarchy level:

```

server {
  authentication-mode mode;
  client-list list-name {
    address ip-address;
  }
  port 862;
}

```

Starting in Junos OS Release 21.3R1, you no longer need to configure the `authentication-mode` statement. The default mode is now `none`, which means that communications with the server are not authenticated.

- To specify the list of allowed control client hosts that can connect to this server, include the `client-list` statement at the `[edit services rpm twamp server]` hierarchy level. Each value you include must be a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can include multiple client lists, each of which can contain a maximum of 64 entries. You must configure at least one client address to enable TWAMP.
- ACX Series routers do not support authentication and encryption modes. The value for `authentication-mode` statement at the `[edit services rpm twamp server]` hierarchy level must be set to `none`.
- TWAMP control connection traffic always arrives on ACX routers with the listening port set as 862. Because this port number for traffic probes can be modified, probes that arrive with a different port number are not recognized and processed by ACX routers correctly. As a result, TWAMP traffic and host-bound packets are dropped in such a scenario.

"[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#)" on page 694 provides information about support for light control of the server.

For Junos OS, you can configure light control for the server (managed control is the default). The Junos OS TWAMP server configuration for light control requires the following minimum configuration at the `[edit services rpm twamp]` hierarchy level:

```
server {
  authentication-mode none;
  light;
  port (862 | 878 | 51000);
}
```

For Junos OS, for a list of restrictions on source addresses, see "[source-address \(TWAMP\)](#)" on page 1408.

For Junos OS Evolved, you can configure either managed or light control for the server. TWAMP server configuration for managed or light control requires the following minimum configuration at the `[edit services monitoring twamp]` hierarchy level, assuming you use the default port for TWAMP (862):

```
server {
  (managed | light);
}
```

For Junos OS Evolved, you cannot use the following addresses for the client-list source IP address used for probes:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

You can configure more than one client, and you can change the TWAMP listening port as long as the change is coordinated with the TWAMP client.

For microkernel-based timestamping in Junos OS, you don't need to configure an si- interface. In this case, the TWAMP connection and sessions are established based on the target address and route.

For inline timestamping in Junos OS, you need to configure si- or sp- services interfaces and the TWAMP server configuration requires the following statements at the [edit interfaces *service-interface-name*] hierarchy level:

```
user@router# show interfaces si-0/0/0
unit 10 {
    rpm twamp-server;
    family inet {
        address 10.10.10.1/24;
    }
}
```

```
user@router# show interfaces sp-0/0/0
unit 10 {
    rpm twamp-server;
    family inet {
        address 10.20.20.1/24;
    }
}
```

NOTE: You cannot configure the TWAMP server on unit 0 of a services interface. If you try, you will receive a configuration error.

(Junos OS only) To configure a TWAMP server on an inline services (si-) interface, configure the amount of bandwidth reserved on each Packet Forwarding Engine for tunnel traffic using inline services by including the `bandwidth (1g | 10g)` statement at the `[edit chassis fpc slot-number pic number inline-services]` hierarchy level. Specify the service PIC (sp-) logical interface that provides the TWAMP service by including the `twamp-server` statement at the `[edit interfaces sp-fpc/pic/port unit logical-unit-number family inet]` hierarchy level.

The `twamp-server` statement is not required for physical interface TWAMP server configuration.

Many other TWAMP server parameters are optional. See the TWAMP server configuration statements for details.

Configure a TWAMP Client

For Junos OS, to configure the TWAMP client service, include the `client` statement and related parameters at the `[edit services rpm twamp]` hierarchy level. For Junos OS Evolved, include the `client` statement and related options at the `[edit services monitoring twamp]` hierarchy level.

There are many options available for TWAMP client configuration. See the configuration statement topics and examples for details.

For microkernel-based timestamping in Junos OS, you don't need to configure an si- interface. In this case, the TWAMP connection and sessions are established based on the target address and route.

For inline timestamping in Junos OS, the si- interfaces are virtual physical interfaces that respond as a TWAMP server. However, you can also configure services interfaces to act as the TWAMP client, which performs the TWAMP controller role.

(Junos OS only) To configure a services interface as a TWAMP client, you configure the service parameters and the service interface as a TWAMP client.

To configure the TWAMP client services interface, include the `rpm twamp-client` statement at the `[edit interfaces si-interface-name]` hierarchy level:

```
user@router# show interfaces si-0/0/0
unit 0 {
    family inet;
}
unit 10 {
    rpm twamp-client;
    family inet {
        address 10.30.30.1/24
    }
}
```

NOTE: You cannot configure the TWAMP client on unit 0 of a service interface. If you try, you will receive a configuration error.

SEE ALSO

[Understand Two-Way Active Measurement Protocol | 686](#)

[Understanding TWAMP Auto-Restart | 713](#)

[Configuring TWAMP Client and TWAMP Server to Reconnect Automatically After TWAMP Server Unavailability | 716](#)

[Example: Configuring TWAMP Client and Server on MX Series Routers | 705](#)

[twamp | 1515](#)

Release History Table

Release	Description
22.4R1-EVO	Starting in Junos OS Evolved Release 22.4R1 for the PTX10001-36MR, PTX10003, PTX10004, PTX10008, and PTX10016 routers, IPv6 source and target addresses (except for link-local addresses) are supported for client lists, control connections, and test sessions.
22.4R1-EVO	Starting in Junos OS Evolved Release 22.4R1 on the PTX10016 router, you can configure SNMP traps for TWAMP.
22.4R1-EVO	Starting in Junos OS Evolved Release 22.4R1 for the QFX5130-32CD, QFX5220, and QFX5700 switches, we support the Two-Way Active Measurement Protocol (TWAMP) managed client and server for IPv4 and IPv6 addresses and the TWAMP Light client and server, as defined in Appendix I of RFC 5357, for IPv4 and IPv6 addresses (including IPv6 link-local addresses). We support both Routing Engine and Packet Forwarding Engine timestamps for TWAMP probes. We also support error reporting through SNMP traps as well as through system log messages.
22.4R1-EVO	Starting in Junos OS Evolved Release 22.4R1 for the ACX7100, ACX7509, and ACX7024 routers, we support inline timestamping, where the timestamping is done in hardware at the generator or the reflector.
22.3R1-EVO	Starting in Junos OS Evolved Release 22.3R1, for TWAMP Light test sessions, you can specify IPv6 link-local addresses for target addresses.

22.3R1	Starting in Junos OS Release 22.3R1 for the MX Series routers with line cards MPC10E, MPC11E, LMIC16-BASE, and LC9600, we support the Two-Way Active Measurement Protocol (TWAMP) Light client and server, as defined in Appendix I of RFC 5357, for IPv6 addresses.
22.3R1	Starting in Junos OS Release 22.3R1 for the ACX710 and ACX5448 Series routers, we support the Two-Way Active Measurement Protocol (TWAMP) managed client and server for IPv4 addresses and the TWAMP Light client and server, as defined in Appendix I of RFC 5357, for IPv4 and IPv6 addresses (except for IPv6 link-local addresses). We also support Packet Forwarding Engine timestamps for TWAMP probes.
22.2R1	Starting in Junos OS Release 22.2R1 for the MX Series and PTX1000, PTX3000, and PTX5000 routers, with line cards MPC10E, MPC11E, LMIC16-BASE, and LC9600, we support the Two-Way Active Measurement Protocol (TWAMP) Light server, as defined in Appendix I of RFC 5357, for IPv4 addresses.
21.4R1-EVO	Starting in Junos OS Evolved Release 21.4R1 for ACX7100 routers, IPv6 source and target addresses (except for link-local addresses) are supported for client lists, control connections, and test sessions.
21.4R1	Starting in Junos OS Release 21.4R1, the Two-Way Active Measurement Protocol (TWAMP) is supported on the EX9200 Series switches.
21.4R1	Starting in Junos OS Release 21.4R1, for TWAMP Light test sessions, you can specify IPv6 link-local addresses for target addresses, and can configure IPv6 addresses for source addresses that correspond to target addresses configured with IPv6 link-local addresses.
21.3R1-EVO	Starting in Junos OS Evolved Release 21.3R1 on PTX Series routers, you can configure SNMP traps for TWAMP.
21.3R1	Starting in Junos OS Release 21.3R1, the Two-Way Active Measurement Protocol (TWAMP) is supported on QFX10000 Series switches.
21.3R1	Starting in Junos OS Release 21.3R1 for the MX Series and PTX1000, PTX3000, and PTX5000 routers, IPv6 target addresses for TWAMP Light test sessions are supported.
21.3R1	Starting in Junos OS Release 21.3R1, you no longer have to configure the authentication-mode statement for the TWAMP server. The default mode is none.
21.2R1-EVO	Starting in Junos OS Evolved 21.2R1, the Two-Way Active Measurement Protocol (TWAMP) is supported on the PTX10004 and ACX7100 routers.

21.1R1-EVO	Starting in Junos OS Evolved 21.1R1, the Two-Way Active Measurement Protocol (TWAMP) is supported on the PTX10001-36MR and the PTX10008 (with the JNP10008-SF3 and either the JNP10K-LC1201 or JNP10K-LC1202-36MR line card) routers.
21.1R1	Starting in Junos OS Release 21.1R1 for the MX Series and PTX1000, PTX3000, and PTX5000 routers, with MPCs up to and including the MPC9E, we support the Two-Way Active Measurement Protocol (TWAMP) Light client and server, as defined in Appendix I of RFC 5357. for IPv4 target addresses. TWAMP Light is a stateless version of TWAMP, where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away. For the MPC10E, MPC11E, LC9600, and LMIC16-BASE line cards, we only support the TWAMP Light client for IPv4 target addresses.
20.3R1-EVO	Starting in Junos OS Evolved 20.3R1, the Two-Way Active Measurement Protocol (TWAMP) is supported on the PTX10003 router.
19.2R1	Starting in Junos OS Release 19.2R1, the Two-Way Active Measurement Protocol (TWAMP) is supported on PTX Series routers.

Example: Configuring TWAMP Client and Server on MX Series Routers

IN THIS SECTION

- [Requirements | 705](#)
- [Overview | 706](#)
- [Configuration for TWAMP client | 706](#)
- [Configuration for TWAMP server | 709](#)
- [Verification | 712](#)

This example shows how to configure the TWAMP client and server and contains the following sections.

Requirements

This example uses the following hardware and software components:

- MX Series routers.
- Junos OS Release 15.1 or later.

Overview

This example explains the Two-Way Active Measurement Protocol (TWAMP). TWAMP is an open protocol for measuring network performance between any two devices supporting the TWAMP protocol. The TWAMP-Control protocol is used to set up performance measurement sessions. The TWAMP-Test protocol is used to send and receive performance measurement probes.

The TWAMP architecture is composed of the following entities that are responsible for starting a monitoring session and exchanging packets:

- The control client initiates all requested test sessions with a start sessions message, and the server acknowledges. When necessary, the control client sends a message to stop all test sessions.
- The session sender and the session reflector exchange test packets according to the TWAMP-Test protocol for each active session. On receiving a TWAMP-Test packet, the session reflector only reflects a measurement packet and does not collect packet statistics in TWAMP.

The TWAMP server is an end system that manages one or more TWAMP sessions and is also capable of configuring per-session ports. The server listens on the TCP port. The session reflector and server make up the TWAMP responder in an IP service-level agreement operation.

For 15.1, both the control client and session sender would be residing on the same Juniper router. The client design does not mandate the server and the session reflector to be on the same system. Hence the Juniper TWAMP client will also be capable of working with a third-party server implementation.

Configuration for TWAMP client

IN THIS SECTION

- [CLI Quick Configuration | 706](#)
- [Configuring TWAMP client | 707](#)

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, of SG1 router.

Configuring Chassis

```
set chassis fpc 4 pic 1 inline-services bandwidth 1g
```

Configuring Interfaces

```
set interfaces si-4/1/0 unit 0 family inet
set interfaces si-4/1/0 unit 10 rpm twamp-client
set interfaces si-4/1/0 unit 10 family inet address 10.60.60.1/32
```

Configuring Services

```
set services rpm twamp client control-connection c1 destination-interface si-4/1/0.10
set services rpm twamp client control-connection c1 history-size 500
set services rpm twamp client control-connection c1 target-address 10.70.70.1
set services rpm twamp client control-connection c1 test-count 1
set services rpm twamp client control-connection c1 test-interval 1
set services rpm twamp client control-connection c1 traps test-iteration-done
set services rpm twamp client control-connection c1 traps control-connection-closed
set services rpm twamp client control-connection c1 test-session t1 target-address 10.70.70.1
set services rpm twamp client control-connection c1 test-session t1 data-fill-with-zeros
set services rpm twamp client control-connection c1 test-session t1 data-size 1400
set services rpm twamp client control-connection c1 test-session t1 probe-count 55
set services rpm twamp client control-connection c1 test-session t1 probe-interval 1
```

Configuring TWAMP client

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

1. Configure the chassis.

```
[edit chassis]
user@router1# set fpc 4 pic 1 inline-services bandwidth 1g
```

2. Configure the interfaces.

```
[edit interfaces]
user@router1# set si-4/1/0 unit 0 family inet
user@router1# set si-4/1/0 unit 10 rpm twamp-client
user@router1# set si-4/1/0 unit 10 family inet address 10.60.60.1/32
```

3. Configure the services.

```
[edit services]
user@router1# set rpm twamp client control-connection c1 destination-interface si-4/1/0.10
user@router1# set rpm twamp client control-connection c1 history-size 500
user@router1# set rpm twamp client control-connection c1 target-address 10.70.70.1
user@router1# set rpm twamp client control-connection c1 test-count 1
user@router1# set rpm twamp client control-connection c1 test-interval 1
user@router1# set rpm twamp client control-connection c1 traps test-iteration-done
user@router1# set rpm twamp client control-connection c1 traps control-connection-closed
user@router1# set rpm twamp client control-connection c1 test-session t1 target-address
10.70.70.1
user@router1# set rpm twamp client control-connection c1 test-session t1 data-fill-with-zeros
user@router1# set rpm twamp client control-connection c1 test-session t1 data-size 1400
user@router1# set rpm twamp client control-connection c1 test-session t1 probe-count 55
user@router1# set rpm twamp client control-connection c1 test-session t1 probe-interval 1
```

Results

From the configuration mode of Router 1, confirm your configuration by entering the `show chassis`, `show interfaces`, and `show services rpm twamp` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@router1# show chassis
fpc 4 {
  pic 1 {
    inline-services {
      bandwidth 1g;
    }
  }
}
```

```
user@router1# show interfaces
si-4/1/0 {
  unit 0 {
    family inet;
  }
  unit 10 {
    rpm twamp-client;
```

```

        family inet {
            address 10.60.60.1/32;
        }
    }
}

```

```

user@router1# show services rpm twamp
client {
    control-connection c1 {
        destination-interface si-4/1/0.10;
        history-size 500;
        target-address 10.70.70.1;
        test-count 1;
        test-interval 1;
        traps {
            test-iteration-done;
            control-connection-closed;
        }
        test-session t1 {
            target-address 10.70.70.1;
            data-fill-with-zeros;
            data-size 1400;
            probe-count 55;
            probe-interval 1;
        }
    }
}

```

Configuration for TWAMP server

IN THIS SECTION

- [CLI Quick Configuration | 710](#)
- [Configuring TWAMP server | 710](#)

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, of SG1 router.

Configuring Chassis

```
set chassis fpc 2 pic 1 inline-services bandwidth 1g
```

Configuring Interfaces

```
set interfaces si-2/1/0 unit 0 family inet
set interfaces si-2/1/0 unit 10 rpm twamp-server
set interfaces si-2/1/0 unit 10 family inet address 10.70.70.1/32
```

Configuring Services

```
set services rpm twamp server authentication-mode none
set services rpm twamp server port 862
set services rpm twamp server client-list Client1 address 10.60.60.1/32
```

Configuring TWAMP server

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

1. Configure the chassis.

```
[edit chassis]
user@router1# set fpc 2 pic 1 inline-services bandwidth 1g
```

2. Configure the interfaces.

```
[edit interfaces]
user@router1# set si-2/1/0 unit 0 family inet
```

```

user@router1#set si-2/1/0 unit 10 rpm twamp-server
user@router1#set si-2/1/0 unit 10 family inet address 10.70.70.1/32

```

3. Configure the services.

```

[edit services]
user@router1# set rpm twamp server authentication-mode none
user@router1# set rpm twamp server port 862
user@router1# set rpm twamp server client-list Client1 address 10.60.60.1/32

```

Results

From the configuration mode of Router 1, confirm your configuration by entering the `show chassis`, `show interfaces`, and `show services rpm twamp server` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```

user@router1# show chassis
fpc 2 {
  pic 1 {
    inline-services {
      bandwidth 1g;
    }
  }
}

```

```

user@router1# show interfaces
si-2/1/0 {
  unit 0 {
    family inet;
  }
  unit 10 {
    rpm twamp-server;
    family inet {
      address 10.70.70.1/32;
    }
  }
}

```

```
}
}

user@router1# show services rpm twamp server
authentication-mode none;
port 862;
client-list Client1 {
  address {
    10.60.60.1/32;
  }
}
```

Verification

IN THIS SECTION

[Verifying TWAMP server sessions | 712](#)

[Verifying TWAMP client sessions | 713](#)

Verifying TWAMP server sessions

Purpose

Verify that the TWAMP server sessions are established.

Action

From operational mode, enter the show services rpm twamp server session command.

```
user@router1> show services rpm twamp server session
```

Session ID	Connection ID	Sender address	Sender port	Reflector address	Reflector port
4	44	10.1.1.1	12345	192.168.219.203	890
78	44	10.22.1.55	345	10.22.2.2	89022
234	423	192.168.219.203	2345	10.2.22.2	3333

5	423	10.4.1.1	82345	10.2.2.2	45909
1	423	192.168.1.1	645	10.32.2.2	2394

Verifying TWAMP client sessions

Purpose

Verify that the TWAMP client sessions are established.

Action

From operational mode, enter the `show services rpm twamp client session` command.

```
user@router1> show services rpm twamp client session
```

Connection	Session	Sender	Sender	Reflector	Reflector
Name	Name	address	port	address	port
c2	t1	10.60.60.1	10008	10.70.70.1	10008

RELATED DOCUMENTATION

| [request services rpm twamp](#) | 1578

Understanding TWAMP Auto-Restart

IN THIS SECTION

- [Benefits](#) | 714
- [TCP Keepalive Support for TWAMP Client and Server](#) | 715

After a network outage or a configuration change, when the Two-Way Active Management Protocol (TWAMP) client goes down, you have to manually start the TWAMP session by using `request services rpm`

`twamp start client` command. Starting in Junos OS Release 19.1R1, the TWAMP client restarts automatically without any manual intervention.

The automatic restart of the TWAMP session enables the TWAMP client to initiate the TCP control connection and UDP test sessions automatically during the following scenarios:

- Immediately after the TWAMP client configuration is committed.
- After the remote operation daemon (rmopd) is started with the valid TWAMP client configuration presence.
- After the TWAMP client configuration is activated.
- Immediately after the TWAMP server is reachable from the TWAMP client, based on the `test-interval`.

When the network fails or the TWAMP server becomes unreachable for any reason, the TWAMP client tries to reconnect to the TWAMP server after every `test-interval` value until it is successful. However, for the client to reconnect to the TWAMP server automatically, the `test-count` value in the `set rpm twamp client control-connection test-count` command must be 0. At the TWAMP server side, the default value of `max-connection-duration` in the `set rpm twamp server max-connection-duration` must also be 0. Thereby, you can retain the connection until it is cleared.

NOTE: Starting in Junos OS Release 19.1R1, the default value of `test-count` at the TWAMP client and `max-connection-duration` at the TWAMP server is 0.

After you configure and commit a TWAMP test, the client runs tests indefinitely—that is, it continues to send probes after the configured test interval even after a test is completed, and even if there is a network or server failure. You can stop the automatic running of tests by changing the value of the `test-count` option to a nonzero value. If you do that, the automatic restart feature is disabled, and you need to manually start the TWAMP client for it to establish connection with the server and start test sessions.

You can maintain and view the statistics related to the previous probes sent during server unavailability. You can Use the `set services rpm twamp client control-connection c1 persistent-results` command to preserve and display the test results after the network recovers or when the TWAMP server is again reachable.

Benefits

- You do not need to restart the TWAMP session manually after the client goes down as a result of a network outage or configuration change.
- You do not need to run an event script to restart TWAMP session from client side.

TCP Keepalive Support for TWAMP Client and Server

Keepalive probes can assert client (peers) when another peer becomes unreachable. If the problem is in the network between two peers, the keepalive action is to wait for some time and then retry sending the keepalive packet before marking the connection as broken.

When the keepalive timer for a TCP connection reaches zero, TCP client sends its peer a keepalive probe packet with no data in it and with the ACK flag turned on. The client receives a reply from the remote host with no data and with the ACK flag set. If the client receives a reply to its keepalive probe, the client can assert that the connection is still up and running. If the peer does not reply to the keepalive probe, you can assert that the connection cannot be considered valid and then take corrective action.

In Junos OS, to detect the TWAMP control connection failures at TWAMP client and TWAMP servers, you need to configure the following parameters:

- `tcp-keepcnt`—Number of unacknowledged probes to send before considering the connection dead and notifying the application layer.
- `tcp-keepidle`—Time interval between the last data packet sent and the first keepalive probe sent.
- `tcp-keepintvl`—Time interval between successive keepalive probes.

Release History Table

Release	Description
19.1R1	Starting in Junos OS Release 19.1R1, the TWAMP client restarts automatically without any manual intervention.

RELATED DOCUMENTATION

tcp-keepcnt 1432
tcp-keepintvl 1435
tcp-keepidle 1434
Configuring TWAMP Client and TWAMP Server to Reconnect Automatically After TWAMP Server Unavailability 716

Configuring TWAMP Client and TWAMP Server to Reconnect Automatically After TWAMP Server Unavailability

You can run TWAMP client automatically without any manual intervention during network failures or configuration changes. In case of a network outage or connection loss between a TWAMP client and TWAMP server, all the affected TWAMP TCP control connections and UDP test-sessions are lost. At each test-interval, the TWAMP client continues to send the control packets to re-establish connectivity with TWAMP server till it is successful. All the statistics will be maintained during that network failure.

This procedure is for Junos OS only. To configure the TWAMP client:

1. Configure the interfaces.

```
[edit interfaces]
user@router1# set si-2/2/0 unit 0 family inet
user@router1# set si-2/2/0 unit 10 rpm twamp-client
user@router1# set si-2/2/0 unit 10 family inet address 192.168.20.1/32
```

2. Configure the chassis.

```
[edit chassis]
user@router1# set fpc 2 pic 2 inline-services bandwidth 1g
```

3. Configure the services.

```
[edit services]
user@router1# set rpm twamp client control-connection c1 destination-interface si-2/2/0.10
user@router1# set rpm twamp client control-connection c1 persistent-results
user@router1# set rpm twamp client control-connection c1 history-size 500
user@router1# set rpm twamp client control-connection c1 routing instance IN
user@router1# set rpm twamp client control-connection c1 target-address 192.0.2.2
user@router1# set rpm twamp client control-connection c1 tcp-keepidle 20
user@router1# set rpm twamp client control-connection c1 tcp-keepintvl 4
user@router1# set rpm twamp client control-connection c1 tcp-keepcnt 10
user@router1# set rpm twamp client control-connection c1 test-interval 4
user@router1# set rpm twamp client control-connection c1 traps control-connection-closed
user@router1# set rpm twamp client control-connection c1 test-session t1 target-address 192.0.2.2
user@router1# set rpm twamp client control-connection c1 test-session t1 data-fill-with-zeros
user@router1# set rpm twamp client control-connection c1 test-session t1 data-size 1400
user@router1# set rpm twamp client control-connection c1 test-session t1 probe-count 20
```

```

user@router1# set rpm twamp client control-connection c1 test-session t1 probe-interval 1
user@router1# set rpm twamp client control-connection c1 test-session t2 target-address
192.0.2.2
user@router1# set rpm twamp client control-connection c1 test-session t1 data-fill-with-zeros
user@router1# set rpm twamp client control-connection c1 test-session t1 data-size 1400
user@router1# set rpm twamp client control-connection c1 test-session t1 probe-count 20
user@router1# set rpm twamp client control-connection c1 test-session t1 probe-interval 1
user@router1# set rpm twamp client control-connection c1 test-session t1 thresholds total-
loss 10
user@router1# set rpm twamp client control-connection c1 test-session t1 thresholds jitter-
gress 20
user@router1# set rpm twamp client control-connection c1 test-session t2 target-address
192.0.3.1
user@router1# set rpm twamp client control-connection c1 test-session t2 data-fill-with-zeros
user@router1# set rpm twamp client control-connection c1 test-session t2 data-size 1400
user@router1# set rpm twamp client control-connection c1 test-session t2 probe-count 15
user@router1# set rpm twamp client control-connection c1 test-session t2 probe-interval 1
user@router1# set rpm twamp client control-connection c1 test-session t2 thresholds total-
loss 10
user@router1# set rpm twamp client control-connection c1 test-session t2 thresholds jitter-
gress 20

```

To configure the TWAMP server:

1. Configure the interfaces.

```

[edit interfaces]
user@router1# set si-1/1/0 unit 30 family inet
user@router1# set si-1/1/0 unit 30 rpm twamp-server
user@router1# set si-1/1/0 unit 30 family inet address 192.02.2/24

```

2. Configure the chassis.

```

[edit chassis]
user@router1# set fpc 1 pic 1 inline-services bandwidth 1g

```

3. Configure the services.

```

[edit services]
user@router1# set rpm twamp server tcp-keepidle 200

```



```

user@router1# set rpm twamp server tcp-keepintvl 20
user@router1# set rpm twamp server tcp-keepcnt 210
user@router1# set rpm twamp server authentication-mode none
user@router1# set rpm twamp server server-inactivity-timeout 5
user@router1# set rpm twamp server reflector-inactivity-timeout 15
user@router1# set rpm twamp server max-connection-duration 0
user@router1# set rpm twamp server maximum-sessions 100
user@router1# set rpm twamp server maximum-sessions-per-connection 50
user@router1# set rpm twamp server maximum-connections 500
user@router1# set rpm twamp server maximum-connections-per-client 500
user@router1# set rpm twamp server port 862
user@router1# set rpm twamp server client-list Client1 address 192.168.20.1/24

```

When the TWAMP server is reachable, the output for Junos OS is as follows. The TWAMP-Server-Status is Connected and the Number-Of-Retries-With-TWAMP-Server is 1

```

user@router1> show services rpm twamp client probe-results | no-more
Jan 11 11:43:42
  Owner: c1, Test: t1
  server-address: 192.0.2.2, server-port: 862, Client address: 192.168.20.1, Client port: 58991
  TWAMP-Server-Status: Connected, Number-Of-Retries-With-TWAMP-Server: 1
  Routing Instance Name: IN
  Destination interface name: si-2/2/0.10
  Test size: 20 probes
  Probe results:
    Response received
    Probe sent time: Fri Jan 11 11:43:41 2019
    Probe rcvd/timeout time: Fri Jan 11 11:43:41 2019
    Rtt: 57 usec, Egress jitter: 1 usec, Ingress jitter: -1 usec, Round trip jitter: 0 usec
    Egress interarrival jitter: 43 usec, Ingress interarrival jitter: 43 usec, Round trip
interarrival jitter: 1 usec
    Results over current test:

.....
.....
  Owner: c1, Test: t2
  server-address: 192.0.2.2, server-port: 862, Client address: 192.168.20.1, Client port: 58991
  TWAMP-Server-Status: Connected, Number-Of-Retries-With-TWAMP-Server: 1
  Routing Instance Name: IN
  Destination interface name: si-2/2/0.10
  Test size: 15 probes
  Probe results:

```

```

Response received
Probe sent time: Fri Jan 11 11:43:36 2019
Probe rcvd/timeout time: Fri Jan 11 11:43:36 2019
Rtt: 58 usec, Egress jitter: 1 usec, Ingress jitter: -1 usec, Round trip jitter: 0 usec
Egress interarrival jitter: 28 usec, Ingress interarrival jitter: 28 usec, Round trip
interarrival jitter: 0 usec
Results over current test:
Probes sent: 15, Probes received: 15, Loss percentage: 0.000000
Measurement: Round trip time
Samples: 15, Minimum: 57 usec, Maximum: 59 usec, Average: 58 usec, Peak to peak: 2 usec,
Stddev: 1 usec, Sum: 866 usec
Measurement: Positive egress jitter
.....
Measurement: Round trip time
Samples: 105, Minimum: 57 usec, Maximum: 59 usec, Average: 58 usec, Peak to peak: 2
usec, Stddev: 1 usec, Sum: 6062 usec
Measurement: Positive egress jitter
Samples: 77, Minimum: 0 usec, Maximum: 398 usec, Average: 12 usec, Peak to peak: 398
usec, Stddev: 63 usec, Sum: 925 usec
Measurement: Negative egress jitter
Samples: 18, Minimum: 16 usec, Maximum: 431 usec, Average: 69 usec, Peak to peak: 415
usec, Stddev: 91 usec, Sum: 1248 usec
Measurement: Positive ingress jitter
Samples: 19, Minimum: 0 usec, Maximum: 431 usec, Average: 66 usec, Peak to peak: 431
usec, Stddev: 90 usec, Sum: 1249 usec
Measurement: Negative ingress jitter
Samples: 76, Minimum: 1 usec, Maximum: 397 usec, Average: 12 usec, Peak to peak: 396
usec, Stddev: 63 usec, Sum: 922 usec
Measurement: Positive round trip jitter
Samples: 79, Minimum: 0 usec, Maximum: 1 usec, Average: 0 usec, Peak to peak: 1 usec,
Stddev: 0 usec, Sum: 26 usec
Measurement: Negative round trip jitter
Samples: 25, Minimum: 1 usec, Maximum: 1 usec, Average: 1 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 25 usec

```

After the server is deactivated using the command `deactivate interfaces si-1/1/0 unit 30`, the output is as follows for Junos OS. The TWAMP-Server-Status is Not Connected and the Number-Of-Retries-With-TWAMP-Server is 12:

```

user@router1> show services rpm twamp client probe-results control-connection c1 | no-more
Jan 11 11:48:24

```

```

Owner: c1, Test: t1
server-address: 192.0.2.2, server-port: 862, Client address: 192.168.20.1, Client port: 58991
TWAMP-Server-Status: Not Connected, Number-Of-Retries-With-TWAMP-Server: 12
Reflector address: 192.0.2.2, Reflector port: 14779, Sender address: 192.168.20.1, sender-
port: 14779
Routing Instance Name: IN
Destination interface name: si-2/2/0.10
Test size: 20 probes
Probe results:
  Response received
  Probe sent time: Fri Jan 11 11:45:38 2019
  Probe rcvd/timeout time: Fri Jan 11 11:45:38 2019
  Rtt: 55 usec, Egress jitter: -17 usec, Ingress jitter: 18 usec, Round trip jitter: 1 usec
  Egress interarrival jitter: 37 usec, Ingress interarrival jitter: 37 usec, Round trip
interarrival jitter: 0 usec
Results over current test:
  Probes sent: 10, Probes received: 10, Loss percentage: 0.000000
  Measurement: Round trip time

.....
  Samples: 17, Minimum: 0 usec, Maximum: 3 usec, Average: 0 usec, Peak to peak: 3 usec, Stddev:
1 usec, Sum: 4 usec
  Measurement: Negative round trip jitter
    Samples: 3, Minimum: 1 usec, Maximum: 3 usec, Average: 2 usec, Peak to peak: 2 usec,
Stddev: 1 usec, Sum: 5 usec
Results over all tests:
  Probes sent: 210, Probes received: 210, Loss percentage: 0.000000

.....

TWAMP-Server-Status: Not Connected, Number-Of-Retries-With-TWAMP-Server: 12
Reflector address: 192.0.2.2, Reflector port: 14778, Sender address: 192.168.20.1, sender-
port: 14778
Routing Instance Name: IN
Destination interface name: si-2/2/0.10
Test size: 15 probes
Probe results:
  Response received
  Probe sent time: Fri Jan 11 11:45:38 2019
  Probe rcvd/timeout time: Fri Jan 11 11:45:38 2019
  Rtt: 58 usec, Egress jitter: -18 usec, Ingress jitter: 19 usec, Round trip jitter: 0 usec

```

```

.....
Results over all tests:
  Probes sent: 160, Probes received: 160, Loss percentage: 0.000000
  Measurement: Round trip time
    Samples: 160, Minimum: 57 usec, Maximum: 59 usec, Average: 58 usec, Peak to peak: 2
usec, Stddev: 1 usec, Sum: 9232 usec
  Measurement: Positive egress jitter
    Samples: 119, Minimum: 0 usec, Maximum: 398 usec, Average: 12 usec, Peak to peak: 398
usec, Stddev: 62 usec, Sum: 1398 usec
  Measurement: Negative egress jitter
    Samples: 27, Minimum: 16 usec, Maximum: 431 usec, Average: 64 usec, Peak to peak: 415
usec, Stddev: 76 usec, Sum: 1723 usec
  Measurement: Positive ingress jitter
    Samples: 28, Minimum: 0 usec, Maximum: 431 usec, Average: 62 usec, Peak to peak: 431
usec, Stddev: 76 usec, Sum: 1727 usec
  Measurement: Negative ingress jitter
    Samples: 118, Minimum: 1 usec, Maximum: 397 usec, Average: 12 usec, Peak to peak: 396
usec, Stddev: 62 usec, Sum: 1400 usec
  Measurement: Positive round trip jitter
    Samples: 120, Minimum: 0 usec, Maximum: 1 usec, Average: 0 usec, Peak to peak: 1 usec,
Stddev: 0 usec, Sum: 39 usec
  Measurement: Negative round trip jitter
    Samples: 39, Minimum: 1 usec, Maximum: 1 usec, Average: 1 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 39 usec

```

After activating the server using the activate interfaces si-1/1/0 unit 30 command the output is as follows for Junos OS. The TWAMP-Server-Status is Connected and the Number-Of-Retries-With-TWAMP-Server is 12.

```

user@router1> show services rpm twamp client probe-results control-connection c1 | no-more
Jan 11 11:48:50
  Owner: c1, Test: t1
  server-address: 192.0.2.2, server-port: 862, Client address: 192.168.20.1, Client port: 58991
  TWAMP-Server-Status: Connected, Number-Of-Retries-With-TWAMP-Server: 12
  Reflector address: 192.0.2.2, Reflector port: 14963, Sender address: 192.168.20.1, sender-
port: 14963
  Routing Instance Name: IN
  Destination interface name: si-2/2/0.10
  Test size: 20 probes
  Probe results:
    Response received
    Probe sent time: Fri Jan 11 11:48:50 2019
    Probe rcvd/timeout time: Fri Jan 11 11:48:50 2019

```

```

.....
Results over all tests:
    Probes sent: 218, Probes received: 218, Loss percentage: 0.000000
    Measurement: Round trip time
        Samples: 218, Minimum: 54 usec, Maximum: 59 usec, Average: 56 usec, Peak to peak: 5
usec, Stddev: 1 usec, Sum: 12160 usec

.....
Owner: c1, Test: t2
server-address: 192.0.2.2, server-port: 862, Client address: 192.168.20.1, Client port: 58991
TWAMP-Server-Status: Connected, Number-Of-Retries-With-TWAMP-Server: 12
Reflector address: 192.0.2.2, Reflector port: 14962, Sender address: 192.168.20.1, sender-
port: 14962
Routing Instance Name: IN
Destination interface name: si-2/2/0.10
Test size: 15 probes
Probe results:
    Response received
    Probe sent time: Fri Jan 11 11:48:50 2019
    Probe rcvd/timeout time: Fri Jan 11 11:48:50 2019
    Rtt: 57 usec, Egress jitter: 2 usec, Ingress jitter: -3 usec,
.....
Results over all tests:
    Probes sent: 168, Probes received: 168, Loss percentage: 0.000000
    Measurement: Round trip time
        Samples: 168, Minimum: 57 usec, Maximum: 59 usec, Average: 58 usec, Peak to peak: 2
usec, Stddev: 1 usec, Sum: 9691 usec
    Measurement: Positive egress jitter
        Samples: 124, Minimum: 0 usec, Maximum: 398 usec, Average: 11 usec, Peak to peak: 398
usec, Stddev: 61 usec, Sum: 1406 usec
    Measurement: Negative egress jitter
        Samples: 29, Minimum: 16 usec, Maximum: 431 usec, Average: 62 usec, Peak to peak: 415
usec, Stddev: 74 usec, Sum: 1806 usec
    Measurement: Positive ingress jitter
        Samples: 30, Minimum: 0 usec, Maximum: 431 usec, Average: 60 usec, Peak to peak: 431
usec, Stddev: 74 usec, Sum: 1811 usec
    Measurement: Negative ingress jitter
        Samples: 123, Minimum: 1 usec, Maximum: 397 usec, Average: 11 usec, Peak to peak: 396
usec, Stddev: 61 usec, Sum: 1410 usec
    Measurement: Positive round trip jitter
        Samples: 125, Minimum: 0 usec, Maximum: 1 usec, Average: 0 usec, Peak to peak: 1 usec,
Stddev: 0 usec, Sum: 42 usec

```

Measurement: Negative round trip jitter
Samples: 42, Minimum: 1 usec, Maximum: 1 usec, Average: 1 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 42 usec

RELATED DOCUMENTATION

tcp-keepcnt 1432
tcp-keepintvl 1435
tcp-keepidle 1434
persistent-results 1296
Understanding TWAMP Auto-Restart 713

Managing License Server for Throughput Data Export

IN THIS CHAPTER

- License Server Management for Throughput Data Export on MX Series Routers for NAT, Firewall, and Inline Flow Monitoring Services | 724
- Guidelines for Configuring an MX Series Router to Transmit Per-Service Throughput to an External Log Collector | 726

License Server Management for Throughput Data Export on MX Series Routers for NAT, Firewall, and Inline Flow Monitoring Services

IN THIS SECTION

- Throughput Measurement and Export | 725

To support our transition to software defined networking (SDN), Juniper Networks supports the Software Business Model Transformation, which includes new licensing, pricing, and branding strategies that make it easier for users to extract value from Juniper software solutions. This value of this approach is known as the Juniper Software Advantage (JSA), which provides the following benefits:

- Simple—Simple to buy, use, and manage rights
- Repeatable—License models which facilitates repeatable use among multiple hardware platforms and usage scenarios.
- Measurable—License fees based on easy to measure usage

Although the licensing of JSA products is trust-based, Juniper Networks might periodically audit the usage of its products. License Measurement Tool (LMT) is a technique that is used to compute the usage

of individual Network Edge Products under JSA. MX Series routers need to define the mechanism for updating the LMT tool with information such as per-service throughput. For example, for services such as carrier-grade NAT and inline flow monitoring, the router needs to calculate per service throughput and update it in LMT.

On MX Series routers, the Routing Engine periodically sends query messages to every Service PIC on which the service, for which throughput collection is being performed, is configured to run. This polling is performed for all the services for which throughput measurement is enabled. Service PICs, upon receiving the query for a particular service, reply with the throughput measured during the last query interval, for that service. If a service PIC hosts multiple services, the Routing Engine sends separate throughput queries to that service PIC for all the services. If a service is configured on multiple services PICs, the Routing Engine aggregates the throughput values received from all of them and exports the aggregated throughput to the log collector in the predefined log format. The LMT application analyzes these values from log collector, performs aggregation on values collected from all routers, and displays them in the LMT application.

You can configure the capability to transmit the throughput details per service for the Junos Address Aware (carrier-grade NAT) and Junos Traffic Vision (previously known as Jflow) in the last time interval to an external log collector. The default time interval at which the throughput data is sent is 300 seconds, which you can configure to suit your network needs. Multiple instances of the same service running on different PICs within a router are supported. If the same service is running on different PICs within a router, the router transmits the consolidated final throughput to the log collector or server. This functionality is supported on MX Series routers with MS-MCPs and MS-MICs, and also in the MX Series Virtual Chassis configuration. To configure the license server properties for throughput data to be transmitted for the defined services, such as NAT or stateful firewall, from the service PIC on the router to the external log collector, include the `license-server` statement at the `[edit]` hierarchy level. To specify the IP address of the license log server, include the `ip-address address` statement at the `[edit license-server]` hierarchy level. To configure the frequency of transmission of throughput data, include the `log-interval seconds` statement at the `[edit license-server]` hierarchy level. To specify the services for which throughput data collection must be performed, include the `services (jflow | cgnat | firewall)` statement at the `[edit license-server]` hierarchy level.

Throughput Measurement and Export

Throughput is defined as: “The network traffic throughput processed by juniper software in a second. It is represented as Mb/Sec (Megabits per second) or GB/sec (Gigabits per second). Throughput is measured as the 95th percentile of all the peaks measured in a quarter.” Service PICs keep track of the amount of data (in bits) processed by the various service plugins running on them. When a throughput query arrives from the Routing Engine, for a particular service, the Service PIC returns the value D/T mbps, in its reply, where:

- D is the amount of data (megabits) processed by that service since the previous query was received. If the query interval happens to be 300 seconds, for example, then D refers to the amount of data that was processed during the last 300 second interval. If the current query happens to be the very

first query, for a particular service, then D represents the cumulative data bits processed so far, by that service.

- T is the time (seconds) that elapsed since the previous query was received. This is the query interval configured using the CLI interface. If the current query happens to be the very first query, for a particular service, then T represents the time that elapsed since that service started processing packets. For all subsequent queries, T equals the query interval.

The Routing Engine aggregates the throughput measured (in mbps) across all the Service PICs on which a particular Service is configured and exports it to the Log collector which performs the 95th percentile calculation.

RELATED DOCUMENTATION

[License Enforcement](#)

[Verifying Junos OS License Installation \(CLI\)](#)

[show system license](#)

Guidelines for Configuring an MX Series Router to Transmit Per-Service Throughput to an External Log Collector

Observe the following guidelines while configuring this functionality on MX Series routers with MS-MPCs and MS-MICs:

- If the syslog server is unreachable, the router cannot send information to the log collector.
- After a graceful Routing Engine switchover (GRES) procedure, the newly functioning active Routing Engine starts sending the data to the server after the configured time interval, which is similar to a reset operation. The time elapsed in the active interval and data before GRES are not preserved.
- The time range must be from 60 through 86400 seconds (24 hours).
- If the timer is not configured, the default value of 300 seconds is assumed.
- The throughput data can be sent only if a service is up and running.
- Only maximum throughput is transmitted for the last 300 seconds or the configured time interval.
- The throughput value must not be less than zero to enable transmission. The data is sent based on the timezone of the router.

- An acknowledgment mechanism for data sent to the log collector is not supported. The router does not receive any acknowledgement regarding whether the data is already written into the log collector.
- The router does not maintain throughput data beyond the configurable time interval.
- No mechanisms exist to track if the log collector is successfully receiving the sent data or if the log server is reachable.
- The time interval and log collector are common for all the services; you cannot configure a different period for collection of logs for each service or a different log collector for each service.
- You cannot clear the system throughput value using a CLI command. It is assumed that the throughput value is not cleared or changed from outside. Throughput must be calculated internally by services and must not be manually modified by a CLI.
- SNMP support for these values is not available.
- The log collector performs a 95 percentile calculation of throughput data. Syslogs are sent even in scaled system conditions to the log collector for the throughput data related to the configured services.
- The following is the format of the syslogs configured to be sent at the prescribed frequency:

```
<Date>    <Time> < time-zone> <Router_name> <Service_name> <Throughput_value> Throughput =
<Unit_Mbps/Gbps> in last <Time_Interval>
```

An example is as follows:

```
Jan  8 08:49:57  America/Adak deuterium  CGNAT Throughput = 1500000 Mbps in last 300Sec
```

Testing the Performance of Network Devices Using RFC 2544-Based Benchmarking

IN THIS CHAPTER

- [Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)
- [Understanding RFC2544-Based Benchmarking Tests for E-LAN and E-Line Services on MX Series Routers | 733](#)
- [Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | 737](#)
- [Configuring an RFC 2544-Based Benchmarking Test | 739](#)
- [Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers | 747](#)
- [Example: Configure an RFC 2544-Based Benchmarking Test on an MX104 Router for Layer 3 IPv4 Services | 749](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for UNI Direction of Ethernet Pseudowires | 763](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for NNI Direction of Ethernet Pseudowires | 775](#)
- [Example: Configuring RFC2544-Based Benchmarking Tests on an MX104 Router for Layer 2 E-LAN Services in Bridge Domains | 788](#)
- [Example: Configuring Benchmarking Tests to Measure SLA Parameters for E-LAN Services on an MX104 Router Using VPLS | 823](#)

Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices

RFC 2544 defines a series of tests that can be used to describe the performance characteristics of a network-interconnecting device, such as a router, and outlines specific formats to report the results of the tests. These tests can be used to benchmark interconnected network devices and devise a guideline or a measurement pattern to analyze the health and efficiency of the network devices. These tests are the standard benchmarking tests for Ethernet networks and are known as RFC 2544-based benchmarking tests. These tests measure throughput, latency, frame loss rate, and bursty frames. The

test methodology enables you to define various parameters such as the different frame sizes to be examined (64, 128, 256, 512, 1024, 1280, and 1518 bytes), the test time for each test iteration (10 seconds through 1,728,000 seconds), and the frame format (UDP-over-IP).

The RFC 2544-based benchmarking test methodology assesses different parameters that are defined in service-level agreements (SLAs). By measuring the performance availability, transmission delay, link bursts, and service integrity, a carrier provider can certify that the working parameters of the deployed Ethernet circuit comply with the SLA and other defined policies.

NOTE: MX Series routers and SRX devices support only the reflector function in RFC 2544-based benchmarking tests.

NOTE: RFC 2544-based benchmarking tests support only UDP over IPv4 test traffic (unicast).

An RFC 2544-based benchmarking test is performed by transmitting test packets from a device that functions as the generator or the initiator (which is also called the originator). These packets are sent to a device that functions as a reflector, which receives and returns the packets to the initiator.

Starting from Junos OS Release 15.1, MX104 Series routers also perform verification of signatures on the received test frames. Starting in Junos OS Release 21.1R1, SRX300 and SRX550HM devices perform verification of signatures on the received test frames. By default, when the router or device receives a test packet that does not have the signature pattern, the packet is dropped. If you generate test traffic using a third-party vendor tool instead of an ACX Series router, you can disable signature verification. To disable signature verification, configure the `disable-signature-check` statement at the `[edit services rpm rfc2544-benchmarking tests test-name test-name]` hierarchy level.

For MX80 and MX104 Universal Routing Platforms and, starting from Junos OS Release 16.1, MX240, MX480, and MX960 Universal Routing Platforms with MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP), support the reflector function and the corresponding benchmarking tests.

Starting from Junos OS Release 17.1R1, the reflector function and the corresponding benchmarking tests are supported on MX Series routers with MPC3E (MX-MPC3E-3D), MPC3E-NG (MX-MPC3E-3D-NG), MPC4E (MPC4E-3D-32XGE-SFP and MPC4E-3D-2CGE-8XGE), MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, and MPC5EQ-100G10G), and MPC6E (MX2K-MPC6E).

Starting in Junos OS Release 20.2R1, the RFC 2544-based benchmarking tests are supported on MX240, MX480, and MX960 routers with the MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line card and on MX2010 and MX2020 routers with the MX2K-MPC11E line card.

Junos OS Release 20.3R1 extends support for the RFC 2544-based benchmarking tests onto the MX240, MX480, and MX960 routers with the MPC7E-MRATE or MPC7E-10G line card, onto the

MX2008, MX2010, and MX2020 routers with the MX2K-MPC8E or MX2K-MPC9E line card, and onto the MX204 and MX10003 (with the LC2103 card) routers.

Starting in Junos OS Release 21.1R1, the IPv4 Layer 3 reflector function and the corresponding benchmarking tests are supported on the SRX300 and SRX550HM devices.

NOTE: To configure RFC2544-based benchmarking tests on MX Series routers, see ["Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers" on page 747](#).

[Table 124 on page 730](#) describes the different MX network topologies in which the benchmarking test is supported.

Table 124: Supported MX Network Topologies for RFC2544 Benchmarking Tests

Service Type	Traffic Direction	Mode	Initial Release on MX104 Series Routers	Initial Release on MX204, MX2008, and MX10003 Series Routers	Initial Release on MX240, MX480, and MX960 Series Routers	Initial Release on MX2010 and MX2020 Series Routers	Whether the Benchmarking Test Is Supported
E-Line (family bridge)	(UNI) Egress (UNI) Ingress	Port Port, VLAN	14.2R1 (E-Line family bridge) 17.1R1	20.3R1 20.3R1	16.1R1 17.1R1	20.2R1 20.2R1	Supported
E-LAN (family bridge and family vpls)	(UNI) Egress (UNI) Ingress	Port Port, VLAN	14.2R1 (E-LAN family bridge) 15.1R1 (E-LAN family vpls) 17.1R1	20.3R1 20.3R1	16.1R1 17.1R1	20.2R1 20.2R1	Supported

Table 124: Supported MX Network Topologies for RFC2544 Benchmarking Tests (Continued)

Service Type	Traffic Direction	Mode	Initial Release on MX104 Series Routers	Initial Release on MX204, MX2008, and MX10003 Series Routers	Initial Release on MX240, MX480, and MX960 Series Routers	Initial Release on MX2010 and MX2020 Series Routers	Whether the Benchmarking Test Is Supported
E-Line (family ccc)	Ingress Egress	Port Port, VLAN	13.3R1 (E-Line pseudowire)	20.3R1 20.3R1	16.1R1	20.2R1	Supported
IP Services (family inet)	NNI	Port Port, VLAN	13.3R1	20.3R1	16.1R1	20.2R1	Supported

NOTE: You can configure a total of four simultaneous active reflection sessions. The four active reflection sessions can be of the same type or can be a combination of the different types of reflection sessions. For instance, you can configure either four IPv4 reflection sessions or one session each for pseudowire reflection, VPLS reflection, Layer 2 reflection, and IPv4 reflection. The maximum reflection bandwidth supported is 4 Gbps in a standalone test condition. Starting in Junos OS Release 20.2R1, MPC10E and MX2K-MPC11E support a maximum reflection bandwidth of 100 Gbps.

Table 125 on page 731 lists the interfaces and the reflection type on which the benchmarking tests are supported.

Table 125: Supported Interfaces for RFC2544 Benchmarking Tests

Type of Reflection	Gigabit Interfaces (ge)	Aggregated Interfaces (ae)	40G/100G interfaces (et) for MPC10E and MX2K-MPC11E	10G Interfaces (xe)	Pseudo Interfaces (irb, lt, vt, lo0, and others)
IPv4	Yes	Yes	Yes	Yes	No

Table 125: Supported Interfaces for RFC2544 Benchmarking Tests (Continued)

Type of Reflection	Gigabit Interfaces (ge)	Aggregated Interfaces (ae)	40G/100G interfaces (et) for MPC10E and MX2K-MPC11E	10G Interfaces (xe)	Pseudo Interfaces (irb, lt, vt, lo0, and others)
Pseudowire ingress	Yes	Yes	Yes	yes	No
Pseudowire egress	Yes	Yes (starting in Junos OS Release 15.1)	Yes	Yes (starting in Junos OS Release 15.1)	No
Layer 2 bridge	Yes	Yes	Yes	Yes	No
Layer 2 VPLS	Yes	Yes	Yes	Yes	No

All active RFC2544-based benchmarking tests are stopped when any of the following events take place:

- System events, such as Packet Forwarding Engine restarts, Routing Engine restarts, and so on.
- Test interface change events, such as deactivation and reactivation of the interface, disabling and enabling of the interface, child link events for aggregated interfaces and so on.

After the benchmarking tests are stopped, the test states of the tests are removed and the user can restart the same test. Other ongoing tests on other interfaces are not interrupted.

NOTE: RFC2544-based benchmarking tests are not supported during an unified in-service software upgrade (ISSU) or a graceful Routing Engine switchover (GRES).

Release History Table

Release	Description
21.1R1	Starting from Junos OS Release 21.1R1, the IPv4 Layer 3 reflector function and the corresponding benchmarking tests are supported on the SRX300 and SRX550HM devices.

20.3R1	Junos OS Release 20.3R1 extends support for the RFC 2544-based benchmarking tests onto the MX240, MX480, and MX960 routers with the MPC7E-MRATE or MPC7E-10G line card, onto the MX2008, MX2010, and MX2020 routers with the MX2K-MPC8E or MX2K-MPC9E line card, and onto the MX204 and MX10003 (with the LC2103 card) routers.
20.2R1	Starting in Junos OS Release 20.2R1, the RFC 2544-based benchmarking tests are supported on MX240, MX480, and MX960 routers with the MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line card and on MX2010 and MX2020 routers with the MX2K-MPC11E line card.
17.1R1	Starting from Junos OS Release 17.1R1, the reflector function and the corresponding benchmarking tests are supported on MX Series routers with MPC3E (MX-MPC3E-3D), MPC3E-NG (MX-MPC3E-3D-NG), MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE), MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, and MPC5EQ-100G10G), and MPC6E (MX2K-MPC6E).
16.1	For MX80 and MX104 Universal Routing Platforms and, starting from Junos OS Release 16.1, MX240, MX480, and MX960 Universal Routing Platforms with MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP), support the reflector function and the corresponding benchmarking tests.
15.1	Starting from Junos OS Release 15.1, MX104 Series routers also perform verification of signatures on the received test frames.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | 737](#)

[Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers | 747](#)

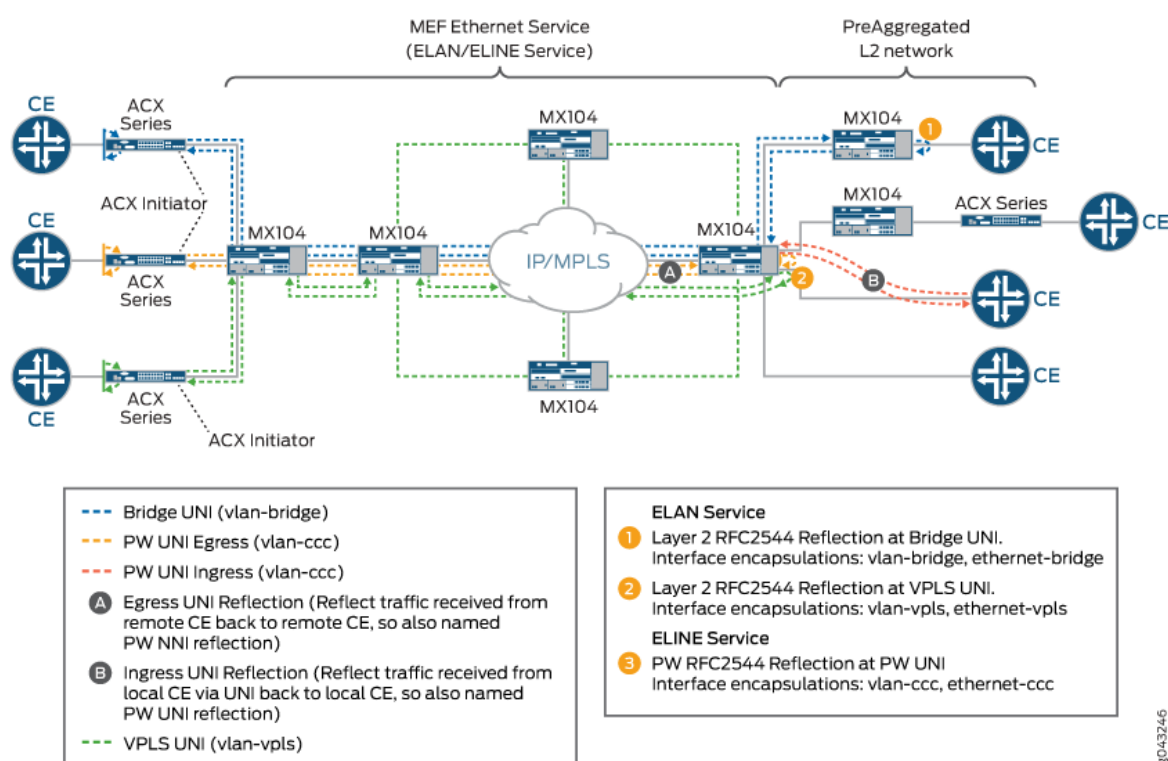
Understanding RFC2544-Based Benchmarking Tests for E-LAN and E-Line Services on MX Series Routers

NOTE: MX Series routers support only the reflector function in RFC2544-based benchmarking tests.

The Metro Ethernet Forum (MEF) defines two Ethernet service types—E-LAN and E-Line—and specifies the associated service attributes and parameters. These services can be supported within the Metro

Ethernet Network (MEN) and also supported over different transport technologies such as SONET, MPLS, and so on. Juniper Networks ACX Series routers, MX80, MX104 Series routers and MX240, MX480, and MX960 Series routers with MPC1, MPC2, and 16-port 10-Gigabit Ethernet MPC provide support for Layer 2 E-LAN and E-Line services reflection. [Figure 62 on page 734](#) shows a sample topology for the E-LAN and E-Line reflection supported on MX104 Series routers.

Figure 62: E-LAN And E-Line Reflection in a metro Solution



Starting in Junos OS Release 15.1, MX104 Series routers support RFC2544-based benchmarking tests for Layer 2 reflection (E-Line service) by using pseudowires (Layer 2 circuit and L2VPN). Starting in Junos OS Release 16.1, MX80 Series routers and MX240, MX480, and MX960 Series routers with MPC1, MPC2, and 16-port 10-gigabit Ethernet MPC support RFC2544-based benchmarking tests for Layer 2 reflection (E-Line service). E-Line provides transparent data transport. You can configure RFC2544-based benchmarking tests for both ingress and egress direction on the customer edge (CE) facing interface of family type CCC for an Ethernet pseudowire.

NOTE: To configure RFC2544-based benchmarking tests on MX240, MX480, MX960 Series routers with MPC1, MPC2, and the 16-port 10-Gigabit Ethernet MPC, see ["Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers" on page 747](#).

Starting in Junos OS Release 15.1, MX104 routers support RFC2544-based benchmarking tests for Layer 2 reflection (E-LAN service) by using VPLS and basic bridge domains. In Junos OS Release 14.2 and earlier, only basic bridge domains are used. Starting in Junos OS Release 16.1, MX80 Series routers and MX240, MX480, and MX960 Series routers with MPC1, MPC2, and 16-port 10-gigabit Ethernet MPC support RFC2544-based benchmarking tests for Layer 2 reflection (E-LAN service) by using VPLS and basic bridge domains. VPLS enables geographically dispersed sites to share an Ethernet broadcast domain by connecting sites across an MPLS network. All sites appear to be in the same Ethernet LAN though traffic travels across the MPLS network. Both LDP-based VPLS and BGP-based VPLS are supported. RFC2544-based benchmarking and performance measurement testing for Layer 2 E-LAN services (bridge/ VPLS) is supported on unicast traffic in egress direction only.

During the benchmarking tests, the initiator or generator transmits a test packet (unicast) to a reflector. The reflector receives and reflects the test packet back to the initiator. The test packet is an UDP-over-IP packet with a source and destination MAC address.

In a E-LAN service, the Layer 2 traffic reflection session is identified by the source MAC address, the destination MAC address, and the egress interface (logical interface). By default, RFC2544-based benchmarking tests are performed when there is no other service traffic. This mode of operation is known as out-of-service mode. The default service mode for the reflecting egress interface for an E-LAN service is also out-of-service mode. In out-of-service mode, while the test is running, all the data traffic (other than test traffic) sent to and from the test interface under test is interrupted. If the test is activated on a logical interface, all the traffic sent to and from the logical interface is interrupted. However, if there are other logical interfaces on the UNI port, the traffic sent to and from those logical interfaces is not interrupted. Control protocol peering is not interrupted whereas pass through control protocol packets such as end-to-end CFM sessions are interrupted. If you do not want the control protocol packets interrupted, you can configure the E-LAN service mode as in-service mode. In the in-service mode, while the test is running, the rest of the data traffic flow sent to and from the UNI port under test on the service is not interrupted. Both peering and pass through control protocols are not interrupted.

In an E-Line service, the reflection session is identified by the egress interface which is the logical interface. On activation of reflection on a logical interface, the traffic received on the logical interface is reflected. You can specify the type of traffic you want reflected by specifying the EtherType (specifies the protocol transported). If you do not specify the EtherType, all traffic is reflected. System does not explicitly block other traffic on the test interface during E-line service. You can block non-test traffic using firewall filters.

By default, for E-LAN services, the reflector swaps MAC addresses. The reflector swaps the source and destination MAC addresses and sends the packet back to the initiator. By default, for E-Line services, the reflector does not swap MAC addresses. [Table 126 on page 736](#) describes the MAC address swapping behavior for the service types.

Table 126: MAC Address Swapping Behavior for E-LAN and E-Line Services

Family	Direction	Default Behavior	User-configurable
bridge	Egress	MAC address swap (E-LAN service type)	No
	Ingress	No MAC address swap (E-Line service type)	Yes
vpls	Egress	MAC address swap (E-LAN service type)	No
ccc	Egress	No MAC address swap	Yes (starting in Junos OS Release 15.1)
	Ingress	MAC address swap	No

By default, the IP addresses and UDP ports are not modified. Optionally, you can configure the reflector to swap the source and destination IP address and the source and destination UDP ports.

You can configure an ACX Series router to operate as an initiator as well as a reflector. The MX104 Series router can be configured to operate only as a reflector.

Starting in Junos OS Release 15.1, MX104 Series routers support the specification of the protocol transported in the Ethernet frame. Starting in Junos OS Release 16.1, MX80 Series routers and MX240, MX480, and MX960 Series routers with MPC1, MPC2, and 16-port 10-gigabit Ethernet MPC also support the specification of the protocol transported in the Ethernet frame. To specify the EtherType (specifies the protocol transported) used for reflection of the test frames, use the `reflect-etype` command. If you do not specify the EtherType, all EtherTypes are reflected.

NOTE: The maximum reflection bandwidth supported is 4 Gbps. Because RFC2544 reflection shares system bandwidth with other loopback services such as tunnel services, you must manage the sharing of bandwidth for performing RFC2544-based performance tests.

NOTE: RFC2544-based benchmarking tests are not supported during unified in-service software upgrade (ISSU) and graceful Routing Engine switchover (GRES).

Release History Table

Release	Description
16.1	Starting in Junos OS Release 15.1, MX104 Series routers support the specification of the protocol transported in the Ethernet frame.
15.1	Starting in Junos OS Release 15.1, MX104 Series routers support RFC2544-based benchmarking tests for Layer 2 reflection (E-Line service) by using pseudowires (Layer 2 circuit and L2VPN).
15.1	Starting in Junos OS Release 15.1, MX104 routers support RFC2544-based benchmarking tests for Layer 2 reflection (E-LAN service) by using VPLS and basic bridge domains.

RELATED DOCUMENTATION

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | 737](#)

[Example: Configuring RFC2544-Based Benchmarking Tests on an MX104 Router for Layer 2 E-LAN Services in Bridge Domains | 788](#)

[disable-signature-check | 1062](#)

[reflect-etype | 1333](#)

Supported RFC 2544-Based Benchmarking Statements on MX Series Routers

Table 127 on page 738 lists the reflector-specific configuration statements that are supported on the MX Series routers. Note that en dash (–) specified in the Initial Release on MX Series routers column denotes that the command is not supported.

Table 127: Supported RFC2544-Based Benchmarking Reflector Statements on MX Series

Statement	Options	Initial Release on MX104 Series Routers	Initial Release on MX240, MX480, MX960 Series Routers
<code>destination-ipv4-address</code>	–	13.3R1	16.1R1
<code>destination-mac-address</code>	–	14.2R1	16.1R1
<code>destination-udp-port</code>	–	13.3R1	16.1R1
<code>direction</code>	(egress ingress)	13.3R1	16.1R1
<code>disable-signature-check</code>	–	15.1R1	16.1R1
<code>family</code>	(ccc inet) (bridge ccc inet) (vpls)	13.3R1 14.2R1 15.1R1	16.1R1
<code>in-service</code>	–	14.2R1	16.1R1
<code>ip-swap</code>	–	14.2R1	16.1R1
<code>mode</code>	reflect	13.3R1	16.1R1
<code>reflect-etype</code>	–	15.1R1	16.1R1
<code>reflect-mode</code>	(mac-swap no-mac-swap)	14.2R1	16.1R1
<code>service-type</code>	(eline elan)	14.2R1	16.1R1
<code>source-ipv4-address</code>	–	13.3R1	16.1R1

Table 127: Supported RFC2544-Based Benchmarking Reflector Statements on MX Series (Continued)

Statement	Options	Initial Release on MX104 Series Routers	Initial Release on MX240, MX480, MX960 Series Routers
source-mac-address	–	14.2R1	16.1R1
source-udp-port	–	13.3R1	16.1R1
test-interface	–	13.3R1	16.1R1
udp-tcp-port-swap	–	14.2R1	16.1R1

RELATED DOCUMENTATION

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Example: Configuring RFC2544-Based Benchmarking Tests on an MX104 Router for Layer 2 E-LAN Services in Bridge Domains | 788](#)

Configuring an RFC 2544-Based Benchmarking Test**IN THIS SECTION**

- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network | 741](#)
- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire | 743](#)
- [Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain | 745](#)

You can configure a benchmarking test to detect and measure performance attributes, such as throughput, latency, frame loss, and bursty or back-to-back frames, of network devices. An RFC 2544-

based benchmarking test is performed by transmitting test packets from a device that functions as the initiator and terminator of the test. These packets are sent to a device that functions as the reflector, which receives and returns the packets back to the initiator.

NOTE: The test configuration is applied only when you start the test. If you update the test configuration during the test, you have to start the test again for the updated configuration to take effect.

The following devices support RFC 2544-based benchmarking tests in either the initiator/terminator or reflector role, according to which families they support:

Table 128: RFC 2544-Based Benchmarking Tests by Role and Family Supported

Platform	family			
	inet	bridge	ccc	vpls
Initiator and Terminator Role				
ACX Series (except for ACX5000 and ACX7000)	x	x	x	
Reflector Role				
ACX Series (except for ACX5000 and ACX7000)	x	x	x	
ACX5000 Series		x	x	
ACX7000 Series	x			
MX Series	x	x	x	x
SRX300 Series and SRX550HM	x			

The family type for the test is configured with the `family name` statement at the `[edit services rpm rfc2544-benchmarking tests test-name name]` hierarchy level.

You must configure a test profile and reference the test profile in a unique test name that defines the parameters for the test to be performed on a certain device. However, the test profile is required when the test mode is configured as initiation and termination. The test-profile parameter is disregarded when the test mode is configured as reflection. MX Series routers and SRX devices support only the reflection function in the RFC 2544-based benchmarking tests. A reflection service does not use the parameters specified in the test profile.

NOTE: To configure RFC2544-based benchmarking tests on MX240, MX480, MX960 Series routers with MPC1, MPC2, and the 16-port 10-Gigabit Ethernet MPC, see ["Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers" on page 747](#).

This chapter describes how to configure a test name for an RFC 2544-based benchmarking test on an MX Series router for Layer 3 IPv4, Ethernet pseudowire, and Layer 2 bridge networks. For SRX devices, you can only configure Layer 3 IPv4 reflection (family inet only).

Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a IPv4 Network

You can configure a test name by including the test-name *test-name* statement at the [edit services rpm rfc2544-benchmarking] hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (type of service, IPv4 or Ethernet), the logical interface, and test duration that are used for a benchmarking test to be run.

To configure a test name and define its attributes for an IPv4 network:

1. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure a instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```


4. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

5. Specify the test mode for the packets that are sent during the benchmarking test. The reflect option causes the test frames to be reflected on the IPv4 network.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

6. Configure the address type family for the benchmarking test. The inet option indicates that the test is run on an IPv4 service.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```

7. Configure the destination IPv4 address for the test packets. This parameter is required only if you configure IPv4 family `inet`. If you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-ipv4-address address
```

8. Specify the UDP port of the destination to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-udp-port port-number
```

9. (Optional) Specify the source IPv4 address to be used in generated test frames. If you do not configure the source IPv4 address for `inet` family, the source address of the interface is used to transmit the test frames.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address address
```

10. Specify the UDP port of the source to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-udp-port port-number
```

11. Specify the logical interface on which the RFC 2544-based benchmarking test is run. If you configure an inet family and the test mode to reflect the frames back on the sender from the other end, then the logical interface is used as the interface to enable the reflection service (reflection is performed on the packets entering the specified interface). If you not configure the logical interface for reflection test mode, then a lookup is performed on the source IPv4 address to determine the interface that hosts the address.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface interface-name
```

Configuring a Test Name for an RFC 2544-Based Benchmarking Test for an Ethernet Pseudowire

You can configure a test name by including the `test-name test-name` statement at the `[edit services rpm rfc2544-benchmarking]` hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (type of service, IPv4 or Ethernet), the logical interface, and test duration, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

To configure a test name and define its attributes for an Ethernet Pseudowire:

1. In configuration mode, go to the `[edit services]` hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure an RPM service instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

5. Specify the test mode for the packets that are sent during the benchmarking test. The `reflect` option causes the test frames to be reflected on the Ethernet pseudowire.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

6. Configure the address type family for the benchmarking test. The `ccc` option indicates that the test is run on a CCC or Ethernet pseudowire service.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

7. Specify the direction of the interface on which the test must be run. This parameter is valid only for a family. To enable the test to be run in the egress direction of the interface (network-to-network interface (NNI)), use the `egress` option. To enable the test to be run in the ingress direction of the interface (user-to-network interface (UNI)), use the `ingress` option.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction (egress | ingress)
```

8. (Optional) Specify the source IPv4 address to be used in generated test frames. If you do not configure the source IPv4 address for family, the default value of 192.168.1.10 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address address
```

9. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface interface-name
```

Configuring a Test Name for an RFC 2544-Based Benchmarking Test for a Layer 2 E-LAN Service in Bridge Domain

You can configure a test name by including the `test-name test-name` statement at the `[edit services rpm rfc2544-benchmarking]` hierarchy level. In the test name, you can configure attributes of the test iteration, such as the address family (bridge), the logical interface, and test duration, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

To configure a test name and define its attributes for a layer 2 E-LAN service in Bridge domains:

1. In configuration mode, go to the `[edit services]` hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure an RPM service instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for the test—for example, `l2b-test1`. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name l2b-test1
```

5. Specify the source and destination MAC addresses of the test packet. Both these parameters are valid only for the bridge family.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set source-mac-address address destination-mac-address address
```

6. Specify the service type under test. This parameter is applicable only for the bridge family.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set service-type elan
```

7. Specify the test mode for the packets that are sent during the benchmarking test. The reflect option causes the test frames to be reflected over the Layer 2 bridge.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set mode reflect
```

8. Configure the address type family for the benchmarking test. The bridge option indicates that the test is run on a E-LAN service over a bridge domain.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set family bridge
```

9. Specify the direction of the interface on which the test must be run. This parameter is valid only for a family. To enable the test to be run in the egress direction of the interface (network-to-network interface (NNI)), use the `egress` option. To enable the test to be run in the ingress direction of the interface (user-to-network interface (UNI)), use the `ingress` option.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set direction egress
```

10. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-test1]
user@host# set test-interface interface-name
```

RELATED DOCUMENTATION

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for UNI Direction of Ethernet Pseudowires | 763](#)

[Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for NNI Direction of Ethernet Pseudowires | 775](#)

[Example: Configuring RFC2544-Based Benchmarking Tests on an MX104 Router for Layer 2 E-LAN Services in Bridge Domains | 788](#)

[Example: Configure an RFC 2544-Based Benchmarking Test on an MX104 Router for Layer 3 IPv4 Services | 749](#)

Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers

NOTE: MX Series routers support only the reflector function in RFC 2544-based benchmarking tests.

RFC 2544 defines a series of tests that can be used to describe the performance characteristics of a network-interconnecting device, such as a router, and outlines specific formats to report the results of the tests. These tests can be used to benchmark interconnected network devices and devise a guideline or a measurement pattern to analyze the health and efficiency of the network devices. These tests are known as RFC 2544-based benchmarking tests and are supported on MX80, MX104, MX240, MX480, MX960, and MX2010 routers with MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP). Starting from Junos OS Release 17.1R1, the RFC 2544-based benchmarking tests are supported on MX Series routers with MPC3E (MX-MPC3E-3D), MPC3E-NG (MX-MPC3E-3D-NG), MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE), MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, and MPC5EQ-100G10G), and MPC6E (MX2K-MPC6E). Starting from Junos OS Release 20.2R1, the RFC 2544-based benchmarking tests are supported on MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line card and MX2010 and MX2020 routers with MX2K-MPC11E line card.

Junos OS Release 20.3R1 extends support for the RFC 2544-based benchmarking tests onto the MX240, MX480, and MX960 routers with the MPC7E-MRATE or MPC7E-10G line card, onto the MX2008, MX2010, and MX2020 routers with the MX2K-MPC8E or MX2K-MPC9E line card, and onto the MX204 and MX10003 (with the LC2103 card) routers.

NOTE: On MX104 and MX80 Series routers that have a single fixed FPC, this configuration is not required.

To enable support for RFC 2544-based benchmarking tests on MX Series routers:

1. In configuration mode, go to the `[edit chassis fpc fpc-slot-number]` hierarchy level.

```
[edit]
user@host# edit chassis fpc fpc-slot-number
```

2. Enable support for service-level agreement (SLA) monitoring services and RFC-based benchmarking tests:

```
[edit chassis fpc fpc-slot-number]
user@host# set slamon-services rfc2544
```

Release History Table

Release	Description
20.3R1	Junos OS Release 20.3R1 extends support for the RFC 2544-based benchmarking tests onto the MX240, MX480, and MX960 routers with the MPC7E-MRATE or MPC7E-10G line card, onto the MX2008, MX2010, and MX2020 routers with the MX2K-MPC8E or MX2K-MPC9E line card, and onto the MX204 and MX10003 (with the LC2103 card) routers.
20.2R1	Starting from Junos OS Release 20.2R1, the RFC 2544-based benchmarking tests are supported on MX240, MX480, and MX960 routers with MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line card and MX2010 and MX2020 routers with MX2K-MPC11E line card.
17.1R1	Starting from Junos OS Release 17.1R1, the RFC 2544-based benchmarking tests are supported on MX Series routers with MPC3E (MX-MPC3E-3D), MPC3E-NG (MX-MPC3E-3D-NG), MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE), MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, and MPC5EQ-100G10G), and MPC6E (MX2K-MPC6E).

RELATED DOCUMENTATION

- [Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)
- [Configuring an RFC 2544-Based Benchmarking Test | 739](#)

Example: Configure an RFC 2544-Based Benchmarking Test on an MX104 Router for Layer 3 IPv4 Services

IN THIS SECTION

- Requirements | 749
- Overview | 750
- Configuration | 750
- Verify the Results of the Benchmarking Test for Layer 3 IPv4 Services | 762

Requirements

NOTE: MX Series routers support only the reflector function in RFC2544-based benchmarking tests. This example uses the MX104 3D Universal Edge Router as the reflector. You can also configure benchmarking tests on MX80 routers and on MX240, MX480, and MX960 routers with MPC1, MPC2, and 16-port 10-Gigabit Ethernet MPC from Junos OS Release 16.1 or later. To configure RFC2544-based benchmarking tests on MX240, MX480, MX960 routers, see ["Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers" on page 747](#).

NOTE: This example is not applicable for ACX7100, ACX5448, ACX5048, and ACX5096 routers because they can only be configured as reflectors, not initiators.

This example uses the following hardware and software components:


- An MX104 router (reflector)
- An ACX Series router (initiator)
- Junos OS Release 13.3 or later

Overview

Consider a sample topology in which a router, Router A (ACX), functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A is connected over a Layer 3 network to another router, Router B (MX104), which functions as a reflector to reflect back the test frames it receives from Router A. IPv4 is used for transmission of test frames over the Layer 3 network. This benchmarking test is used to compute the IPv4 service parameters between Router A and Router B. Logical interfaces on both the routers are configured with IPv4 addresses to measure the performance attributes, such as throughput, latency, frame loss, and bursty frames, of network devices for the IPv4 service.

Figure 63 on page 750 shows the sample topology to perform an RFC 2544 test for a Layer 3 IPv4 Service.

Figure 63: RFC 2544-Based Benchmarking Test for a Layer 3 IPv4 Service



Configuration

IN THIS SECTION

- CLI Quick Configuration | 751
- Configure Benchmarking Test Parameters on Router B | 752
- Configure Benchmarking Test Parameters on Router A | 755
- Results | 760

In this example, you configure the benchmarking test for a Layer 3 IPv4 service that is between interface ge-0/0/0 on Router A and interface ge-0/0/4 on Router B to detect and analyze the performance of the interconnecting routers. You do not configure a test profile on Router B, because it operates as a reflector. You must configure the reflector (Router B) before you configure the initiator (Router A), because the reflector needs to be already configured and the tests running before you start tests on the initiator. If you start the tests on the initiator first, then all the packets sent are lost until you start the tests on the reflector.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configure Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 unit 0 family inet address 192.0.2.2/24
set services rpm rfc2544-benchmarking tests test-name test1 test-interface ge-0/0/4.0
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 destination-ipv4-address 192.0.2.2
set services rpm rfc2544-benchmarking tests test-name test1 destination-udp-port 4001
set services rpm rfc2544-benchmarking tests test-name test1 source-ipv4-address 192.0.2.1
```

Configure Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.1/24
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 1000
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 test-interface ge-0/0/0.0
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 destination-ipv4-address 192.0.2.2
set services rpm rfc2544-benchmarking tests test-name test1 destination-udp-port 4001
set services rpm rfc2544-benchmarking tests test-name test1 source-ipv4-address 192.0.2.1
```

Configure Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following 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 the test parameters on Router B:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

```
[edit]
user@RouterB# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@RouterB# edit ge-0/0/4
```

3. Configure a logical unit and specify the protocol family as inet.

```
[edit interfaces ge-0/0/4]
user@RouterB# edit unit 0 family inet
```

4. Specify the address for the logical interface.

```
[edit interfaces ge-0/0/4 unit 0 family inet]
user@RouterB# set address 192.0.2.2/24
```

5. Go to the top level of the configuration mode.

```
[edit interfaces ge-0/0/4 unit 0]
user@RouterB# top
```

6. In configuration mode, go to the [edit services rpm rfc2544-benchmarking] hierarchy level.

```
[edit]
user@RouterB# edit services rpm rfc2544-benchmarking
```

7. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@RouterB# edit tests test-name test1
```

8. Specify the logical interface, ge-0/0/4.0, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set test-interface ge-0/0/4.0
```

9. Specify reflect as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set mode reflect
```

10. Configure the address type family, inet, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set family inet
```

11. Configure the destination IPv4 address for the test packets as 192.0.2.2. The destination IPv4 address configured on the reflector must match the destination IPv4 address configured on the initiator. If you configure 192.0.2.1 instead, you get this error message: error: test test1 - Could not determine local interface for address 192.0.2.1.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-ipv4-address 192.0.2.2
```

12. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set destination-udp-port 4001
```

13. Configure the source IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address 192.0.2.1
```

14. Go to the top level of the configuration mode.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# top
```

15. Commit the configuration.

```
[edit]
user@RouterB# commit
```

16. Confirm the configuration. If the output does not contain the configuration below, repeat the configuration instructions in this example to correct it.

```
[edit interfaces]
ge-0/0/4 {
  unit 0 {
    family inet {
      address 192.0.2.2/24;
    }
  }
}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name test1 {
```

```

        test-interface ge-0/0/4.0;
        mode reflect;
        family inet;
        destination-ipv4-address 192.0.2.2;
        destination-udp-port 4001;
        source-ipv4-address 192.0.2.1
    }
}
}

```

17. Exit to operational mode.

```

[edit]
user@RouterB# exit
user@RouterB>

```

18. Start the benchmarking test on the reflector.

```

user@host> test services rpm rfc2544-benchmarking test test1 start

```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the `test services rpm rfc2544-benchmarking test test1 stop` command.

Configure Benchmarking Test Parameters on Router A

Step-by-Step Procedure

The following 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 the test parameters on Router A:

1. In configuration mode, go to the `[edit interfaces]` hierarchy level:

```

[edit]
user@RouterA# edit interfaces

```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@RouterA# edit ge-0/0/0
```

3. Configure a logical unit and specify the protocol family.

```
[edit interfaces ge-0/0/0]
user@RouterA# edit unit 0 family inet
```

4. Specify the address for the logical interface.

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@RouterA# set address 192.0.2.1/24
```

5. Go to the top level of the configuration mode.

```
[edit interfaces ge-0/0/0 unit 0]
user@RouterA# top
```

6. In configuration mode, go to the [edit services rpm rfc2544-benchmarking] hierarchy level.

```
[edit]
user@RouterA# edit services rpm rfc2544-benchmarking
```

7. Define a name for a test profile—for example, throughput.

```
[edit services rpm rfc2544-benchmarking]
user@RouterA# edit profiles test-profile throughput
```

8. Configure the type of test to be performed as throughput.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# set test-type throughput
```

9. Specify the size of the test packet as 64 bytes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# set packet-size 64
```

10. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1,000 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# set bandwidth-kbps 1000
```

11. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# up
```

12. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@RouterA# up
```

13. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@RouterA# edit tests test-name test1
```

14. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set test-profile throughput
```


15. Specify the logical interface, ge-0/0/0.0, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set test-interface ge-0/0/0.0
```

16. Specify the test mode for the packets that are sent during the benchmarking test as initiate and terminate.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set mode initiate-and-terminate
```

17. Configure the address type family, inet, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set family inet
```

18. Configure the destination IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set destination-ipv4-address 192.0.2.2
```

19. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set destination-udp-port 4001
```

20. Configure the source IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set source-ipv4-address 192.0.2.1
```

21. Go to the top level of the configuration mode.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# top
```

22. Commit the configuration.

```
[edit]
user@RouterA# commit
```

23. Confirm the configuration. If the output does not contain the configuration below, repeat the configuration instructions in this example to correct it.

```
[edit]
user@RouterA# show
[edit interfaces]
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.0.2.1/24;
      }
    }
  }

[edit services rpm]
  rfc2544-benchmarking {
    profiles {
      test-profile throughput {
        test-type throughput
        packet-size 64;
        bandwidth-kbps 1000;
      }
    }

    tests {
      test-name test1 {
        test-profile throughput;
        interface ge-0/0/0.0;
        mode initiate-and-terminate;
        family inet;
```

```

        destination-ipv4-address 192.0.2.2
        destination-udp-port 4001;
        source-ipv4-address 192.0.2.1
    }
}
}

```

24. Exit to operational mode.

```

[edit]
user@RouterA# exit
user@RouterA>

```

25. Start the benchmarking test on the initiator.

```

user@RouterA> test services rpm rfc2544-benchmarking test test1 start

```

After the test successfully completes, it automatically stops at the initiator. Once the test is successfully completed at the initiator, you can stop the test at the reflector by entering the `test services rpm rfc2544-benchmarking test test1 stop` command on Router B in operational mode.

Results

If you have not done so already, confirm your configuration on Router A and Router B by entering the `show` command in configuration mode at the `[edit interfaces]` and `[edit services rpm]` hierarchy levels. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Configuration for Benchmarking Test Parameters on Router A:

```

[edit interfaces]
ge-0/0/0 {
    unit 0 {
        family inet {
            address 192.0.2.1/24;
        }
    }
}

[edit services rpm]

```

```

rfc2544-benchmarking {
  profiles {
    test-profile throughput {
      test-type throughput
      packet-size 64;
      bandwidth-kbps 1000;
    }
  }

  tests {
    test-name test1 {
      test-profile throughput;
      test-interface ge-0/0/0.0;
      mode initiate-and-terminate;
      family inet;
      destination-ipv4-address 192.0.2.2
      destination-udp-port 4001;
      source-ipv4-address 192.0.2.1
    }
  }
}

```

Configuration for Benchmarking Test Parameters on Router B:

```

[edit interfaces]
ge-0/0/4 {
  unit 0 {
    family inet {
      address 192.0.2.2/24;
    }
  }
}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name test1 {
      test-interface ge-0/0/4.0;
      mode reflect;
      family inet;
    }
  }
}

```

```

        destination-ipv4-address 192.0.2.2;
        destination-udp-port 4001;
        source-ipv4-address 192.0.2.1
    }
}

```

Verify the Results of the Benchmarking Test for Layer 3 IPv4 Services

IN THIS SECTION

- [Verify the Benchmarking Test Results | 762](#)

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

Verify the Benchmarking Test Results

Purpose

Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests | summary)` command, on either the initiator or the reflector, to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.

RELATED DOCUMENTATION

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for UNI Direction of Ethernet Pseudowires

IN THIS SECTION

- Requirements | 763
- Overview | 763
- Configuration | 764
- Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service | 774

This example shows how to configure the benchmarking test for the user-to-network interface (UNI) direction of an Ethernet pseudowire service.

Requirements

NOTE: MX Series routers support only the reflector function in RFC2544-based benchmarking tests. This example uses the MX104 3D Universal Edge Router as the reflector. You can also configure benchmarking tests on MX80 Series routers and MX240, MX480, and MX960 Series routers with MPC1, MPC2, and 16-port 10-gigabit Ethernet MPC from Junos OS Release 16.1 or later. To configure RFC2544-based benchmarking tests on MX240, MX480, MX960 Series routers, see ["Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers" on page 747](#).

This example uses the following hardware and software components:

- An MX104 (reflector)
- An ACX Series router (initiator)
- Junos OS Release 13.3 or later

Overview


Consider a sample topology in which a router, Router A (MX104), functions as a reflector of the test frames for an RFC 2544-based benchmarking test. The logical customer edge (CE)-facing interface and inet family are configured on Router A. Router A is not part of a pseudowire and therefore, a Layer 3 family configuration is required on it. Router A, which is a customer edge device CE1 is connected to Router B (ACX), which functions as a provider edge device PE1 over an Ethernet pseudowire in the UNI

direction with EtherType or Layer 2 Ethernet payload. The logical interface, family, and UNI direction are configured on Router B. Router B or PE1 is connected over an Ethernet pseudowire in the NNI direction to a provider edge device at the remote site, PE2. The link between CE1 and PE1 is an Ethernet Layer 2 network and it can be configured with any EtherType value. The link between PE1 and PE2 is an Ethernet line (E-Line) or an Ethernet Private Line (EPL) that has Layer 2 payload and Layer 3 transport sent over it. Router B or PE1 functions as an initiator and terminator of the test frames that are sent to Router A and reflected back from it.

This benchmarking test is used to compute the performance attributes in the user-to-network interface (UNI) direction of an Ethernet pseudowire service between Router A and Router B. Data traffic arriving from a network-to-network interface (NNI) toward the customer edge is ignored while the test is in progress. Packets from the CE are not sent toward the NNI because all packets are assumed to be test probes.

Figure 64 on page 764 shows the sample topology to perform an RFC 2544 test for the UNI direction of an Ethernet pseudowire service.

Figure 64: RFC 2544-Based Benchmarking Test for UNI Direction of an Ethernet Pseudowire



Configuration

IN THIS SECTION

- CLI Quick Configuration | 765
- Configuring Benchmarking Test Parameters on Router A | 766
- Configuring Benchmarking Test Parameters on Router B | 770

In this example, you configure the benchmarking test for the UNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 vlan-id 101
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.1/24
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration 20m
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate,terminate
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 dest-address 192.0.2.2
set services rpm rfc2544-benchmarking tests test-name test1 udp-port 4001
```

Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
```



```
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction uni
```

Configuring Benchmarking Test Parameters on Router A

Step-by-Step Procedure

The following require 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 the test parameters on Router A:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/0
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
```

4. Configure a logical unit and specify the protocol family as inet.

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0 family inet
```

5. Specify the address for the logical interface.

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# set address 192.0.2.1/24
```

6. Configure the VLAN ID on the logical interface as 101.

```
[edit interfaces ge-0/0/0 unit 0]  
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/0 unit 0]  
user@host# top
```

8. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]  
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]  
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]  
user@host# edit rfc2544-benchmarking
```

11. Define a name for a test profile—for example, throughput.

```
[edit services rpm rfc2544-benchmarking]  
user@host# edit profiles test-profile throughput
```

12. Configure the type of test to be performed as throughput.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# set test-type throughput
```

13. Specify the size of the test packet as 64 bytes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type packet-size 64
```

14. Specify the period for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds). In this example, you configure the period as 20 minutes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```

15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```

16. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# up
```

17. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```

18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set test-profile throughput
```

20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, inet, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set family inet
```

23. Configure the destination IPv4 address for the test packets as 192.0.2.2.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set dest-address 192.0.2.2
```

24. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set udp-port 4001
```

Configuring Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following require 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 the test parameters on Router B:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID as 101 on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# top
```

8. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]  
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]  
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]  
user@host# edit rfc2544-benchmarking
```

11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]  
user@host# edit tests test-name test1
```

12. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]  
user@host# set test-interface ge-0/0/4.1
```

13. Specify reflect as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

14. Configure the address type family, ccc, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

15. Specify the direction of the interface on which the test must be run, which is UNI in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction uni
```

16. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the test services rpm rfc2544-benchmarking test test1 stop command.

Results

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

Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
  vlan-tagging;
  unit 0 {
    vlan-id 101;
    family inet {
      address 192.0.2.1/24;
    }
  }
}
```

```

    }
}

[edit services rpm]
rfc2544-benchmarking {
  profiles {
    test-profile throughput {
      test-type throughput
      packet-size 64;
      test-duration 20m;
      bandwidth-kbps 500;
    }
  }

  tests {
    test-name test1 {
      interface ge-0/0/0.1;
      test-profile throughput;
      mode initiate,terminate;
      family inet;
      dest-address 192.0.2.2
      udp-port 4001;
    }
  }
}

```

Benchmarking Test Parameters on Router B:

```

[edit interfaces]
ge-0/0/4 {
  vlan-tagging;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 101;
  }
}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name test1 {

```



```

        interface ge-0/0/4.1;
        mode reflect;
        family ccc;
        direction uni;
    }
}

```

After you have configured the device, enter the `commit` command in configuration mode.

Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service

IN THIS SECTION

- [Verifying the Benchmarking Test Results | 774](#)

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

Verifying the Benchmarking Test Results

Purpose

Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the `show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` in the [CLI Explorer](#).

RELATED DOCUMENTATION

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

Example: Configuring an RFC 2544-Based Benchmarking Test on an MX104 Router for NNI Direction of Ethernet Pseudowires

IN THIS SECTION

- [Requirements | 775](#)
- [Overview | 776](#)
- [Configuration | 777](#)
- [Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service | 787](#)

This example shows how to configure the benchmarking test for a network-to-network interface (NNI) direction of an Ethernet pseudowire service.

Requirements

NOTE: MX Series routers support only the reflector function in RFC2544-based benchmarking tests. This example uses the MX104 3D Universal Edge Router as the reflector. You can also configure benchmarking tests on MX80 Series routers and MX240, MX480, and MX960 Series routers with MPC1, MPC2, and 16-port 10-gigabit Ethernet MPC from Junos OS Release 16.1 or later. To configure RFC2544-based benchmarking tests on MX240, MX480, MX960 Series

routers, see ["Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers" on page 747.](#)

This example uses the following hardware and software components:

- An MX104 (reflector)
- An ACX Series router (initiator)
- Junos OS Release 13.3 or later

Overview

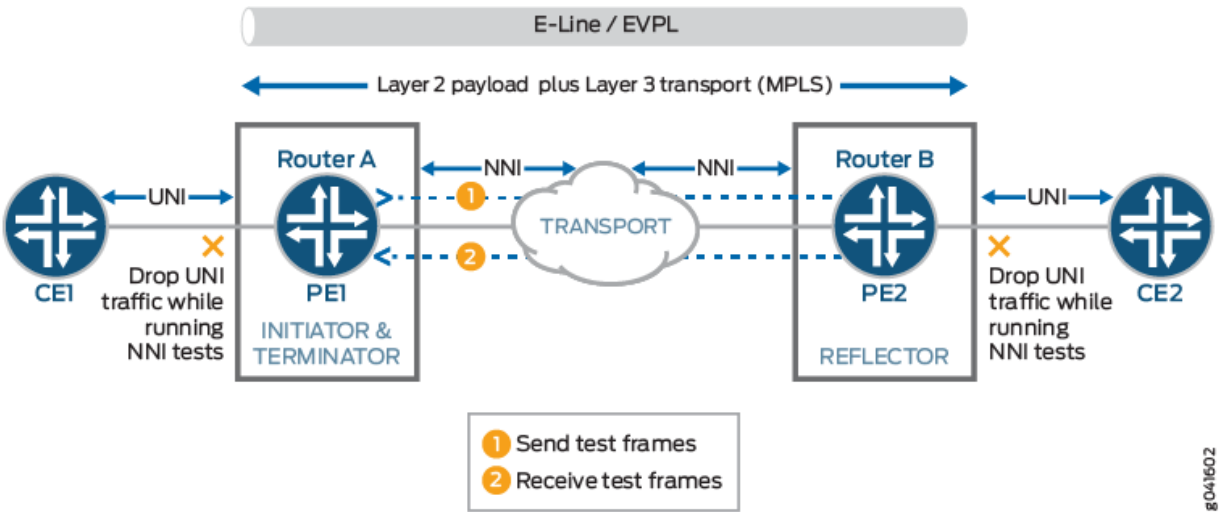
Consider a sample topology in which a router, Router A (ACX), functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A operates as a provider edge device PE1, which is connected to a customer edge device CE1 on one side and over an Ethernet pseudowire to another router Router B (MX104), which functions as a reflector to reflect back the test frames it receives from Router A. Router B operates as a provider edge device, PE2, which is the remote router located at the other side of the service provider core. The UNI direction of CE1 is connected to the NNI direction of PE1. An MPLS tunnel connects PE1 and PE2 over the Ethernet pseudowire or the Ethernet line (E-Line).

NOTE: When pseudowire reflection is enabled on an interface, the router does not block the ingress or egress traffic through the test interface. To block other data traffic, you must explicitly configure firewall filters.

This benchmarking test is used to compute the performance attributes in the network-to-network interface (NNI) direction of an Ethernet pseudowire service between Router A and Router B. The logical interface under test on Router A is the CE1 interface with UNI as the direction, and the logical interface under test on Router B is the CE2 interface with NNI as the direction. Data traffic arriving from UNI toward NNI is ignored while the test is in progress. Packets from NNI are not sent toward the customer edge because all packets are assumed to be test frames. The family and NNI direction are configured on routers A and B.

Figure 65 on page 777 shows the sample topology to perform an RFC 2544 test for the NNI direction of an Ethernet pseudowire service.

Figure 65: RFC 2544-Based Benchmarking Test for NNI Direction of an Ethernet Pseudowire



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 777](#)
- [Configuring Benchmarking Test Parameters on Router | 778](#)
- [Configuring Benchmarking Test Parameters on Router B | 782](#)
- [Results | 785](#)

In this example, you configure the benchmarking test for the NNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/0 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration 20
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name test1 family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction egress
```

Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction ingress
```

Configuring Benchmarking Test Parameters on Router

Step-by-Step Procedure

The following require 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 the test parameters on Router A:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/0
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```

8. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

11. Define a name for a test profile—for example, throughput.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile throughput
```

12. Configure the type of test to be performed as throughput.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type throughput
```

13. Specify the size of the test packet as 64 bytes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type packet-size 64
```

14. Specify the period—for example, 20 minutes—for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds).

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```

15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```

16. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# up
```

17. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```

18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-profile throughput
```


20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, ccc, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

23. Specify the direction of the interface on which the test must be run, which is egress in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction egress
```

Configuring Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following require 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 the test parameters on Router B:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]  
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]  
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]  
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/4 unit 0]  
user@host# top
```

8. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]  
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

12. Specify the logical interface, ge-0/0/4.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/4.1
```

NOTE: When pseudowire reflection is enabled on an interface, the router does not block the ingress or egress traffic through the test interface. To block other data traffic, you must explicitly configure firewall filters.

13. Specify reflect as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

14. Configure the address type family, ccc, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

15. Specify the direction of the interface on which the test must be run, which is ingress in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction ingress
```

16. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test test1 start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the `test services rpm rfc2544-benchmarking test test1 stop` command.

Results

In configuration mode, confirm your configuration on Router A and Router B by entering the `show` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
  vlan-tagging;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 101;
  }
}

[edit services rpm]
rfc2544-benchmarking {
  profiles {
    test-profile throughput {
      test-type throughput
      packet-size 64;
      test-duration 20m;
      bandwidth-kbps 500;
    }
  }
}
```

```

tests {
    test-name test1 {
        interface ge-0/0/0.1;
        test-profile throughput;
        mode initiate-and-terminate;
        family ccc;
        direction egress;
    }
}

```

Benchmarking Test Parameters on Router B:

```

[edit interfaces]
ge-0/0/4 {
    vlan-tagging;
    unit 0 {
        encapsulation vlan-ccc;
        vlan-id 101;
    }
}

[edit services rpm]
rfc2544-benchmarking {
    # Note, When in reflector mode, test profile is not needed
    tests {
        test-name test1 {
            interface ge-0/0/4.1;
            mode reflect;
            family ccc;
            direction egress;
        }
    }
}

```

After you have configured the device, enter the `commit` command in configuration mode.

Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service

IN THIS SECTION

- [Verifying the Benchmarking Test Results | 787](#)

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

Verifying the Benchmarking Test Results

Purpose

Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the `show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` in the [CLI Explorer](#).

RELATED DOCUMENTATION

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)
[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

Example: Configuring RFC2544-Based Benchmarking Tests on an MX104 Router for Layer 2 E-LAN Services in Bridge Domains

IN THIS SECTION

- Requirements | 788
- Overview | 788
- Configuration | 790
- Verifying the Results of the Benchmarking Tests for Layer 2 Services (E-LAN) in Bridge Domains | 810

This example shows how to configure benchmarking tests for the Layer 2 E-LAN services in bridge domains. The example covers the four basic tests: throughput, frame-loss, back-to-back, and latency.

Requirements

NOTE: MX Series routers support only the reflector function in RFC2544-based benchmarking tests. This example uses the MX104 3D Universal Edge Router as the reflector. You can also configure benchmarking tests on MX80 Series routers and MX240, MX480, and MX960 Series routers with MPC1, MPC2, and 16-port 10-gigabit Ethernet MPC from Junos OS Release 16.1 or later. To configure RFC2544-based benchmarking tests on MX240, MX480, MX960 Series routers, see ["Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers" on page 747](#).

This example uses the following hardware and software components:

- An MX104 (reflector)
- An ACX Series router (initiator)
- Junos OS Release 14.2 or later for MX Series routers


Overview

Consider a sample topology in which an ACX Series router functions as an initiator and terminator of the test frames for an RFC2544-based benchmarking test. ACX Series router is connected to a customer edge device CE1, on one side and is connected over a Layer 2 network to an MX104 Series router. The MX104 Series router functions as a reflector to reflect the test frames it receives from the ACX Series initiator back to the initiator. The MX04 Series router is also connected to a customer edge device CE2.

NOTE: When Layer 2 reflection is enabled on an interface, filters are configured internally to block the ingress and egress traffic except test traffic through the test interface.

Figure 66 on page 789 shows the sample topology to perform all four RFC2544-based benchmarking tests (throughput, back-to-back frames, latency, and frame-loss) for the UNI direction on a Layer 2 bridge network.

Figure 66: Layer 2 Reflection Simple Topology



On the ACX Series router, ge-1/2/1.0 is the Layer 2 NNI interface and ge-1/1/3.0 is the Layer 2 UNI interface. On the MX104 Series router, ge-1/1/6.0 is the Layer 2 NNI interface and ge-1/1/5.0 is the Layer 2 UNI interface. The benchmarking tests are used to compute the performance attributes for an E-LAN service on a bridge domain.

NOTE: Test packets can be identified using the destination MAC address, source MAC address, and test interface. Both tagged and untagged interfaces are supported. For tagged interfaces, the test interface is the VLAN sub interface. For untagged interfaces, the physical port represents the test interface. Traffic through other VLAN sub interfaces, present in the same physical port, is not affected when you configure the benchmarking test on one of the sub interfaces.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 790](#)
- [Configuring Throughput Benchmarking Test Parameters on the ACX Series Router | 793](#)
- [Configuring Back-to-Back Frames Benchmarking Test Parameters on the ACX Series Router | 795](#)
- [Configuring Latency Benchmarking Test Parameters on the ACX Series Router | 797](#)
- [Configuring Frame Loss Benchmarking Test Parameters on the ACX Series Router | 799](#)
- [Configuring Other Benchmarking Test Parameters on the ACX Series Router | 802](#)
- [Configuring Benchmarking Test Parameters on the MX104 Router | 803](#)
- [Configuring Other Benchmarking Test Parameters on the MX104 Router | 804](#)
- [Results | 806](#)

In this example, you configure the benchmarking tests for the UNI direction for an E-LAN service on a Layer 2 bridge domain that is enabled between two routers to detect and analyze the performance of the interconnected routers. In this example, we start by configuring the ACX Series router. On the ACX Series router, you first configure each test by specifying the test profile, the test attributes, and then define the test by associating the test with the test profile with the relevant attributes. You can then configure the interface. On the MX104 Series router, you perform the same steps. However, a few attributes such as the outer VLAN ID, source UDP port, destination UDP port, the duration of each iteration, and their values are only applicable to the initiator or the ACX Series router.

NOTE: When you configure the Layer 2 reflection, you can specify the service type under test as ELINE if you want to simulate an ELINE service using bridge encapsulation.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configuring Benchmarking Test Parameters on the ACX Series Router

```

set services rpm rfc2544-benchmarking profiles test-profile tput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile tput packet-size 128
set services rpm rfc2544-benchmarking profiles test-profile tput bandwidth-kbps 900000
set services rpm rfc2544-benchmarking profiles test-profile b2bt test-type back-back-frames
set services rpm rfc2544-benchmarking profiles test-profile b2bt packet-size 512
set services rpm rfc2544-benchmarking profiles test-profile b2bt bandwidth-kbps 950000
set services rpm rfc2544-benchmarking profiles test-profile lty test-type latency
set services rpm rfc2544-benchmarking profiles test-profile lty packet-size 512
set services rpm rfc2544-benchmarking profiles test-profile lty bandwidth-kbps 1000000
set services rpm rfc2544-benchmarking profiles test-profile frloss test-type frame-loss
set services rpm rfc2544-benchmarking profiles test-profile frloss packet-size 1600
set services rpm rfc2544-benchmarking profiles test-profile frloss bandwidth-kbps 1000000
set services rpm rfc2544-benchmarking tests test-name tput-test test-profile tput
set services rpm rfc2544-benchmarking tests test-name tput-test source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name tput-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name tput-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name tput-test service-type elan
set services rpm rfc2544-benchmarking tests test-name tput-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name tput-test family bridge
set services rpm rfc2544-benchmarking tests test-name tput-test direction egress
set services rpm rfc2544-benchmarking tests test-name tput-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name tput-test destination-udp-port 200
set services rpm rfc2544-benchmarking tests test-name tput-test test-iterator-duration 20
set services rpm rfc2544-benchmarking tests test-name tput-test test-interface ge-1/1/3.0
set services rpm rfc2544-benchmarking tests test-name b2b-test test-profile b2bt
set services rpm rfc2544-benchmarking tests test-name b2b-test source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name b2b-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name b2b-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name b2b-test service-type elan
set services rpm rfc2544-benchmarking tests test-name b2b-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name b2b-test family bridge
set services rpm rfc2544-benchmarking tests test-name b2b-test direction egress
set services rpm rfc2544-benchmarking tests test-name b2b-test test-iterator-duration 20
set services rpm rfc2544-benchmarking tests test-name b2b-test test-interface ge-1/1/3.0
set services rpm rfc2544-benchmarking tests test-name lty-test test-profile lty
set services rpm rfc2544-benchmarking tests test-name lty-test source-mac-address

```

```

00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name lty-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name lty-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name lty-test service-type elan
set services rpm rfc2544-benchmarking tests test-name lty-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name lty-test family bridge
set services rpm rfc2544-benchmarking tests test-name lty-test direction egress
set services rpm rfc2544-benchmarking tests test-name lty-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name lty-test destination-udp-port 200
set services rpm rfc2544-benchmarking tests test-name lty-test test-iterator-duration 20
set services rpm rfc2544-benchmarking tests test-name lty-test test-interface ge-1/1/3.0
set services rpm rfc2544-benchmarking tests test-name frloss-test test-profile frloss
set services rpm rfc2544-benchmarking tests test-name frloss-test source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name frloss-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name frloss-test ovlan-id 400
set services rpm rfc2544-benchmarking tests test-name frloss-test service-type elan
set services rpm rfc2544-benchmarking tests test-name frloss-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name frloss-test family bridge
set services rpm rfc2544-benchmarking tests test-name frloss-test direction egress
set services rpm rfc2544-benchmarking tests test-name frloss-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name frloss-test destination-udp-port 200
set services rpm rfc2544-benchmarking tests test-name frloss-test test-iterator-duration 20
set services rpm rfc2544-benchmarking tests test-name frloss-test test-interface ge-1/1/3.0
set interfaces ge-1/2/1 flexible-vlan-tagging
set interfaces ge-1/2/1 mtu 9192
set interfaces ge-1/2/1 encapsulation flexible-ethernet-services
set interfaces ge-1/2/1 unit 0 encapsulation vlan-bridge
set interfaces ge-1/2/1 unit 0 vlan-id 400
set interfaces ge-1/1/3 flexible-vlan-tagging
set interfaces ge-1/1/3 mtu 9192
set interfaces ge-1/1/3 encapsulation flexible-ethernet-services
set interfaces ge-1/1/3 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/3 unit 0 vlan-id 400
set bridge-domains bd1 vlan-id 600
set bridge-domains bd1 interface ge-1/2/1.0
set bridge-domains bd1 interface ge-1/1/3.0

```

Configuring Benchmarking Test Parameters on the MX104 Router

```

set services rpm rfc2544-benchmarking tests test-name l2b-reflector source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name l2b-reflector destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name l2b-reflector service-type elan
set services rpm rfc2544-benchmarking tests test-name l2b-reflector mode reflect
set services rpm rfc2544-benchmarking tests test-name l2b-reflector family bridge
set services rpm rfc2544-benchmarking tests test-name l2b-reflector direction egress
set services rpm rfc2544-benchmarking tests test-name l2b-reflector test-interface ge-1/1/5.0
set interfaces ge-1/1/6 flexible-vlan-tagging
set interfaces ge-1/1/6 mtu 9192
set interfaces ge-1/1/6 encapsulation flexible-ethernet-services
set interfaces ge-1/1/6 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/6 unit 0 vlan-id 400
set interfaces ge-1/1/5 flexible-vlan-tagging
set interfaces ge-1/1/5 mtu 9192
set interfaces ge-1/1/5 encapsulation flexible-ethernet-services
set interfaces ge-1/1/5 unit 0 encapsulation vlan-bridge
set interfaces ge-1/1/5 unit 0 vlan-id 400
set bridge-domains bd1 domain-type bridge
set bridge-domains bd1 vlan-id 500
set bridge-domains bd1 interface ge-1/1/6.0
set bridge-domains bd1 interface ge-1/1/5.0

```

Configuring Throughput Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the throughput test and reference the test-profile in a unique test-name. The test-name defines the parameters for the throughput test to be performed on the ACX Series router.

To configure the throughput test parameters on the ACX Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```

[edit]
user@host# edit services rpm rfc2544-benchmarking

```

2. Define a name for the first test profile—for example, `tput` for the throughput test profile.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile tput
```

3. Configure the type of test to be performed as throughput, specify the packet size as 128 bytes, and define the theoretical maximum bandwidth for the test in kilobits per second (Kbps), with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile tput]
user@host# set test-type throughput packet-size 128 bandwidth-kbps 900000
```

4. Enter the `up` command twice to go to the `[edit services rpm rfc2544-benchmarking]` level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile tput ]
user@host# up
user@host# up
```

5. Define a name for the throughput test—for example, `tput-test`. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name tput-test
```

6. Specify the name of the test profile, `tput`, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set test-profile tput
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address  
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test to be E-LAN.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set ovlan-id 400 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, bridge, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set family bridge direction egress source-udp-port 200 destination-udp-port 200
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds, and specify the logical interface, ge-0/2/1.0, on which the RFC2544-benchmarking tests are run.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

Configuring Back-to-Back Frames Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the back to back frames test and reference the test-profile in a unique test-name. The test-name defines the parameters for the back to back frames test to be performed on the ACX Series router.

To configure the back-to-back frames test parameters on the ACX Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the back-to-back test profile—for example, b2bt.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile b2bt
```

3. Configure the type of test to be performed as back-to-back frames, specify the packet size as 128 bytes, and define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile b2bt]
user@host# set test-type back-to-back-frames packet-size 4444 bandwidth-kbps 950000
```

4. Enter the up command twice to go to the [edit services rpm rfc2544-benchmarking] level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile b2bt ]
user@host# up
user@host# up
```

5. Define a name for the back-to-back frames test—for example, b2bt-test. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name b2bt-test
```

6. Specify the name of the test profile, b2bt, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set test-profile b2bt
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set ovlan-id 400 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, bridge, for the benchmarking test and specify the direction, egress.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set family bridge direction egress
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/1.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

Configuring Latency Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the latency test and reference the test-profile in a unique test-name. The test-name defines the parameters for the latency test to be performed on the ACX Series router.

To configure the latency test parameters on the ACX Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the latency test profile—for example, lty.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile lty
```

3. Configure the type of test to be performed as latency, specify the packet size of the test packet, and define the maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# set test-profile lty test-type latency packet-size 512 bandwidth-kbps 1000000
```

4. Enter the up command twice to go to the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile lty]
user@host# up
user@host# up
```

5. Define a name for the latency test—for example, lty-test. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name lty-test
```

6. Specify the name of the test profile, lty, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set test-profile lty
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set ovlan-id 400 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, bridge, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set family bridge direction egress source-udp-port 200 destination-udp-port 200
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/1.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

Configuring Frame Loss Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the frame loss test and reference the test-profile in a unique test-name. The test-name defines the parameters for the frame loss test to be performed on the ACX Series router.

To configure the frame loss test parameters on the ACX Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the frame loss test profile—for example, frloss.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile frloss
```

3. Configure the type of test performed as frame loss, specify the packet size of the test packet, and define the maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# set test-profile frloss test-type frame-loss packet-size 1600 bandwidth-kbps 1000000
```

4. Enter the up command to go to the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles ]
user@host# up
```

5. Define a name for the frame loss test—for example, frloss-test. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name frloss-test
```

6. Specify the name of the test profile, frloss, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set test-profile frloss
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set ovlan-id 400 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, bridge, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set family bridge direction egress source-udp-port 200 destination-udp-port 200
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/1.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set test-iterator-duration 20 test-interface ge-1/1/3.0
```

12. Enter the exit command to go to the [edit] hierarchy level.

```
[edit services rpm rfc2544-benchmarking tests test-name test4 ]
user@host# exit
```

Configuring Other Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the interface and bridge domain on the ACX Series router:

1. Configure the Layer 2 NNI interface on which the tests must be run from the [edit] hierarchy level.

```
[edit]
user@host# edit interfaces ge-1/2/1
```

2. Configure flexible VLAN tagging for the transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-1/2/1]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation flexible-ethernet-services
```

3. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-1/2/1]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```

4. Configure the Layer 2 UNI interface.

```
[edit]
user@host# edit interfaces ge-1/1/3
```

5. Configure flexible VLAN tagging for transmission of non-tagged frames or 802.1Q single-tag and dual-tag frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-1/1/3]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation flexible-ethernet-services
```

6. Configure a logical unit for the interface and specify the encapsulation and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-1/1/3]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```

7. Configure the bridge domain, bd1, and specify the VLAN ID associated with the bridge domain and the associated interfaces from the [edit] hierarchy level.

```
[edit]
user@host# set bridge-domains bd1 vlan-id 600 interface ge-1/2/1.0
user@host# set bridge-domains bd1 vlan-id 600 interface ge-1/1/3.0
```

Configuring Benchmarking Test Parameters on the MX104 Router

Step-by-Step Procedure

The following configuration requires you to configure a unique test-name for the benchmarking test on the MX104 Series router. The test-name defines the parameters for the benchmarking test to be performed. Because the test interface and test MAC addresses are the same, you can create a single test configuration at the reflector (MX104).

To configure the benchmarking test parameters on the MX104 Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the test—for example, l2b-reflector. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name l2b-reflector
```

3. Specify the source and destination MAC addresses of the test packet.

```
[edit services rpm rfc2544-benchmarking test-name l2b-reflector]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address 00:00:5e:00:53:22
```

4. Specify the service type under test and the mode, which is reflect, at the reflector.

```
[edit services rpm rfc2544-benchmarking test-name l2b-reflector]
user@host# set service-type elan
```

5. Specify the mode which is reflect at the reflector.

```
[edit services rpm rfc2544-benchmarking test-name l2b-reflector]
user@host# set mode reflect
```

6. Configure the family type, bridge, and specify the direction, egress, for the benchmarking test. Also, specify the logical interface, ge-1/1/5.0, on which the RFC2544-based benchmarking test is being run.

```
[edit services rpm rfc2544-benchmarking tests test-name l2b-reflector]
user@host# set family bridge direction egress test-interface ge-1/1/5.0
```

Configuring Other Benchmarking Test Parameters on the MX104 Router

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the interface and bridge domain on the MX104 Series router:

1. Configure the Layer 2 NNI interface on which the tests must be run.

```
[edit]
user@host# edit interfaces ge-1/1/6
```

2. Configure flexible VLAN tagging for transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-1/1/6]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation flexible-ethernet-services
```

3. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interface.

```
[edit interfaces ge-1/1/6]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```

4. Configure the Layer 2 NNI interface.

```
[edit]
user@host# edit interfaces ge-1/1/5
```

5. Configure flexible VLAN tagging for transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-1/1/5]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation flexible-ethernet-services
```

6. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-1/1/5]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 400
```


7. Configure the bridge domain, bd1, and specify the VLAN ID associated with the bridge domain, and the associated interfaces from the [edit] hierarchy level.

```
[edit]
user@host# set bridge-domains bd1 vlan-id 500 interface ge-1/1/6.0
user@host# set bridge-domains bd1 vlan-id 500 interface ge-1/1/5.0
```

8. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test l2b-reflector start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the test services rpm rfc2544-benchmarking test l2b-reflector stop command.

Results

In configuration mode, confirm your configuration on the ACX Series router and the MX104 Series router by entering the show command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on the ACX Series router :

```
[edit interfaces]
ge-1/2/1 {
  flexible-vlan-tagging;
  mtu 9192;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 400;
  }
}
ge-1/1/3 {
  flexible-vlan-tagging;
  mtu 9192;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 400;
  }
}
```

```

[edit bridge-domains]
bd1 {
    vlan-id 600;
    interface ge-1/2/1.0;
    interface ge-1/1/3.0;
}

[edit services rpm]
rfc2544-benchmarking {
    profiles {
        test-profile tput {
            test-type throughput
            packet-size 128;
            bandwidth-kbps 900000;
        }
        test-profile b2bt {
            test-type back-back-frames
            packet-size 512;
            bandwidth-kbps 950000;
        }
        test-profile lty {
            test-type latency
            packet-size 512;
            bandwidth-kbps 100000;
        }
        test-profile frloss {
            test-type frameloss
            packet-size 1600;
            bandwidth-kbps 1000000;
        }
    }
    tests {
        test-name tput-test {
            interface ge-1/1/3.0;
            test-profile tput;
            mode initiate-and-terminate;
            source-mac-address 00:00:5e:00:53:11;
            destination-mac-address 00:00:5e:00:53:22;
            ovlan-id 400;
            service-type elan;
            family bridge;
            direction egress;
        }
    }
}

```

```

        source-udp-port 200;
        destination-udp-port 200;
        test-iterator-duration 20;
    }
    test-name b2b-test {
        interface ge-1/1/3.0;
        test-profile b2bt;
        mode initiate-and-terminate;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        ovlan-id 400;
        service-type elan;
        family bridge;
        direction egress;
        test-iterator-duration 20;
    }
    test-name lty-test {
        interface ge-1/1/3.0;
        test-profile lty;
        mode initiate-and-terminate;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        ovlan-id 400;
        service-type elan;
        family bridge;
        direction egress;
        source-udp-port 200;
        destination-udp-port 200;
        test-iterator-duration 20;
    }
    test-name frloss-test {
        interface ge-1/1/3.0;
        test-profile frloss;
        mode initiate-and-terminate;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        ovlan-id 400;
        service-type elan;
        family bridge;
        direction egress;
        source-udp-port 200;
        destination-udp-port 200;
        test-iterator-duration 20;
    }

```

```

    }
  }
}

```

Benchmarking Test Parameters on the MX104 Series router:

```

[edit interfaces]
  ge-1/1/6 {
    flexible-vlan-tagging;
    mtu 9192;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id 400;
    }
  }

  ge-1/1/5 {
    flexible-vlan-tagging;
    mtu 9192;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id 400;
    }
  }
}

[edit bridge-domains]
  bd1 {
    vlan-id 500;
    interface ge-1/1/6.0;
    interface ge-1/1/5.0;
  }

[edit services rpm]
  rfc2544-benchmarking {
    # Note, When in reflector mode, test profile is not needed
    tests {
      test-name l2b-reflector {
        interface ge-1/1/5.0;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        family bridge;
      }
    }
  }

```

```

        mode reflect;
        service-type elan;
        family bridge;
        direction egress;
    }
}
}

```

Verifying the Results of the Benchmarking Tests for Layer 2 Services (E-LAN) in Bridge Domains

IN THIS SECTION

- [Verifying the Throughput Benchmarking Test Results | 810](#)
- [Verifying the Back-to-Back Benchmarking Test Results | 813](#)
- [Verifying the Frame Loss Benchmarking Test Results | 816](#)
- [Verifying the Latency Benchmarking Test Results | 819](#)

Examine the results of the benchmarking tests that are performed on the configured service between the ACX Series router and the MX104 Series router. Start the test on the reflector first and then start the test on the initiator.

Verifying the Throughput Benchmarking Test Results

Purpose

Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX Series router and the MX104 Series router.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the ACX Series router.

```

user@host> show services rpm rfc2544-benchmarking test-id 1 detail
Test information :
    Test id: 1, Test name: tput_test, Test type: Throughput

```

Test mode: Initiate-and-Terminate
 Test packet size: 128
 Test state: TEST_STATE_COMPLETED
 Status: Test-Completed
 Test start time: 2014-09-24 22:21:09 PDT
 Test finish time: 2014-09-24 22:21:33 PDT
 Counters last cleared: Never

Test-profile Configuration:

Test-profile name: tput
 Test packet size: 128
 Theoretical max bandwidth : 900000 kbps

Test Configuration:

Test mode: Initiate-and-Terminate
 Duration in seconds: 20
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %

Bridge family Configuration:

Interface : ge-1/1/3.0
 Test direction: Egress
 Source mac address: 00:00:5e:00:53:11
 Destination mac address: 00:00:5e:00:53:22
 Outer vlan-id: 400
 Outer vlan priority: 0
 Outer vlan cfi: 0
 Outer tag protocol id: 0x8100
 Source ipv4 address: 192.168.1.10
 Destination ipv4 address: 192.168.1.20
 Source udp port: 200
 Destination udp port: 200

Rfc2544 throughput test information :

Initial test load percentage : 100.00 %
 Test iteration mode : Binary
 Test iteration step : 50.00 %
 Theoretical max bandwidth : 900000 kbps

Test packet size: 128

Iteration	Internal Overhead	Duration (sec)	Elapsed time	----- Theoretical	Throughput Transmit Measured
1	0	20	20	100.00 %	100.00 % 100.00 %

Result of the iteration runs : Throughput Test complete for packet size 128

Best iteration : 1, Best iteration (pps) : 760135

Best iteration throughput : 100.00 %

RFC2544 Throughput test results summary:

Packet Size (kbps)	Internal overhead	Theoretical rate (pps)	Transmit pps	Tx Packets	Rx Packets	Measured throughput %	Measured bandwidth
128	0	760135	760135	15202700	15202700	100.00 %	900000

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the MX104 Series router.

```
user@host> show services rpm rfc2544-benchmarking test-id 1 detail
```

Test information :

Test id: 1, Test name: l2b-reflector, Test type: Reflect

Test mode: Reflect

Test packet size: 0

Test state: TEST_STATE_RUNNING

Status: Running

Test start time: 2014-09-24 22:20:54 PDT

Test finish time: TEST_RUNNING

Counters last cleared: Never

Test Configuration:

Test mode: Reflect

Duration in seconds: 864000

Test finish wait duration in seconds: 1

Test family: Bridge

Test iterator pass threshold: 0.50 %

Test receive failure threshold: 0.00 %

Test transmit failure threshold: 0.50 %

Bridge family Configuration:

Interface : ge-1/1/5.0

Test direction: Egress

```
Source mac address: 00:00:5e:00:53:11
Destination mac address: 00:00:5e:00:53:22
Service type: Elan
```

Elapsed time	Reflected Packets	Reflected Bytes
61	15202700	1945945600

You can also use the `show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the run `show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` in the CLI Explorer.

Verifying the Back-to-Back Benchmarking Test Results

Purpose

Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX Series router and the MX104 Series router.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the ACX Series router.

```
user@host> show services rpm rfc2544-benchmarking test-id 4 detail
Test information :
  Test id: 4, Test name: b2b-test, Test type: Back-Back-Frames
  Test mode: Initiate-and-Terminate
  Test packet size: 128 512
  Test state: TEST_STATE_COMPLETED
  Status: Test-Completed
  Test start time: 2014-09-24 22:30:16 PDT
  Test finish time: 2014-09-24 22:31:03 PDT
  Counters last cleared: Never
```


Test-profile Configuration:

Test-profile name: b2bt
 Test packet size: 128 512
 Theoretical max bandwidth : 950000 kbps

Test Configuration:

Test mode: Initiate-and-Terminate
 Duration in seconds: 20
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %

Bridge family Configuration:

Interface : ge-1/1/3.0
 Test direction: Egress
 Source mac address: 00:00:5e:00:53:11
 Destination mac address: 00:00:5e:00:53:22
 Outer vlan-id: 400
 Outer vlan priority: 0
 Outer vlan cfi: 0
 Outer tag protocol id: 0x8100
 Source ipv4 address: 192.168.1.10
 Destination ipv4 address: 192.168.1.20
 Source udp port: 4040
 Destination udp port: 4041

Rfc2544 Back-Back test information :

Initial burst length: 20 seconds at 950000 kbps
 Test iteration mode : Binary
 Test iteration step : 50.00 %

Test packet size: 128

Iteration	Theoretical burst length (packets)	Transmit burst length (packets)	Internal overhead	Duration time	Elapsed
1	16047280	16047280	0	20	20

Result of the iteration runs : Back-Back Test complete for packet size 128

Best iteration : 1

Measured burst (num sec) : 20 sec

Measured burst (num pkts) : 16047280 packets

```

Test packet size: 512
Iteration   Theoretical    Transmit    Internal Duration Elapsed
            burst length burst length overhead  time
            (packets)   (packets)
      1    4464280      4464280         0      20      20

```

Result of the iteration runs : Back-Back Test complete for packet size 512

Best iteration : 1

Measured burst (num sec) : 20 sec

Measured burst (num pkts) : 4464280 packets

RFC2544 Back-Back test results summary:

Packet Size	Measured Burst length (Packets)	Time (seconds)
128	16047280	20
512	4464280	20

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the MX104 Series router.

```

user@host> show services rpm rfc2544-benchmarking test-id 4 detail

```

Test information :

Test id: 4, Test name: l2b-reflector, Test type: Reflect

Test mode: Reflect

Test packet size: 0

Test state: TEST_STATE_RUNNING

Status: Running

Test start time: 2014-09-24 22:30:07 PDT

Test finish time: TEST_RUNNING

Counters last cleared: Never

Test Configuration:

Test mode: Reflect

Duration in seconds: 864000

Test finish wait duration in seconds: 1

Test family: Bridge

Test iterator pass threshold: 0.50 %

Test receive failure threshold: 0.00 %

Test transmit failure threshold: 0.50 %

```

Bridge family Configuration:
  Interface : ge-1/1/5.0
  Test direction: Egress
  Source mac address: 00:00:5e:00:53:11
  Destination mac address: 00:00:5e:00:53:22
  Service type: Elan

```

Elapsed time	Reflected Packets	Reflected Bytes
58	20511560	4339763200

You can also use the `show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the run `show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` in the CLI Explorer.

Verifying the Frame Loss Benchmarking Test Results

Purpose

Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX Series router and the MX104 Series router.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the ACX Series router.

```

user@host> show services rpm rfc2544-benchmarking test-id 3 detail
Test information :
  Test id: 3, Test name: frloss-test, Test type: Frame-Loss
  Test mode: Initiate-and-Terminate
  Test packet size: 1600
  Test state: TEST_STATE_COMPLETED
  Status: Test-Completed

```

Test start time: 2014-09-24 22:26:45 PDT
 Test finish time: 2014-09-24 22:27:55 PDT
 Counters last cleared: Never

Test-profile Configuration:

Test-profile name: frloss
 Test packet size: 1600
 Theoretical max bandwidth : 1000000 kbps

Test Configuration:

Test mode: Initiate-and-Terminate
 Duration in seconds: 20
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %

Bridge family Configuration:

Interface : ge-1/1/3.0
 Test direction: Egress
 Source mac address: 00:00:5e:00:53:11
 Destination mac address: 00:00:5e:00:53:22
 Outer vlan-id: 400
 Outer vlan priority: 0
 Outer vlan cfi: 0
 Outer tag protocol id: 0x8100
 Source ipv4 address: 192.168.1.10
 Destination ipv4 address: 192.168.1.20
 Source udp port: 200
 Destination udp port: 200

Rfc2544 frame-loss test information :

Initial test load percentage : 100.00 %
 Test iteration mode : step-down
 Test iteration step : 10 %
 Theoretical max bandwidth : 1000000 kbps

Test packet size: 1600

Iteration	Internal Overhead	Duration (sec)	Elapsed time	----- Theoretical	Throughput Transmit Measured	----- Frame-loss rate %
1	0	20	20	100.00 %	100.00 % 100.00 %	0.00 %
2	0	20	20	100.00 %	100.00 % 100.00 %	0.00 %

```

3      0      20      20      100.00 %      100.00 % 100.00 % 0.00 %

```

Result of the iteration runs : Frame-loss test complete for packet size 1600

Percentage throughput transmitted: 100.00 %

Frame-loss rate (percent) : 0.00 %

RFC2544 Frame-loss test results summary:

Packet Loss Size percent	Internal overhead	Theoretical rate (pps)	Transmit pps	Transmit throughput	Tx Packets	Rx Packets	Frame rate
1600	0	77160	77160	100.00 %	1543200	1543200	0.00 %

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the MX104 Series router.

```

user@host> show services rpm rfc2544-benchmarking test-id 3 detail

```

Test information :

Test id: 3, Test name: l2b-reflector, Test type: Reflect

Test mode: Reflect

Test packet size: 0

Test state: TEST_STATE_RUNNING

Status: Running

Test start time: 2014-09-24 22:25:36 PDT

Test finish time: TEST_RUNNING

Counters last cleared: Never

Test Configuration:

Test mode: Reflect

Duration in seconds: 864000

Test finish wait duration in seconds: 1

Test family: Bridge

Test iterator pass threshold: 0.50 %

Test receive failure threshold: 0.00 %

Test transmit failure threshold: 0.50 %

Bridge family Configuration:

Interface : ge-1/1/5.0

Test direction: Egress

Source mac address: 00:00:5e:00:53:11

Destination mac address: 00:00:5e:00:53:22

Service type: Elan

Elapsed time	Reflected Packets	Reflected Bytes
95	1624361	2598977600

You can also use the `show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the run `show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` in the CLI Explorer.

Verifying the Latency Benchmarking Test Results

Purpose

Verify that the necessary and statistical values are displayed for the benchmarking tests that are run on the configured service between the ACX Series router and the MX104 Series router.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the ACX Series router.

```
user@host> show services rpm rfc2544-benchmarking test-id 5 detail
```

Test information :

Test id: 5, Test name: lty-test, Test type: Latency

Test mode: Initiate-and-Terminate

Test packet size: 512

Test state: TEST_STATE_COMPLETED

Status: Test-Completed

Test start time: 2014-09-24 22:33:05 PDT

Test finish time: 2014-09-24 22:40:46 PDT

Counters last cleared: Never

Test-profile Configuration:

Test-profile name: lty
 Test packet size: 512
 Theoretical max bandwidth : 1000000 kbps

Test Configuration:

Test mode: Initiate-and-Terminate
 Duration in seconds: 20
 Test finish wait duration in seconds: 1
 Test family: Bridge
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %

Bridge family Configuration:

Interface : ge-1/1/3.0
 Test direction: Egress
 Source mac address: 00:00:5e:00:53:11
 Destination mac address: 00:00:5e:00:53:22
 Outer vlan-id: 400
 Outer vlan priority: 0
 Outer vlan cfi: 0
 Outer tag protocol id: 0x8100
 Source ipv4 address: 192.168.1.10
 Destination ipv4 address: 192.168.1.20
 Source udp port: 200
 Destination udp port: 200

Rfc2544 latency test information :

Theoretical max bandwidth : 1000000 kbps
 Initial test load percentage : 100.00 %
 Duration in seconds: 20
 Measurement unit for timestamp: Nanoseconds

Test packet size: 512

Iteration	Duration	Elapsed	Theoretical	Transmit	Throughput	-----
Latency	-----					
	(sec)	time	rate (pps)	pps	percent	Minimum
Average	Maximum	Probe				
1	20	20	234962	234962	100.00 %	44008
45253	47424	45096				
2	20	20	234962	234962	100.00 %	44008
45237	47456	45256				
3	20	20	234962	234962	100.00 %	43864

45198	46976	45144				
4	20	20	234962	234962	100.00 %	43832
45243	47088	45096				
5	20	20	234962	234962	100.00 %	44072
45261	46976	45176				
6	20	20	234962	234962	100.00 %	43784
45214	46864	45032				
7	20	20	234962	234962	100.00 %	44024
45259	47216	45240				
8	20	20	234962	234962	100.00 %	44072
45290	46864	45192				
9	20	20	234962	234962	100.00 %	43976
45272	46792	45208				
10	20	20	234962	234962	100.00 %	44024
45206	46976	45112				
11	20	20	234962	234962	100.00 %	44040
45198	47088	45176				
12	20	20	234962	234962	100.00 %	44008
45223	46976	45160				
13	20	20	234962	234962	100.00 %	44088
45257	47408	45176				
14	20	20	234962	234962	100.00 %	43976
45183	46832	45080				
15	20	20	234962	234962	100.00 %	44024
45198	47088	45112				
16	20	20	234962	234962	100.00 %	43864
45206	46912	45208				
17	20	20	234962	234962	100.00 %	44056
45209	46960	45176				
18	20	20	234962	234962	100.00 %	44008
45198	46912	45112				
19	20	20	234962	234962	100.00 %	43816
45175	47248	45000				
20	20	20	234962	234962	100.00 %	43912
45202	46992	45192				

Result of the iteration runs : Latency Test complete for packet size 512

Internal overhead per packet: 0

Avg (min) Latency : 43972

Avg (avg) latency : 45224

Avg (Max) latency : 47052

Avg (probe) latency : 45147

RFC2544 Latency test results summary:

Packet	Internal	Theoretical	Transmit	Tx	Rx	----- Latency	
Size	overhead	rate (pps)	pps	Packets	Packets	Minimum	Average
Maximum	Probe						
512	0	234962	234962	93984800	93984800	43972	45224
47052	45147						

In operational mode, enter the `show services rpm rfc2544-benchmarking test-id test-id-number detail` command on the MX104 Series router.

```
user@host> show services rpm rfc2544-benchmarking test-id 5 detail
```

Test information :

Test id: 5, Test name: l2b-reflector, Test type: Reflect

Test mode: Reflect

Test packet size: 0

Test state: TEST_STATE_RUNNING

Status: Running

Test start time: 2014-09-24 22:32:55 PDT

Test finish time: TEST_RUNNING

Counters last cleared: Never

Test Configuration:

Test mode: Reflect

Duration in seconds: 864000

Test finish wait duration in seconds: 1

Test family: Bridge

Test iterator pass threshold: 0.50 %

Test receive failure threshold: 0.00 %

Test transmit failure threshold: 0.50 %

Bridge family Configuration:

Interface : ge-1/1/5.0

Test direction: Egress

Source mac address: 00:00:5e:00:53:11

Destination mac address: 00:00:5e:00:53:22

Service type: Elan

Elapsed

Reflected

Reflected

time	Packets	Bytes
426	84586320	43308195840

You can also use the `show services rpm rfc2544-benchmarking (aborted-test | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC2544-based benchmarking tests for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the `show services rpm rfc2544-benchmarking operational mode` command, see `show services rpm rfc2544-benchmarking` topic in the CLI Explorer.

RELATED DOCUMENTATION

[Understanding RFC2544-Based Benchmarking Tests for E-LAN and E-Line Services on MX Series Routers | 733](#)

[Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | 737](#)

Example: Configuring Benchmarking Tests to Measure SLA Parameters for E-LAN Services on an MX104 Router Using VPLS

IN THIS SECTION

- [Requirements | 824](#)
- [Overview | 824](#)
- [Configuration | 825](#)
- [Verifying the Results of the Benchmarking Test for Layer 2 ELAN Services Using VPLS | 853](#)

This example shows how to configure benchmarking tests for the E-LAN services using BGP-based VPLS. The example covers the four benchmarking tests: throughput, frame loss, back-to-back frames, and latency.

Requirements

NOTE: MX Series routers support only the reflector function in RFC2544-based benchmarking tests.

This example uses the following hardware and software components:

- An MX104 3D Universal Edge Router (reflector)
- Any MX Series router
- Any ACX Series router (initiator)
- Junos OS Release 15.1 or later for MX Series routers


Overview

Consider a sample topology in which an ACX Series router functions as an initiator and terminator of the test frames for an RFC2544-based benchmarking test. The ACX Series router is connected to a provider edge router, PE1 (an MX Series router). The PE1 router is configured with a VPLS routing instance and is connected over a Layer 2 network to another provider edge router, PE2 (an MX104 Series router). A simple VPLS network with BGP signaling is created between routers PE1 and PE2. The MX104 Series router also functions as a reflector to reflect the test frames it receives from the ACX Series router back to the initiator.

Benchmarking tests compute the performance attributes in the user-to-network interface (UNI) direction of the Layer 2 E-LAN service between the ACX Series router and the MX104 Series router. To measure SLA parameters for E-LAN services using VPLS, configure specific benchmarking tests. In this example, all four benchmarking tests (throughput, back-to-back frames, latency, and frame-loss) are configured.

Figure 67 on page 825 shows the sample topology to perform all four RFC2544-based benchmarking tests for the UNI direction on a Layer 2 network using VPLS.

Figure 67: Layer 2 Reflection with Simple BGP-based VPLS Topology



On the ACX Series router, ge-0/2/1.0 is the Layer 2 NNI interface and ge-0/2/0.0 is the Layer 2 UNI interface. For each benchmarking test configured on the ACX Series router, specify the source MAC address as 00:00:5e:00:53:11 and 00:00:5e:00:53:22 as the destination MAC address. Also, specify the VLAN ID as 512. On the MX Series router, ge-0/3/0.0 is the Layer 2 NNI interface and ge-0/2/1.0 is the UNI interface. On the MX104 Series router, ge-0/2/5.0 is the Layer 2 NNI interface and ge-0/3/1.0 is the Layer 2 UNI interface. The benchmarking tests are used to compute the performance attributes for an E-LAN service using VPLS.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 826](#)
- [Configuring Throughput Benchmarking Test Parameters on the ACX Series Router \(Initiator\) | 831](#)
- [Configuring Back-to-Back Frames Benchmarking Test Parameters on the ACX Series Router | 833](#)
- [Configuring Latency Benchmarking Test Parameters on the ACX Series Router | 835](#)
- [Configuring Frame Loss Benchmarking Test Parameters on the ACX Series Router | 837](#)
- [Configuring Other Benchmarking Test Parameters on the ACX Series Router | 840](#)
- [Configuring the VPLS Parameters on the MX Series Router \(PE1\) | 841](#)
- [Configuring Benchmarking Test Parameters on the MX104 Router \(Reflector\) | 843](#)
- [Configuring Other Benchmarking Test Parameters on the MX104 Router \(Reflector\) | 845](#)
- [Configuring VPLS Parameters on the MX104 Router \(Reflector\) | 846](#)
- [Results | 849](#)

In this example, you configure the benchmarking tests for the UNI direction for a Layer 2 E-LAN service using VPLS between two routers (initiator and reflector) to detect and analyze the performance of the interconnected routers. The initiator and reflector routers are not directly connected to each other. The initiator is connected to a provider edge router (PE1), which is in turn connected to the reflector. In this example, the ACX Series router is the initiator, an MX Series router is PE1, and the MX104 router is the other provider edge router (PE2) and reflector. Start by configuring the initiator. On the ACX Series router, you first configure each test by specifying the test profile and the test attributes, and then define the test by associating the test with the test profile with the relevant attributes. You can then configure the interface. On the MX Series router, configure the VPLS parameters to enable VPLS on the router. On the MX104 Series router, configure the benchmarking parameters and the VPLS parameters.

NOTE: When you configure Layer 2 reflection, you can specify the service type under test as ELINE if you want to simulate an Eline service by using bridge encapsulation.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configuring Benchmarking Test Parameters on the ACX Series Router (Initiator)

```
set services rpm rfc2544-benchmarking profiles test-profile tput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile tput packet-size 256
set services rpm rfc2544-benchmarking profiles test-profile tput bandwidth-kbps 600000
set services rpm rfc2544-benchmarking profiles test-profile b2bt test-type back-back-frames
set services rpm rfc2544-benchmarking profiles test-profile b2bt packet-size 9104
set services rpm rfc2544-benchmarking profiles test-profile b2bt bandwidth-kbps 600000
set services rpm rfc2544-benchmarking profiles test-profile lty test-type latency
set services rpm rfc2544-benchmarking profiles test-profile lty packet-size 1024
set services rpm rfc2544-benchmarking profiles test-profile lty bandwidth-kbps 6000000
set services rpm rfc2544-benchmarking profiles test-profile frloss test-type frame-loss
set services rpm rfc2544-benchmarking profiles test-profile frloss packet-size 1600
set services rpm rfc2544-benchmarking profiles test-profile frloss bandwidth-kbps 6000000
set services rpm rfc2544-benchmarking profiles test-profile frloss step-percent 5
set services rpm rfc2544-benchmarking tests test-name tput-test test-profile tput
set services rpm rfc2544-benchmarking tests test-name tput-test source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name tput-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name tput-test ovlan-id 512
```

```

set services rpm rfc2544-benchmarking tests test-name tput-test service-type elan
set services rpm rfc2544-benchmarking tests test-name tput-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name tput-test family bridge
set services rpm rfc2544-benchmarking tests test-name tput-test direction egress
set services rpm rfc2544-benchmarking tests test-name tput-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name tput-test destination-udp-port 400
set services rpm rfc2544-benchmarking tests test-name tput-test test-iterator-duration 250
set services rpm rfc2544-benchmarking tests test-name tput-test test-interface ge-0/2/0.0
set services rpm rfc2544-benchmarking tests test-name b2bt-test test-profile b2bt
set services rpm rfc2544-benchmarking tests test-name b2bt-test source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name b2bt-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name b2bt-test ovlan-id 512
set services rpm rfc2544-benchmarking tests test-name b2bt-test service-type elan
set services rpm rfc2544-benchmarking tests test-name b2bt-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name b2bt-test family bridge
set services rpm rfc2544-benchmarking tests test-name b2bt-test direction egress
set services rpm rfc2544-benchmarking tests test-name b2bt--test destination-udp-port 400
set services rpm rfc2544-benchmarking tests test-name b2bt-test test-iterator-duration 10
set services rpm rfc2544-benchmarking tests test-name b2b-test test-interface ge-0/2/0.0
set services rpm rfc2544-benchmarking tests test-name lty-test test-profile lty
set services rpm rfc2544-benchmarking tests test-name lty-test source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name lty-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name lty-test ovlan-id 512
set services rpm rfc2544-benchmarking tests test-name lty-test service-type elan
set services rpm rfc2544-benchmarking tests test-name lty-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name lty-test family bridge
set services rpm rfc2544-benchmarking tests test-name lty-test direction egress
set services rpm rfc2544-benchmarking tests test-name lty-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name lty-test destination-udp-port 400
set services rpm rfc2544-benchmarking tests test-name lty-test test-iterator-duration 10
set services rpm rfc2544-benchmarking tests test-name lty-test test-interface ge-0/2/0.0
set services rpm rfc2544-benchmarking tests test-name frloss-test test-profile frloss
set services rpm rfc2544-benchmarking tests test-name frloss-test source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name frloss-test destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name frloss-test ovlan-id 512
set services rpm rfc2544-benchanrking tests test-name frloss-test ovlan-priority 7
set services rpm rfc2544-benchanrking tests test-name frloss-test ovlan-cfi 1

```

```

set services rpm rfc2544-benchmarking tests test-name frloss-test service-type elan
set services rpm rfc2544-benchmarking tests test-name frloss-test mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name frloss-test family bridge
set services rpm rfc2544-benchmarking tests test-name frloss-test direction egress
set services rpm rfc2544-benchmarking tests test-name frloss-test source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name frloss-test destination-udp-port 400
set services rpm rfc2544-benchmarking tests test-name frloss-test test-iterator-duration 30
set services rpm rfc2544-benchmarking tests test-name frloss-test test-interface ge-0/2/0.0
set interfaces ge-0/2/0 flexible-vlan-tagging
set interfaces ge-0/2/0 mtu 9192
set interfaces ge-0/2/0 encapsulation flexible-ethernet-services
set interfaces ge-0/2/0 unit 0 encapsulation vlan-bridge
set interfaces ge-0/2/0 unit 0 vlan-id 512
set interfaces ge-0/2/1 flexible-vlan-tagging
set interfaces ge-0/2/1 mtu 9192
set interfaces ge-0/2/1 encapsulation flexible-ethernet-services
set interfaces ge-0/2/1 unit 0 encapsulation vlan-bridge
set interfaces ge-0/2/1 unit 0 vlan-id 512
set bridge-domains bd1 vlan-id 10
set bridge-domains bd1 interface ge-0/2/1.0
set bridge-domains bd1 interface ge-0/2/0.0

```

Configuring VPLS Parameters on the MX Router (Provider Edge Router PE1)

```

set chassis fpc 0 pic 2 tunnel-services
set interfaces ge-0/2/1 flexible-vlan-tagging
set interfaces ge-0/2/1 mtu 9192
set interfaces ge-0/2/1 encapsulation vlan-vpls
set interfaces ge-0/2/1 unit 0 encapsulation vlan-vpls
set interfaces ge-0/2/1 unit 0 vlan-id 512
set interfaces ge-0/3/0 mtu 9192
set interfaces ge-0/3/0 unit 0 family inet address 192.0.2.1/24
set interfaces ge-0/3/0 unit 0 family mpls
set interfaces lo0 unit 0 family inet address 198.51.100.1/32
set routing-options router-id 198.51.100.1
set routing-options autonomous-system 65100
set protocols mpls interface ge-0/3/0.0
set protocols bgp group test type internal
set protocols bgp group test local-address 198.51.100.1
set protocols bgp group test family l2vpn signaling
set protocols bgp group test neighbor 198.51.100.2
set protocols ospf traffic-engineering

```

```

set protocols ospf reference-bandwidth 1g
set protocols ospf area 0.0.0.0 interface ge-0/3/0.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp interface ge-0/3/0.0 set protocols ldp interface lo0.0
set routing-instances vpls-pe1 instance-type vpls
set routing-instances vpls-pe1 interface ge-0/2/1.0
set routing-instances vpls-pe1 no-local-switching
set routing-instances vpls-pe1 route-distinguisher 198.51.100.1:101
set routing-instances vpls-pe1 vrf-target target:1:2
set routing-instances vpls-pe1 protocols vpls site-range 8
set routing-instances vpls-pe1 protocols vpls no-tunnel-services
set routing-instances vpls-pe1 protocols vpls site HUB site-identifier 1
set routing-instances vpls-pe1 protocols vpls vpls-id 1
set routing-instances vpls-pe1 protocols vpls neighbor 198.51.100.2

```

Configuring Benchmarking Test Parameters and VPLS Parameters on the MX104 Router (Provider Edge Router PE2)

```

set services rpm rfc2544-benchmarking tests test-name l2v-reflector source-mac-address
00:00:5e:00:53:11
set services rpm rfc2544-benchmarking tests test-name l2v-reflector destination-mac-address
00:00:5e:00:53:22
set services rpm rfc2544-benchmarking tests test-name l2v-reflector service-type elan
set services rpm rfc2544-benchmarking tests test-name l2v-reflector in-service
set services rpm rfc2544-benchmarking tests test-name l2v-reflector ip-swap
set services rpm rfc2544-benchmarking tests test-name l2v-reflector udp-tcp-port-swap
set services rpm rfc2544-benchmarking tests test-name l2v-reflector mode reflect
set services rpm rfc2544-benchmarking tests test-name l2v-reflector family vpls
set services rpm rfc2544-benchmarking tests test-name l2v-reflector reflect-etype 2048
set services rpm rfc2544-benchmarking tests test-name l2v-reflector direction egress
set services rpm rfc2544-benchmarking tests test-name l2v-reflector source-udp-port 200
set services rpm rfc2544-benchmarking tests test-name l2v-reflector destination-udp-port 200
set services rpm rfc2544-benchmarking tests test-name l2v-reflector test-interface ge-0/3/1.0
set interfaces ge-0/2/5 mtu 9192
set interfaces ge-0/2/5 unit 0 family inet address 203.0.113.1/24
set interfaces ge-0/2/5 unit 0 family mpls
set interfaces ge-0/3/1 flexible-vlan-tagging
set interfaces ge-0/3/1 mtu 9192
set interfaces ge-0/3/1 encapsulation vlan-vpls
set interfaces ge-0/3/1 unit 0 encapsulation vlan-vpls
set interfaces ge-0/3/1 unit 0 vlan-id 512
set interfaces ge-0/3/1 unit 0 family vpls filter input portmirror

```



```

set interfaces ge-0/3/1 unit 0 family vpls filter output portmirror
set interfaces ge-0/3/2 flexible-vlan-tagging
set interfaces ge-0/3/2 mtu 9192
set interfaces ge-0/3/2 encapsulation vlan-vpls
set interfaces ge-0/3/2 unit 0 encapsulation vlan-vpls
set interfaces ge-0/3/2 unit 0 vlan-id 512
set interfaces lo0 unit 0 family inet address 198.51.100.2/32
set forwarding-options port-mirroring input rate 1
set forwarding-options port-mirroring family vpls output interface ge-0/3/3.0
set forwarding-options port-mirroring family vpls output no-filter-check
set forwarding-options port-mirroring instance pm1 input rate 10000
set forwarding-options port-mirroring instance pm1 family vpls output interface ge-0/3/3.0
set routing-options router-id 198.51.100.2
set routing-options autonomous-system 65100
set protocols mpls interface ge-0/2/5.0
set protocols bgp group test type internal
set protocols bgp group test local-address 198.51.100.2
set protocols bgp group test family l2vpn signaling
set protocols bgp group test neighbor 198.51.100.1
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface ge-0/2/5.0
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ldp interface ge-0/2/5.0
set protocols ldp interface lo0.0
set firewall family vpls filter portmirror term 1 then count pm1
set firewall family vpls filter portmirror term 1 then accept
set firewall family vpls filter portmirror term 1 then port-mirror
set routing-instances vpls-pe2 instance-type vpls
set routing-instances vpls-pe2 interface ge-0/3/1.0
set routing-instances vpls-pe2 interface ge-0/3/3.0
set routing-instances vpls-pe2 no-local-switching
set routing-instances vpls-pe2 route-distinguisher 198.51.100.2:102
set routing-instances vpls-pe2 vrf-target target:1:2
set routing-instances vpls-pe2 protocols vpls site-range 8
set routing-instances vpls-pe2 protocols vpls no-tunnel-services
set routing-instances vpls-pe2 protocols vpls site SPOKE site-identifier 2
set routing-instances vpls-pe2 protocols vpls vpls-id 1
set routing-instances vpls-pe2 protocols vpls neighbor 198.51.100.1

```

Configuring Throughput Benchmarking Test Parameters on the ACX Series Router (Initiator)

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the throughput test and reference the test profile in a unique test name. The test name defines the parameters for the throughput test to be performed on the ACX Series router.

To configure the throughput test parameters on the ACX Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the first test profile—for example, tput—for the throughput test profile.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile tput
```

3. Configure the type of test to be performed as throughput, specify the packet size as 256 bytes, and define the theoretical maximum bandwidth for the test as 600000 Kbps. You can specify any value from 1 Kbps through 1,000,000 Kbps for the maximum bandwidth.

```
[edit services rpm rfc2544-benchmarking profiles test-profile tput]
user@host# set test-type throughput packet-size 256 bandwidth-kbps 600000
```

4. Enter the up command twice to go to the [edit services rpm rfc2544-benchmarking] level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile tput ]
user@host# up
user@host# up
```

5. Define a name for the throughput test—for example, `tput-test`. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name tput-test
```

6. Specify the name of the test profile, `tput`, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set test-profile tput
```

7. Configure the source and destination MAC addresses for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test to be E-LAN.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set ovlan-id 512 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, `bridge`, for the benchmarking test and specify the direction, `egress`. Also, specify the source and destination UDP ports to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set family bridge direction egress source-udp-port 200 destination-udp-port 400
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds, and specify the logical interface, ge-0/2/0.0, on which the RFC2544-benchmarking tests are run.

```
[edit services rpm rfc2544-benchmarking tests test-name tput-test]
user@host# set test-iterator-duration 250 test-interface ge-0/2/0.0
```

Configuring Back-to-Back Frames Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the back-to-back frames test and reference the test profile in a unique test name. The test name defines the parameters for the back-to-back frames test to be performed on the ACX Series router.

To configure the back-to-back frames test parameters on the ACX Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the back-to-back test profile—for example, b2bt.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile b2bt
```

3. Configure the type of test to be performed as back-to-back frames, specify the packet size as 9104 bytes, and specify the theoretical maximum bandwidth for the test as 600000 Kbps. You can specify any value from 1 Kbps through 1,000,000 Kbps as the maximum bandwidth.

```
[edit services rpm rfc2544-benchmarking profiles test-profile b2bt]
user@host# set test-type back-to-back-frames packet-size 9104 bandwidth-kbps 600000
```

4. Enter the `up` command twice to go to the `[edit services rpm rfc2544-benchmarking]` level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile b2bt ]
user@host# up
user@host# up
```

5. Define a name for the back-to-back frames test—for example, `b2bt-test`. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name b2bt-test
```

6. Specify the name of the test profile, `b2bt`, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set test-profile b2bt
```

7. Configure the source and destination MAC addresses for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test as E-LAN.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set ovlan-id 512 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, bridge, for the benchmarking test and specify the direction, egress.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set family bridge direction egress
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/0.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name b2bt-test]
user@host# set test-iterator-duration 10 test-interface ge-0/2/0.0
```

Configuring Latency Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the latency test and reference the test-profile in a unique test-name. The test-name defines the parameters for the latency test to be performed on the initiator (ACX Series router).

To configure the latency test parameters on the initiator:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the latency test profile—for example, lty.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile lty
```

3. Configure the type of test to be performed as latency, specify the packet size of the test packet as 1024, and specify the maximum bandwidth for the test in Kbps, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# set test-profile lty test-type latency packet-size 1024 bandwidth-kbps 600000
```

4. Enter the up command twice to go to the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile lty]
user@host# up
user@host# up
```

5. Define a name for the latency test—for example, lty-test. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name lty-test
```

6. Specify the name of the test profile, lty, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set test-profile lty
```

7. Configure the source and destination MAC addresses for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID for the test frames and specify the service type under test.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set ovlan-id 512 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, bridge, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set family bridge direction egress source-udp-port 200 destination-udp-port 400
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/0.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name lty-test]
user@host# set test-iterator-duration 10 test-interface ge-0/2/0.0
```

Configuring Frame Loss Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following configuration requires you to configure a test profile for the frame loss test and reference the test-profile in a unique test-name. The test-name defines the parameters for the frame loss test to be performed on the ACX Series router.

To configure the frame loss test parameters on the ACX Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```


2. Define a name for the frame loss test profile—for example, `frloss`.

```
[edit services rpm rfc2544-benchmarking]
user@host# set profiles test-profile frloss
```

3. Configure the type of test performed as frame loss, specify the packet size of the test packet, and define the maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# set test-profile frloss test-type frame-loss packet-size 1600 bandwidth-kbps
600000
```

4. Enter the `up` command to go to the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```

5. Define a name for the frame loss test—for example, `frloss-test`. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name frloss-test
```

6. Specify the name of the test profile, `frloss`, to be associated with the test name.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set test-profile frloss
```

7. Configure the source and destination MAC address for the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address
00:00:5e:00:53:22
```

8. Configure the outer VLAN ID, priority, and the canonical format indicator (cfi) value for the test frames. Together, the four added bytes, priority (3 bits) and canonical format indicator (1 bit) form the VLAN tag. Also, specify the service type under test.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set ovlan-id 512 ovlan-priority 7 ovlan-cfi 1 service-type elan
```

9. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set mode initiate-and-terminate
```

10. Configure the family type, bridge, for the benchmarking test and specify the direction, egress. Also, specify the source and destination UDP port to be used in the UDP headers of the test packet.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set family bridge direction egress source-udp-port 200 destination-udp-port 400
```

11. Specify the duration of each iteration in seconds, with a value from 10 seconds to 1,728,000 seconds. Also, specify the logical interface, ge-0/2/1.0, on which the RFC2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name frloss-test]
user@host# set test-iterator-duration 30 test-interface ge-0/2/0.0
```

12. Enter the exit command to go to the [edit] hierarchy level.

```
[edit services rpm rfc2544-benchmarking tests test-name test4 ]
user@host# exit
```

Configuring Other Benchmarking Test Parameters on the ACX Series Router

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the interface and bridge domain on the ACX Series router:

1. Configure the Layer 2 NNI interface on which the tests must be run from the [edit] hierarchy level.

```
[edit]
user@host# edit interfaces ge-0/2/1
```

2. Configure flexible VLAN tagging for the transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-0/2/1]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation flexible-ethernet-services
```

3. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-0/2/1]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 512
```

4. Configure the Layer 2 UNI interface.

```
[edit]
user@host# edit interfaces ge-0/2/0
```

5. Configure flexible VLAN tagging for transmission of non-tagged frames or 802.1Q single-tag and dual-tag frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-0/2/0]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation flexible-ethernet-services
```

6. Configure a logical unit for the interface and specify the encapsulation and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-0/2/0]
user@host# set unit 0 encapsulation vlan-bridge vlan-id 512
```

7. Configure the bridge domain, bd1, and specify the VLAN ID associated with the bridge domain and the associated interfaces from the [edit] hierarchy level.

```
[edit]
user@host# set bridge-domains bd1 vlan-id 10 interface ge-0/2/1.0
user@host# set bridge-domains bd1 vlan-id 10 interface ge-0/2/0.0
```

Configuring the VPLS Parameters on the MX Series Router (PE1)

Step-by-Step Procedure

The following configuration requires you to enable a simple VPLS topology between the PE1 and PE2 routers. In this example, PE1 is a MX Series router. On the PE1 router, configure the tunnel services interface and prepare the router for VPLS by configuring the BGP, MPLS, OSPF protocols.

To configure the VPLS parameters on the MX Series router:

1. Configure tunnel services.

```
[edit]
user@host# set chassis fpc 0 pic 2 tunnel-services
```

2. Configure the VPLS VLAN encapsulation on the router.

```
[edit]
user@host# set interfaces ge-0/2/1 flexible-vlan-tagging
user@host# set interfaces ge-0/2/1 mtu 9192
user@host# set interfaces ge-0/2/1 encapsulation vlan-vpls
user@host# set interfaces ge-0/2/1 unit 0 encapsulation vlan-vpls
user@host# set interfaces ge-0/2/1 unit 0 vlan-id 512
```

3. Configure the routing interface and the loopback interface on the router.

```
[edit]
user@host# set interfaces ge-0/3/0 mtu 9192
user@host# set interfaces ge-0/3/0 unit 0 family inet address 192.0.2.1/24
user@host# set interfaces ge-0/3/0 unit 0 family mpls
user@host# set interfaces lo0 unit 0 family inet address 198.51.100.1/32
```

4. Configure the routing options on the router.

```
[edit]
user@host# set routing-options router-id 198.51.100.1
user@host# set routing-options autonomous-system 65100
```

5. Configure MPLS on the router to advertise the Layer 2 VPN interface that communicates with the PE2 router.

```
[edit]
user@host# set protocols mpls interface ge-0/3/0.0
```

6. Configure BGP as the signaling protocol on the router to enable carrying of Layer 2 VPLS messages.

```
[edit]
user@host# set protocols bgp group test type internal
user@host# set protocols bgp group test local-address 198.51.100.1
user@host# set protocols bgp group test family l2vpn signaling
user@host# set protocols bgp group test neighbor 198.51.100.2
```

7. Configure OSPF on the router to enable exchange of routing information.

```
[edit]
user@host# set protocols ospf traffic-engineering
user@host# set protocols ospf reference-bandwidth 1g
user@host# set protocols ospf area 0.0.0.0 interface ge-0/3/0.0
user@host# set protocols ospf area 0.0.0.0 interface lo0.0
```

8. Configure LDP on the router to enable LDP for all connections

```
[edit]
user@host# set protocols ldp interface ge-0/3/0.0
user@host# set protocols ldp interface lo0.0
```

9. Create and configure the VPLS routing interface vpls-pe1.

```
[edit]
user@host# set routing-instances vpls-pe1 instance-type vpls
user@host# set routing-instances vpls-pe1 interface ge-0/2/1.0
user@host# set routing-instances vpls-pe1 no-local-switching
user@host# set routing-instances vpls-pe1 route-distinguisher 198.51.100.1:101
user@host# set routing-instances vpls-pe1 vrf-target target:1:2
user@host# set routing-instances vpls-pe1 protocols vpls site-range 8
user@host# set routing-instances vpls-pe1 protocols vpls no-tunnel-services
user@host# set routing-instances vpls-pe1 protocols vpls site HUB site-identifier 1
user@host# set routing-instances vpls-pe1 protocols vpls vpls-id 1
user@host# set routing-instances vpls-pe1 protocols vpls neighbor 198.51.100.2
```

Configuring Benchmarking Test Parameters on the MX104 Router (Reflector)

Step-by-Step Procedure

The following configuration requires you to configure a unique test-name for the benchmarking test on the MX104 Series router. The test-name defines the parameters for the benchmarking test to be performed. Because the test interface and test MAC addresses are the same, you can create a single test configuration at the reflector (MX104).

To configure the benchmarking test parameters on the MX104 Series router:

1. In configuration mode, at the [edit] hierarchy level, configure a real-time performance monitoring service (RPM) instance and an RFC2544-based benchmarking test for the RPM instance.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the test—for example, l2v-reflector. The test name can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking ]
user@host# set tests test-name l2v-reflector
```

3. Specify the source and destination MAC addresses of the test packet.

```
[edit services rpm rfc2544-benchmarking test-name l2v-reflector]
user@host# set source-mac-address 00:00:5e:00:53:11 destination-mac-address 00:00:5e:00:53:22
```

4. Specify the service type under test and the mode in which the test is executed, which is in-service, at the reflector. Also, specify if the IP address, TCP and UDP port must be swapped.

```
[edit services rpm rfc2544-benchmarking test-name l2v-reflector]
user@host# set service-type elan in-service ip-swap udp-tcp-port-swap
```

5. Specify the mode, which is reflect, at the reflector.

```
[edit services rpm rfc2544-benchmarking test-name l2v-reflector]
user@host# set mode reflect
```

6. Configure the family type, vpls, specify the direction, egress, and specify the protocol being transported in the Ethernet frame, for the benchmarking test. Also, specify the source and destination UDP ports and specify the logical interface, ge-0/3/1.0, on which the RFC2544-based benchmarking test is being run.

```
[edit services rpm rfc2544-benchmarking tests test-name l2v-reflector]
user@host# set family vpls direction egress source-udp-port 200 destination-udp-port 200 test-interface ge-0/3/1.0
```

Configuring Other Benchmarking Test Parameters on the MX104 Router (Reflector)

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the interface and bridge domain on the MX104 Series router:

1. Configure the Layer 2 NNI interface on which the tests must be run.

```
[edit]
user@host# edit interfaces ge-0/3/1.0
```

2. Configure flexible VLAN tagging for transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-0/3/1.0]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation vlan-vpls
```

3. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/3/1.0]
user@host# set unit 0 encapsulation vlan-vpls vlan-id 512
```

4. Configure the Layer 2 UNI interface.

```
[edit]
user@host# edit interfaces ge-0/3/2.0
```


5. Configure flexible VLAN tagging for transmission of untagged frames or 802.1Q single-tagged and dual-tagged frames on the logical interface. You can also specify the maximum transmission unit (MTU) size for the interface and the encapsulation.

```
[edit interfaces ge-0/3/2.0]
user@host# set flexible-vlan-tagging mtu 9192 encapsulation vlan-vpls
```

6. Configure a logical unit for the interface, specify the encapsulation, and configure the VLAN ID on the logical interfaces.

```
[edit interfaces ge-0/3/2.0]
user@host# set unit 0 encapsulation vlan-vpls vlan-id 512
```

7. Configure the bridge domain, bd1, and specify the VLAN ID associated with the bridge domain, and the associated interfaces from the [edit] hierarchy level.

```
[edit]
user@host# set bridge-domains bd1 vlan-id 500 interface ge-1/1/6.0
user@host# set bridge-domains bd1 vlan-id 500 interface ge-1/1/5.0
```

8. Start the benchmarking test on the reflector.

```
user@host> test services rpm rfc2544-benchmarking test l2v-reflector start
```

After the test is successfully completed at the initiator, you can stop the test at the reflector by entering the `test services rpm rfc2544-benchmarking test l2v-reflector stop` command.

Configuring VPLS Parameters on the MX104 Router (Reflector)

Step-by-Step Procedure

The following configuration requires you to enable a simple VPLS topology between the PE1 and PE2 routers. In this example, PE2 is a MX104 Series router. On the PE2 router, configure the tunnel services interface and prepare the router for VPLS by configuring the BGP, MPLS, OSPF protocols to complement the configuration on PE1.

1. Configure tunnel services.

```
[edit]
user@host# set chassis fpc 0 pic 2 tunnel-services
```

2. Configure the VPLS VLAN encapsulation on the router.

```
[edit]
user@host# set interfaces ge-0/2/5 flexible-vlan-tagging
user@host# set interfaces ge-0/2/5 mtu 9192
user@host# set interfaces ge-0/2/5 encapsulation vlan-vpls
user@host# set interfaces ge-0/2/5 unit 0 encapsulation vlan-vpls
user@host# set interfaces ge-0/2/5 unit 0 vlan-id 512
```

3. Configure the routing interface and the loopback interface on the router.

```
[edit]
user@host# set interfaces ge-0/3/0 mtu 9192
user@host# set interfaces ge-0/3/0 unit 0 family inet address 192.0.2.1/24
user@host# set interfaces ge-0/3/0 unit 0 family mpls
user@host# set interfaces lo0 unit 0 family inet address 198.51.100.1/32
```

4. Configure the routing options on the router.

```
[edit]
user@host# set routing-options router-id 198.51.100.1
user@host# set routing-options autonomous-system 100
```

5. Configure MPLS on the router to advertise the Layer 2 VPN interface that communicates with the PE1 router.

```
[edit]
user@host# set protocols mpls interface ge-0/2/5.0
```

6. Configure BGP as the signaling protocol on the router to enable carrying of Layer 2 VPLS messages.

```
[edit]
user@host# set protocols bgp group test type internal
user@host# set protocols bgp group test local-address 198.51.100.1
user@host# set protocols bgp group test family l2vpn signaling
user@host# set protocols bgp group test neighbor 198.51.100.2
```

7. Configure OSPF on the router to enable exchange of routing information.

```
[edit]
user@host# set protocols ospf traffic-engineering
user@host# set protocols ospf reference-bandwidth 1g
user@host# set protocols ospf area 0.0.0.0 interface ge-0/2/5.0
user@host# set protocols ospf area 0.0.0.0 interface lo0.0
```

8. Configure LDP on the router to enable LDP for all interfaces.

```
[edit]
user@host# set protocols ldp interface ge-0/2/5.0
user@host# set protocols ldp interface lo0.0
```

9. Create and configure the VPLS routing interface, vpls-pe2.

```
[edit]
user@host# set routing-instances vpls-pe2 instance-type vpls
user@host# set routing-instances vpls-pe2 interface ge-0/3/1.0
user@host# set routing-instances vpls-pe2 no-local-switching
user@host# set routing-instances vpls-pe2 route-distinguisher 198.51.100.1:101
user@host# set routing-instances vpls-pe2 vrf-target target:1:2
user@host# set routing-instances vpls-pe2 protocols vpls site-range 8
user@host# set routing-instances vpls-pe2 protocols vpls no-tunnel-services
user@host# set routing-instances vpls-pe2 protocols vpls site SPOKE site-identifier 1
user@host# set routing-instances vpls-pe2 protocols vpls vpls-id 1
user@host# set routing-instances vpls-pe2 protocols vpls neighbor 198.51.100.2
```

Results

In configuration mode, confirm your configuration on the ACX Series router, the MX Series router, and the MX104 Series router by entering the `show` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Benchmarking Test Parameters on the ACX Series router:

```
[edit interfaces]
  ge-0/2/0 {
    flexible-vlan-tagging;
    mtu 9192;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id 512;
    }
  }
  ge-0/2/1 {
    flexible-vlan-tagging;
    mtu 9192;
    encapsulation flexible-ethernet-services;
    unit 0 {
      encapsulation vlan-bridge;
      vlan-id 512;
    }
  }

[edit bridge-domains]
  bd1 {
    vlan-id 600;
    interface ge-0/2/1.0;
    interface ge-0/2/0.0;
  }

[edit services rpm]
  rfc2544-benchmarking {
    profiles {
      test-profile tput {
        test-type throughput
        packet-size 256;
        bandwidth-kbps 600000;
```

```

}
test-profile b2bt {
    test-type back-back-frames
    packet-size 9104;
    bandwidth-kbps 600000;
}
test-profile lty {
    test-type latency
    packet-size 1024;
    bandwidth-kbps 600000;
}
test-profile frloss {
    test-type frameloss
    packet-size 1600;
    bandwidth-kbps 6000000;
}
tests {
    test-name tput-test {
        interface ge-0/2/0.0;
        test-profile tput;
        mode initiate-and-terminate;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        ovlan-id 512;
        service-type elan;
        family bridge;
        direction egress;
        source-udp-port 200;
        destination-udp-port 400;
        test-iterator-duration 250;
    }
    test-name b2b-test {
        interface ge-0/2/0.0;
        test-profile b2bt;
        mode initiate-and-terminate;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        ovlan-id 512;
        service-type elan;
        family bridge;
        direction egress;
        destination-udp-port 400;
        test-iterator-duration 10;
    }
}

```

```

    }
    test-name lty-test {
        interface ge-0/2/0.0;
        test-profile lty;
        mode initiate-and-terminate;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        ovlan-id 512;
        service-type elan;
        family bridge;
        direction egress;
        source-udp-port 200;
        destination-udp-port 400;
        test-iterator-duration 10;
    }
    test-name frloss-test {
        interface ge-0/2/0.0;
        test-profile frloss;
        mode initiate-and-terminate;
        source-mac-address 00:00:5e:00:53:11;
        destination-mac-address 00:00:5e:00:53:22;
        ovlan-id 512;
        service-type elan;
        family bridge;
        direction egress;
        source-udp-port 200;
        destination-udp-port 400;
        test-iterator-duration 30;
    }
}

```

VPLS Parameters on the MX Series router:

```

[edit routing-instances]
vpls-instance vpls-pe1{
    instance-type vpls;
    interface ge-0/2/1.0;
    route-distinguisher 198.51.100.1:101;
    vrf-target target:1:2;
}
[edit]

```

```

protocols {
  vpls {
    vpls-id 1;
    neighbor 198.51.100.2;
    site-range 8;
    no-tunnel-services;
    site HUB {
      site-identifier 1;
    }
  }
}

```

Benchmarking Test Parameters and VPLS Parameters on the MX104 Series router:

```

[edit interfaces]
ge-0/3/1 {
  flexible-vlan-tagging;
  mtu 9192;
  encapsulation vlan-vpls;
  unit 0 {
    encapsulation vlan-vpls;
    vlan-id 512;
  }
}

ge-0/2/5 {
  flexible-vlan-tagging;
  mtu 9192;
  unit 0 {
    family inet address 203.0.113.1/24;
    family mpls;
  }
}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name l2v-reflector {
      interface ge-0/3/1.0;
      source-mac-address 00:00:5e:00:53:11;
      destination-mac-address 00:00:5e:00:53:22;
    }
  }
}

```

```

        mode reflect;
        service-type elan;
        in-service;
        ip-swap;
        udp-tcp-port swap;
        family vpls;
        reflect-etype 2048;
        direction egress;
        source-udp-port 200;
        destination-udp-port 200;
    }
}
}

[edit routing-instances]
  vpls-instance vpls-pe2 {
    instance-type vpls;
    interface ge-0/3/1;
    route-distinguisher 198.51.100.2:102;
    vrf-target target:1:2;
  }
[edit]
  protocols {
    vpls {
      vpls-id 1;
      neighbor 198.51.100.1;
      site-range 8;
      no-tunnel-services;
      site SPOKE {
        site-identifier 2;
      }
    }
  }
}

```

After you have configured the device, enter the `commit` command, in configuration mode.

Verifying the Results of the Benchmarking Test for Layer 2 ELAN Services Using VPLS

IN THIS SECTION

- [Verifying the Benchmarking Test Results | 854](#)

Examine the results of the benchmarking test that is performed on the configured service between the ACX Series router and the MX104 Series router.

Verifying the Benchmarking Test Results

Purpose

Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between the ACX Series router and the MX104 Series router.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the `show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` topic in the CLI Explorer.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Example: Configuring RFC2544-Based Benchmarking Tests on an MX104 Router for Layer 2 E-LAN Services in Bridge Domains | 788](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | 737](#)

Configuring RFC 2544-Based Benchmarking Tests on ACX Series

IN THIS CHAPTER

- [RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)
- [Layer 2 and Layer 3 RFC 2544-Based Benchmarking Test Overview | 859](#)
- [Configuring RFC 2544-Based Benchmarking Tests | 864](#)
- [Configuring Ethernet Loopback for RFC 2544-Based Benchmarking Tests | 882](#)
- [RFC 2544-Based Benchmarking Test States | 885](#)
- [Example: Configure an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services | 887](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires | 901](#)
- [Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires | 913](#)
- [Configuring a Service Package to be Used in Conjunction with PTP | 926](#)

RFC 2544-Based Benchmarking Tests for ACX Routers Overview

RFC 2544 defines a series of tests that can be used to describe the performance characteristics of network interconnecting devices. RFC2544-based benchmarking test methodology can be applied to a single device under test (DUT), or a network service (set of devices working together to provide end-to-end service). When applied to a service, the RFC2544 test results can characterize the Service-Level-Agreement (SLA) parameters.

RFC 2544 tests are performed by transmitting test packets from a device that functions as the generator or the initiator. These packets are sent to a device that functions as the reflector, which receives and returns the packets back to the initiator.

ACX Series routers support RFC 2544 tests to measure the following:

- Throughput

- Latency
- Frame loss rate
- Back-to-back frames

With embedded RFC 2544, an ACX Series router can be configured as an initiator and another ACX Series router as a reflector.

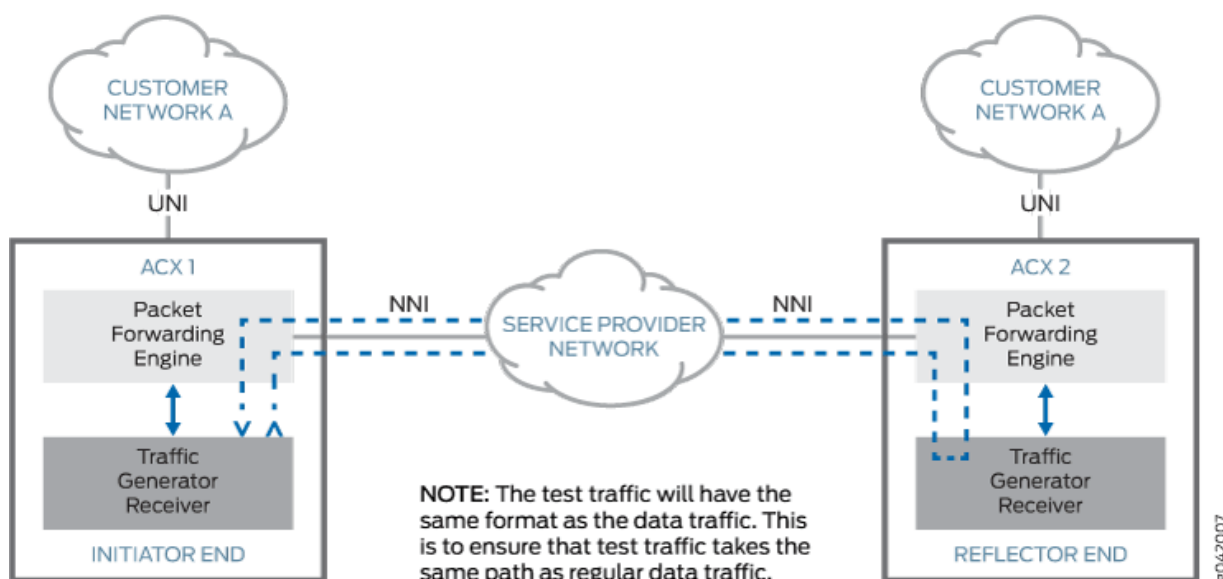
NOTE: Prior to Junos OS Evolved 22.4R1, ACX7100 routers can be configured only as a Layer 3 reflector (family inet). Starting in Junos OS Evolved 22.4R1, ACX7100 routers can also be configured as a Layer 2 reflector (family ccc or ethernet-switching.)

Starting in Junos OS Evolved 22.4R1, ACX7509 and ACX7024 routers can be configured as Layer 2 or Layer 3 reflectors.

ACX5448, ACX5048, and ACX5096 routers can be configured only as a Layer 2 reflector (family bridge or ccc). ACX5048 and ACX5096 routers support only E-Line services.

Figure 68 on page 856 shows the components, role of initiator and reflector, and the flow of test packets in an RFC 2544-based benchmarking test.

Figure 68: RFC 2544-Based Benchmarking Test Methodology



To run RFC 2544-based tests, you need a router to generate service test traffic and a router to reflect the service test traffic back. You need to:

1. Identify two service endpoints between which the RFC2544-based test needs to be run.

2. Configure the reflector end and start reflection.
3. Configure the initiator end and initiate the test.
4. Review the results after the test is complete. Test results are reported in a specific format.

On ACX Series routers, you can run the following RFC 2544-based performance measurement tests:

- Throughput test:
 - Sends a specific number of frames at a specified rate from the initiator through the network service or a DUT. The test starts with a user-configured theoretical maximum rate.
 - Counts the number of transmitted frames and the number of received frames.
 - If the number of frames received is less than those transmitted, the test is repeated with a 50 percent reduced frame rate.
 - Throughput is the maximum rate at which the count of test frames received is equal to the number of test frames transmitted through the network service.

You can repeat throughput tests for different frame sizes.

- Latency test:

NOTE: To run a latency test, you need to determine the throughput for DUT or a network service at each of the specified frame sizes.

- Starts with a stream of frames at a particular frame size through the DUT at the determined throughput rate.
- Sends an identifying tag in one frame after 60 seconds and calculate the latency when the frame with the same tag is received by the initiator.
- Is repeated for at least 20 times with the reported latency value being the average of the recorded values.

You can repeat latency tests for different frame sizes.

- Frame loss rate test:
 - Involves sending a specific number of frames at a specified rate through the DUT or a network service to be tested and counting the frames that are transmitted.
 - Calculates frame loss rate at each point using the equation:

$$((\text{input_count} - \text{output_count}) \times 100) / \text{input_count}.$$

- Runs a trial for the frame rate that corresponds to 100 percent of the configured maximum theoretical rate.
- Is repeated for the frame rate that corresponds to 90 percent of the maximum rate used and then for 80 percent of the maximum rate until a certain trial result shows no lost frames.

You repeat the frame loss rate tests for different frame sizes.

- Back-to-back frames test:
 - Involves sending a burst of frames with minimum interframe gaps through the DUT or a network service and counting the number of frames forwarded.
 - Is rerun with an increased length of burst frames if the count of transmitted frames is equal to the number of frames forwarded.
 - Is rerun with a reduced length of burst frames if the count of forwarded frames is less than the number of frames transmitted.

The back-to-back value is the number of frames in the longest burst that the DUT or a network service can handle without the loss of any frames.

You can repeat back-to-back frame tests for different frame sizes.

Starting in Junos OS Evolved 21.1R1, you can configure RFC 2544-based benchmarking tests on ACX7100 routers. To configure these tests, configure the `rfc2544` statement at the `[edit services monitoring]` hierarchy level.

To configure RFC2544 benchmarking tests for Junos OS, configure the `rfc2544-benchmarking` statement at the `[edit services rpm]` hierarchy level.

The ACX5448 router supports:

- RFC2544 egress Layer 2 reflection functionality for family bridge.
- Multiple RFC2544 reflection sessions.
- Reflection on 1G/10G/40G/Ch10G/Ch25G/100G ports.
- Ethernet Layer 2 frames to carry IP/UDP packets for RFC2544 reflection.

ACX5448 routers do not support the following RFC2544 features:

- Any interface in the bridge domain matching the bridge VLAN identifier.
- Multiple simultaneous sessions with multiple VLAN bridges.
- Multiple test sessions cannot exceed 100G bandwidth.
- IPv6 reflection.

- IPv6 filter support to identify the loopback stream.
- RFC 2544 reflection functionality for family `ccc` (PWE reflection) and family `inet` (Layer 3 IPv4 reflection).
- Reflection without MAC swap and MAC overwrite.
- Reflection on E-Line and E-LAN services.

Release History Table

Release	Description
22.4R1-EVO	Starting in Junos OS Evolved 22.4R1, we've added support for Layer 2 reflection (bridge, L2CKT, L2VPN, EVPN-VPWS, EVPN-FXC, EVPN-MPLS, and VPLS) with family <code>ccc</code> or family <code>ethernet-switching</code> to the ACX7100, ACX7509, and ACX7024 routers. We've also added support for Layer 3 reflection (IPv4, L3VPN) with family <code>inet</code> to the ACX7509 and ACX7024 routers.
21.1R1-EVO	Starting in Junos OS Evolved 21.1R1, we've added support for Layer 3 reflection (IPv4, L3VPN) with family <code>inet</code> for the ACX7100 routers.

RELATED DOCUMENTATION

[Layer 2 and Layer 3 RFC 2544-Based Benchmarking Test Overview](#) | 859

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[show services rpm rfc2544-benchmarking](#) | 1791

[show services rpm rfc2544-benchmarking test-id](#) | 1800

Layer 2 and Layer 3 RFC 2544-Based Benchmarking Test Overview

In ACX Series routers, RFC 2544-based benchmark tests can be run to measure the performance characteristics of the E-Line, E-LAN, and EVPL services.

NOTE: Prior to Junos OS Evolved 22.4R1, ACX7100 routers can be configured only as a Layer 3 reflector (family `inet`). Starting in Junos OS Evolved 22.4R1, ACX7100 routers can also be configured as a Layer 2 reflector (family `ccc` or `ethernet-switching`.)

Starting in Junos OS Evolved 22.4R1, ACX7509 and ACX7024 routers can be configured as Layer 2 or Layer 3 reflectors.

ACX5448, ACX5048, and ACX5096 routers can be configured only as a Layer 2 reflector (family bridge or ccc). ACX5048 and ACX5096 routers support only E-Line services.

- You can configure the test on the following underlying services:
 - Between two IPv4 endpoints—In this mode, the generator sends test packets to a user-configured IP destination or UDP port (which is on the reflector).
 - Between two user-to-network interfaces (UNIs) of Ethernet Virtual Connection (EVC), Ethernet Private Line (EPL, also called E-Line), Ethernet Virtual Private Line (EVPL), EVC (EPL, EVPL)—One end is configured as the generator or initiator and the other end acts as the reflector. The generator receives the test packets that are returned from the reflector and computes the test results.

NOTE: Benchmarking tests are not supported for IPv6-based services.

- You cannot perform multiple simultaneous RFC 2544-based benchmarking tests on the same pseudowire.
- Interoperation of the RFC 2544 benchmarking tests with other third-party customer premises equipment (CPE) that provides embedded or dedicated benchmarking test capability is not supported.
- Fragmented test-frames and one-way measurements of frames are not supported. You must configure one end or the source device to initiate and terminate test frames and the other end or the destination device to reflect the received frames back to the initiator.
- RFC 2544 generator and reflector are supported with testing bandwidth up to 1 Gbps. ACX5048 and ACX5096 routers supports test bandwidth of up to 40 Gbps.
- RFC2544 Layer 2 reflection supports these Layer 2 services : L2 (Bridge), L2CKT, L2VPN, EVPN-VPWS, EVPN-FXC, EVPN-MPLS, and VPLS. You can configure Layer 2 reflection only in the egress direction. Layer 2 reflection occurs at the UNI interface for unicast traffic only. The MAC addresses are always swapped after reflection. You can configure swapping for IP addresses or UDP ports using the `ip-swap` or `udp-tcp-port-swap` statements.
- RFC2544 Layer 3 reflection supports IPv4 or L3VPN traffic. You can configure Layer 3 reflection only in the ingress direction. Only traffic destined for the host is reflected; transit traffic is not affected. The IP addresses and UDP addresses are swapped after reflection.

- Supported transport mechanisms for ELAN include:
 - Multipoint Q-in-Q over provider bridged networks
 - Provider Backbone Bridge (Mac-in-Mac)
 - VPLS over IP/MPLS
 - Ethernet VPN, EVPN-MPLS, and EVPN-VXLAN
- Supported transport mechanisms for E-Line include:
 - Ethernet pseudo-wires
 - Q-in-Q
 - Provider Backbone Bridge (Mac-in-Mac)
 - Bridge domains with two logical interfaces; one for service provider and one for customer VLANs.
- The test session is supported in out-of-service mode for the underlying service. You must not transmit any traffic to the UNI port, configured as a generator or a reflector, that is being tested during the duration of the test. However, other services that are not configured for the testing session are not impacted.
- Devices embedded with benchmarking test capabilities (generators and reflectors) interoperate with other Juniper Networks devices that support the RFC 2544-based generator or reflector functionality.
- RFC 2544 generator traffic undergoes the same traffic classifier and policer or shaper processing as the ingress customer traffic from the UNI port.
- RFC 2544 generator produces a report with clear details of pass or fail for each critical testing metric, based on the configured thresholds.
- The testing packets can be configured and the format of the packet depends on the underlying service on which the test is configured. For IP-based service, the IP or port values can be configured. For Ethernet-based service, unicast untagged or VLAN ID-tagged dot1p formats (IEEE 802.1p or packet classification Layer 2 headers) are supported. The Ethernet destination address and source address that you configured are used.
- You can run RFC 2544 benchmarking `inet` tests on Layer 3 VPNs or virtual routers.
- For an `inet` service, each test session needs to use a unique UDP port. On the initiator device, the source UDP port that you specify by using the `source-udp-port` statement must be unique and not used by other UDP services that terminate at the initiator. On the reflector device, the UDP port of the destination to be used in the UDP header for the generated frames by using the `destination-udp-port` statement must be unique and not used by other UDP services that terminate at the reflector.

- You must start the test on the router that operates as the reflector before you start the test on router that functions as the initiator.
- You must configure the size of the test packet based on the configured MTU of the packets.
- For computation of the test results for a user-to-network interface (UNI) or ingress direction of an Ethernet pseudowire service, the customer edge (CE) device that is configured as a reflector for `inet` must have the reflected destination address resolved using ARP or a statically configured route must be present on the CE device to connect to the initiator.
- For benchmarking tests on the UNI direction of an Ethernet pseudowire service, if reflection mode is configured, you must configure a static ARP entry. Otherwise, the tests fail when test frames on the UNI interface are reflected. ARP resolution does not enable a successful reflection of test frames for UNI interfaces.
- For a CCC family and with the test performed in the egress or network-to-network interface (NNI) direction, the tests stop on the initiator and reflector when the pseudowire goes down.
- For an RFC 2544 test that is run in the egress or network-to-network interface (NNI) direction of an Ethernet service for a CCC family, the ingress features are not applied.
- In ACX5048 and ACX5096 routers, for a CCC family, the pseudowire has to be opened prior to the start of the RFC 2544 test and during the course of the test.
- The configured packet size denotes the untagged packet size. Any additional VLAN in the payload causes the packet length to be increased correspondingly.
- For an `inet` service, if you configure an interface on an initiator for the RFC 2544-based benchmarking test to be run without specifying the source IPv4 address for the test frames, the primary IP address of the interface is used for the test frames. If the primary IP address is not configured, the first IPv4 address of the interface is used. Similarly, for an unnumbered interface on an initiator on which the RFC 2544 test is run, the primary or the first IP address of the donor loopback interface is retrieved and used in the test frames. You must explicitly configure the source IPv4 address for the test frames by using the `source-ipv4-address` statement if you want a particular address to be used.
- RFC 2544 test generates packets for performance benchmarking testing. The packets can be destined for known or unknown unicast MAC addresses, and they can be either tagged or untagged frames. UDP/IP packet is used as the frame payload. Refer to ["Configuring RFC 2544-Based Benchmarking Tests" on page 864](#) for the frame fields that can be configured.
- Supported outer TPIDs for tagged frames are 0x8100, 0x88a8, 0x9100, and 0x9200.
- RFC 2544 benchmark tests can be run in **out-of-service** and in **in-service** modes.

NOTE: In **out-of-service** mode, while the test is running, all the data traffic sent to and from the UNI port under test on the service is interrupted. Control protocol packets are not interrupted.

In **in-service** mode, while the test is running, only the data traffic corresponding to the test session is interrupted, rest of the data traffic flow sent to and from the UNI port under test on the service are not affected. Control protocol packets are not interrupted.

- The source MAC address, destination MAC address, and the UNI port under test configured uniquely identifies the RFC 2544 benchmark test session (or test stream).
- You can run only one test at a time. Multiple simultaneous tests cannot be run at a time.
- The maximum theoretical test bandwidth supported by ACX Series routers for RFC 2544 test initiator or reflector is 1 Gbps. On ACX5048 and ACX5096 routers, the maximum theoretical test bandwidth supported for RFC 2544 reflector is 40 Gbps.
- RFC 2544 tests can be run with different frame sizes. In ACX Series routers, the supported frame sizes are 64, 68, 72, 128, 256, 512, 768, 1024, 1280, 1518, 1522, 1600, 1728, 2496, 3584, 4016, 9104, and 9136 bytes.
- The test uses round-trip traffic for performance measurement.
- A history of the test results is stored in memory.
- The test results can be copied to the local file system or a remote file system, optionally.

NOTE: RFC 2544 tests cannot compute the performance attributes of multicast or broadcast traffic streams.

For ACX routers running Junos OS Evolved, the number of RFC 2544 test sessions supported varies according to the interface speed as shown in [Table 129 on page 863](#).

Table 129: Number of RFC2544 Sessions Supported

Interface Speed	ACX7509	ACX7100	ACX7024
10G	16 sessions	16 sessions	4 sessions
25G	12 sessions	16 sessions	4 sessions

Table 129: Number of RFC2544 Sessions Supported (*Continued*)

Interface Speed	ACX7509	ACX7100	ACX7024
50G	6 sessions	16 sessions	Not supported
100G	3 sessions	16 sessions	1 session
400G	None	3 sessions	None

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[show services rpm rfc2544-benchmarking | 1791](#)

[show services rpm rfc2544-benchmarking test-id | 1800](#)

Configuring RFC 2544-Based Benchmarking Tests

IN THIS SECTION

- [Test Profile and Test Name Overview | 865](#)
- [Configure a Test Profile for an RFC 2544-Based Benchmarking Test | 871](#)
- [Configure a Test Name for an RFC 2544-Based Benchmarking Test For an Initiator | 874](#)
- [Configure a Test Name for an RFC 2544-Based Benchmarking Test For a Reflector | 878](#)
- [Start and Stop the RFC 2544-Based Benchmarking Test | 881](#)
- [Copying an RFC 2544-Based Benchmarking Test Result | 881](#)

This topic describes how to configure a test-profile and a test-name, start and stop a RFC2544-benchmark test, and copy the test result to a local or a remote file.

Test Profile and Test Name Overview

To configure a RFC 2544 benchmark test on an initiator, you must first configure a test-profile and reference the test-profile in a unique test-name. The test-name defines the parameters for the tests to be performed.

To configure a test-profile, include the test-profile *profile-name* statement at the [edit services rpm rfc2544-benchmarking] hierarchy level. Test profile is applicable only for the initiator, not the reflector.

To configure a test-name, include the test-name *test-name* statement at the [edit services rpm rfc2544-benchmarking] (Junos OS) or [edit services monitoring rfc2544 tests] (Junos OS Evolved) hierarchy level.

(Junos OS) To configure Ethernet loopback as the test mode on a logical interface, include the Ethernet-loopback statement at the [edit services rpm rfc2544-benchmarking] hierarchy level.

The table below lists the parameters for configuring the test profile at the initiator.

Table 130: Parameters for test-profile Configuration at the Initiator

Parameters	Description
test-type	RFC 2544 test type (throughput latency frame-loss back-back-frames).
packet-size	Size of the test packet. The valid packet sizes are 64, 68, 72, 128, 256, 512, 768, 1024, 1280, 1518, 1522, 1600, 1728, 2496, 3584, 4016, 9104, and 9136 bytes.
bandwidth-kbps	Define the maximum bandwidth limit, in kilobits per second (kbps). Range: 1,000 kpbs through 1,000,000 kbps.
step-percent	Specify the step percentage for frame-loss tests. Default: 10 percent Range: 1 through 100 percent

The table below lists the parameters for configuring a test-name at initiator and reflector.

Table 131: Parameters for test-name Configuration at the Initiator and the Reflector

Parameters	Description
check-test-interface-mtu	<p>When the check-test-interface-mtu parameter is configured, the software validates the MTU size of the test packets with the MTU size configured on the interface and the following would be the behavior for initiator and reflector modes:</p> <ul style="list-style-type: none"> On the initiator, if the MTU size of the test packet is larger than the MTU size configured on the interface, then the RFC2544-based benchmarking test fails to start. On the reflector, if the test packets coming to the reflector does not confirm to the MTU size configured on the interface, then these test packets do not get reflected and are dropped.
destination-ipv4-address	<p>Specify the destination IPv4 address.</p> <p>This parameter is mandatory when family inet is specified and optional when family ccc is specified.</p> <p>If a value is not specified, then by default 192.168.1.20 is used.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
destination-mac-address	<p>Specify the destination MAC address. For example, 0011.2233.4455.</p> <p>This parameter cannot be used when family inet is specified.</p> <p>This parameter is optional when family ccc and ethernet-switching is specified. If not specified, then the default value of 0x00:0x11:0xAE:0x92:0x2F:0x28 is used.</p>
destination-udp-port	<p>Specify the destination UDP port number for the test frames. Default: 4041.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
direction	<p>Specify the test direction (egress ingress). This parameter is valid only when family ccc, ethernet-switching and bridge.</p> <p>This parameter is mandatory for mode ethernet-loopback</p>
disable-signature-check	<p>Disable signature verification on the received test frames.</p>

Table 131: Parameters for test-name Configuration at the Initiator and the Reflector (Continued)

Parameters	Description
dscp-code-points	<p>Specify the value of the Differentiated Services (DiffServ) field. For example, 001111.</p> <p>If a value is not specified, then '0' is used in IP header.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
family	<p>Configure the test family (bridge ccc ethernet-switching inet).</p> <p>This parameter is mandatory for mode ethernet-loopback</p>
forwarding-class	Specify the forwarding class to be used for test frames.
halt-on-prefix-down	<p>If specified, a prefix that moves to the down state causes the corresponding tests to be stopped.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
ignore-test-interface-state	When the ignore-test-interface-state parameter is configured for RFC2544 benchmarking tests, the test continues to run even if there are any occurrences of interface up or down events. This is applicable to both initiator and reflector test modes.
in-service	<p>If specified, only the data traffic corresponding to the test session is interrupted, rest of the data traffic flow sent to and from the UNI port under test on the service are not affected.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
ivlan-cfi	<p>CFI bit used in the inner VLAN tag.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
ivlan-id	<p>Configure inner VLAN ID for the test frames.</p> <p>This parameter is valid only for family ccc mode.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>

Table 131: Parameters for test-name Configuration at the Initiator and the Reflector (Continued)

Parameters	Description
ivlan-priority	<p>Configure the priority value for the IEEE 802.1p bit in the inner VLAN tag.</p> <p>Range: 0 through 7.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
mode	<p>Specify the test mode (ethernet-loopback, initiate-and-terminate, or reflect).</p> <ul style="list-style-type: none"> ethernet-loopback—Test frames are loopbacked to the measuring device after the source MAC address and the destination MAC addresses are swapped. initiate-and-terminate—Test frames are initiated and terminated at the same end. If you specify this mode, then a reflector should be configured on the peer end to bring back the test frames. reflect—Test frames are reflected on the chosen service.
outer-tag-protocol-id	<p>TPID to be used in the outer VLAN tag.</p> <p>Supported values are 0x8100, 0x88a8, 0x9100, 0x9200.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
ovlan-cfi	<p>CFI bit used in the outer VLAN tag.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
ovlan-id	<p>Configure the outer VLAN ID for the test frames.</p> <p>Range: 0 through 4094</p> <p>This parameter is valid only for family ccc mode.</p>
ovlan-priority	<p>Configure the priority value for the IEEE 802.1p bit in the outer VLAN tag.</p> <p>Range: 0 through 7</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>

Table 131: Parameters for test-name Configuration at the Initiator and the Reflector (Continued)

Parameters	Description
packet-loss-priority	<p>Specify the packet loss priority (PLP) value.</p> <p>If a value is not configured, then the default value of low is used.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
reflect-etype	<p>Specify the EtherType ID to be used for reflection of test frames. This parameter is valid only in mode reflect. If not specified, then all EtherTypes are reflected.</p> <p>Range: 1 through 65,535.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
reflect-mode	<p>Specify the reflection mode (mac-rewrite mac-swap no-mac-swap).</p> <ul style="list-style-type: none"> • mac-rewrite—MAC values specified in source-mac-address and destination-mac-address would be used. • mac-swap—Swaps the source-mac-address and destination-mac-address in the test frame. This is the default behavior. • no-mac-swap—Does not swap MAC addresses. Test frames are returned back as-is.
reflector-port	<p>Port used to configure reflector functionality for RFC 2544 test. The range of ports that can be used based on the front panel port number are:</p> <ul style="list-style-type: none"> • On ACX5048 [16 through 53] • On ACX5096 [64 through 95, 100 through 103].
service-type	<p>Specify the service type (E-Line or E-LAN)</p>
skip-arp-iteration	<p>This parameter is valid only in family inet mode. ARP iteration is a 3-second iteration that is run for all inet tests. The results of ARP iteration are ignored in test result calculations. The primary use of sending test frames for 3 seconds is to ensure that all devices on the path to destination build their ARP entries.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>

Table 131: Parameters for test-name Configuration at the Initiator and the Reflector (Continued)

Parameters	Description
source-ipv4-address	<p>Specify the source IPv4 address used for the test frames. If a value is not specified for this parameter, then:</p> <ul style="list-style-type: none"> For family ccc, if a value is not specified, then by default 192.168.1.10 is used. For family inet, the source address of the interface is used to send out test frames. <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
source-mac-address	<p>Specify the source MAC address. For example, 0011.2233.4455</p> <p>This parameter cannot be used when family inet is specified.</p> <p>This parameter is optional when family ccc or ethernet-switching is specified. If not specified, then the default value of 0x00:0x60:0x67:0x71:0xC6:0x62 is used.</p>
source-udp-port	<p>Specify the source UDP port number for the test frames.</p> <p>Default: 4040</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
test-finish-wait-duration	<p>Number of seconds to wait after transmitting the last frame and before concluding that the test as complete.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
test-iterator-duration	<p>Specify the duration of each iteration in seconds.</p> <p>Range: 10 through 120 seconds</p> <p>The default value for test types throughput, back-to-back frames and frame loss rate is 20 seconds. The default value for test type latency is 120 seconds.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>

Table 131: Parameters for test-name Configuration at the Initiator and the Reflector (Continued)

Parameters	Description
test-interface	<p>Specify the name of the logical interface (UNI) on which the test needs to be run.</p> <p>When you specify the family as inet and mode as initiate-and-terminate the test-interface is ignored. Instead, the test is run on the egress logical interface that is determined by the route lookup on the specified destination-ipv4-address.</p> <p>When you specify the family as inet and mode as reflect, the test-interface is used as the interface to enable reflection service. If test-interface is not configured, a lookup is performed on the source-ipv4-address parameter to determine the interface hosting the address.</p> <p>This parameter is mandatory for mode ethernet-loopback.</p>
test-profile	<p>Specify the name of the test-profile to be used for the test.</p> <p>The test-profile parameter is ignored when mode reflect is used.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
vlan-cfi	<p>CFI bit used in the VLAN tag.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
vlan-id	<p>Configure the VLAN ID for the test frames.</p> <p>This parameter is valid only for mode ethernet-loopback.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>
vlan-priority	<p>Configure the VLAN priority value.</p> <p>Range: 0 through 7.</p> <p>NOTE: This parameter is not supported on ACX5048 and ACX5096 routers.</p>

Configure a Test Profile for an RFC 2544-Based Benchmarking Test

You can configure a test profile by including the test-profile *profile-name* statement at the [edit services rpm rfc2544-benchmarking] hierarchy level.

This procedure shows how to configure a test profile for Junos OS. Routers running Junos OS Evolved only support reflector mode, and so you cannot configure a test profile on these routers.

To configure a test profile:

1. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]
user@host# edit services
```

2. Configure an RPM service instance.

```
[edit services]
user@host# edit rpm
```

3. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

4. Define a name for a test profile—for example, profile1.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit profiles test-profile profile1
```

5. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1,000 Kbps through 1,000,000 Kbps. Specify a complete decimal number.

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile1]
user@host# set bandwidth-kbps kbps
```

6. Specify the size of the test packet in bytes, with a value from 64 through 9136, to be used for each test iteration. You can specify up to 10 packet sizes, separated by a space, that are used sequentially for the test. The valid packet sizes are 64, 68, 72, 128, 256, 512, 768, 1024, 1280, 1518, 1522, 1600, 1728, 2496, 3584, 4016, 9104, and 9136 bytes. If you specify a packet size other than the ones listed here as valid sizes, the configuration is saved when you commit the setting and no error message is displayed. However, when you start the test by entering the test services rpm rfc2544-benchmarking test *test-name* start command, an error message is displayed specifying that you configured an invalid packet size in the test profile associated with the test name.

NOTE:

- The minimum frame size for untagged frames should be 64.
- The minimum frame size for single-tagged frames should be 68.
- The minimum frame size for dual-tagged frames should be 72.

These values are not applicable for inet.

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile1]
user@host# set packet-size bytes
```

7. Specify the step percentage for frame-loss tests with a value from 1 through 100. This parameter is not applicable for other test types.

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile1]
user@host# set step-percent percent-value
```

8. Configure the type of test to be performed.

- To configure a throughput test, use the throughput option with the test-type statement.

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile1]
user@host# set test-type throughput
```

- To configure a latency test, use the latency option with the test-type statement.

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile1]
user@host# set test-type latency
```

- To configure a frame-loss test, use the frame-loss option with the test-type statement.

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile1]
user@host# set test-type frame-loss
```

- To configure a back-to-back frames test, use the `back-back-frames` option with the `test-type` statement.

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile1]
user@host# set test-type back-back-frames
```

Configure a Test Name for an RFC 2544-Based Benchmarking Test For an Initiator

You can configure a test name by including the `test-name test-name` statement at the `[edit services rpm rfc2544-benchmarking]` (Junos OS) or `[edit services monitoring rfc2544]` (Junos OS Evolved) hierarchy level.

Routers running Junos OS Evolved support only reflector mode.

(Junos OS) To configure a test name and define its attributes for initiator mode:

1. Navigate to the correct hierarchy level in configuration mode.

```
[edit]
user@host# edit services rpm rfc2544-benchmarking
```

2. Define a name for the test—for example, `test1`.

The test name identifier can be up to 32 characters in length. This step sets the correct hierarchy level for the rest of the steps in this procedure.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

3. Configure the destination IPv4 address for the test packets.

This parameter is required only if you configure an IPv4 family `inet`. If you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.

```
user@host# set destination-ipv4-address address
```

4. (Optional) Specify the source MAC address used in generated test frames.

You configure this statement for family `ccc`; you cannot configure it for an `inet` family. If you specify this parameter for an `inet` family, a commit error occurs when you commit the configuration. If you

do not configure the source MAC address, the default value of 0x00:0x60:0x67:0x71:0xC6:0x62 is used.

```
user@host# set source-mac-address address
```

5. Specify the destination MAC address used in generated test frames.

```
user@host# set destination-mac-address address
```

6. Specify the logical interface on which the RFC 2544-based benchmarking test is run.
This interface is a local user-to-network interface (UNI) on behalf of which the test frames are generated when the test direction is egress.

```
user@host# set test-interface interface-name
```

7. Specify the family for the benchmarking test.
The `inet` option indicates that the test is run on an IPv4 service. The `ccc` option indicates that the test is run on an CCC or Ethernet pseudowire service. The `bridge` option indicates that the test is run on a Layer 2 service.

```
user@host# set family bridge
```

8. Specify the `initiate-and-terminate` mode for the packets that are sent during the benchmarking test.
The `initiate-and-terminate` option causes the test frames to be initiated from one end and terminated at the same end. The initiation and termination mode requires a reflector to be configured at the peer end to return the test frames from the peer to the originator.

```
user@host# set mode initiate-and-terminate
```

9. Specify the direction (`egress` | `ingress`) of the interface on which the test must be run.
The `egress` option causes the test to be run in the egress direction of the interface (traffic sent from user-to-network interface (UNI) toward network-to-network interface (NNI)). The `ingress` option causes the test to be run in the ingress direction of the interface (traffic sent on user-to-network interface (UNI)). You cannot configure `ingress` for a bridge family.

```
user@host# set direction egress
```

10. Configure the outer VLAN ID for the test frames.

This statement is valid only for a CCC or an Ethernet pseudowire family.

```
user@host# set ovlan-id number
```

11. Configure the inner VLAN ID for the test frames.

This statement is valid only for a CCC or an Ethernet pseudowire family.

```
user@host# set ivlan-id number
```

12. Configure the priority value for the IEEE 802.1p bit in the outer VLAN tag.

The priority value is configured when the UNI interface is dual-tagged.

```
user@host# set ovlan-priority value
```

13. (Optional) Configure the priority value for the IEEE 802.1p bit in the inner VLAN tag.

```
user@host# set ivlan-priority value
```

14. (Optional) Configure the CFI value for the outer VLAN tag.

```
user@host# set ovlan-cfi value
```

15. (Optional) Specify the source IPv4 address to be used in generated test frames.

If you do not configure the source-ipv4-address for an `inet` family, the source address of the interface is used to transmit the test frames. If you do not configure the source-ipv4-address for a `ccc` family, the default value of 192.168.1.10 is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-ipv4-address address
```

16. Specify the destination IPv4 address to be used in generated test frames.

```
user@host# set destination-ipv4-address address
```

17. Specify the source UDP port to be used in the UDP header for the generated frames.

If you do not specify the UDP port, the default value of 4040 is used.

```
user@host# set source-udp-port port-number
```

18. Specify the destination UDP port to be used in the UDP header for the generated frames.

If you do not specify the UDP port, the default value of 4041 is used.

```
user@host# set destination-udp-port port-number
```

19. Specify the value of the Differentiated Services (DiffServ) field within the IP header of the test frames.

The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern. If you do not specify this value, 0 is used in the DSCP fields in the IP header.

```
user@host# set dscp-code-points dscp-code-bits
```

20. Specify the forwarding class to be used for test frames. The forwarding class specifies the manner in which the test frames are processed by the Packet Forwarding Engine of the router. If you do not configure this parameter, test frames are treated as best-effort traffic.

```
user@host# set forwarding-class class-name
```

21. Specify the halt-on-prefix-down option to enable a prefix that moves to the down state to cause the corresponding tests to be stopped.

The show command output for the test displays that the test was terminated because the prefix went down. By default, the RFC 2544-based benchmarking test ignores a prefix-down event (when the prefix associated with the test goes down) and continues to run.

```
user@host# set halt-on-prefix-down
```

22. Specify the duration of each iteration in seconds.

If you configure this value, the default value of each iteration depends on the type of test being run. For throughput, back-back-frames, and frame-loss types of tests, the default value is 20 seconds. For latency tests, the default value is 120 seconds.

```
user@host# set test-iterator-duration seconds
```

23. Specify the name of the test profile to be associated with a particular test name.

You must have previously configured the profile by using the `test-profile profile1` statement at the `[edit services rpm rfc2544-benchmarking]` hierarchy level. The test profile is required when the test mode is configured as initiation and termination. The `test-profile profile1` parameter is disregarded when the test mode is configured as reflection. A reflection service does not use the parameters specified in the test profile because the reflection service uses the same parameters for the test frames as the received test frames when it returns the frames to the initiator.

```
user@host# set test-profile profile1
```

Configure a Test Name for an RFC 2544-Based Benchmarking Test For a Reflector

To configure a test name and define its attributes for reflector mode:

NOTE: In ACX5048 and ACX5096 routers, while performing a RFC 2544 benchmark test, you must ensure that there are no configurations associated with the reflector port.

1. Navigate to the correct hierarchy level in configuration mode:

a. For Junos OS:

```
[edit]
user@host# edit services rpm rfc-benchmarking
```

b. For Junos OS Evolved:

```
[edit]
user@host# edit services monitoring rfc2544
```

2. Define a name for the test—for example, test1.

The test name identifier can be up to 32 characters in length. This step sets the correct hierarchy level for the rest of the steps in this procedure.

a. For Junos OS:

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

- b. For Junos OS Evolved:

```
[edit services monitoring rfc2544]
user@host# edit tests test-name test1
```

3. Specify the test mode for the packets that are sent during the benchmarking test.

The reflect option causes the test frames to be reflected back to the initiator end.

```
user@host# set mode reflect
```

4. Specify the family for the benchmarking test.

Configure the bridge option for Junos OS or the ethernet-switching option for Junos OS Evolved.

- a. For Junos OS:

```
user@host# set family bridge
```

- b. For Junos OS Evolved:

```
user@host# set family ethernet-switching
```

5. Specify the direction (egress | ingress) of the interface on which the test must be run.

The egress option causes the test to be run in the egress direction of the interface (traffic sent from user-to-network interface (UNI) toward network-to-network interface (NNI)). The ingress option causes the test to be run in the ingress direction of the interface (traffic sent on user-to-network interface (UNI)). You cannot configure ingress for a bridge or ethernet-switching family.

```
user@host# set direction egress
```

6. Configure the destination IPv4 address for the test packets.

You configure this statement only if you configure the IPv4 family `inet` option. This option is not required if you specify circuit cross-connect (CCC) or ethernet-switching as the family. If you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.

```
user@host# set destination-ipv4-address address
```

7. Specify the source MAC address used in generated test frames.

You configure this statement for a `ccc` or `ethernet-switching` family and not for an `inet` family. If you specify this parameter for an `inet` family, a commit error occurs when you commit the configuration. This parameter is optional. If you do not configure the source MAC address, the default value of `0x00:0x60:0x67:0x71:0xC6:0x62` is used.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set source-mac-address address
```

8. Specify the destination MAC address used in generated test frames.

```
user@host# set destination-mac-address address
```

9. Specify the logical interface on which the RFC 2544-based benchmarking test is run.
This interface is a local user-to-network interface (UNI) on behalf of which the test frames are generated when the test direction is egress. You cannot configure this statement for Layer 3 reflection (family `inet`).

```
user@host# set test-interface interface-name
```

10. Specify the service type as E-Line or E-LAN.

```
user@host# set service-type eline / elan
```

11. (Junos OS) Specify the forwarding class to be used for test frames.
The forwarding class specifies the manner in which the test frames are processed by the Packet Forwarding Engine of the router. If you do not configure this parameter, test frames are treated as best-effort traffic.

```
user@host# set forwarding-class class-name
```

12. (Optional) Specify the EtherType to be used for reflection of the test frames.
The EtherType is a two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. If you do not configure this statement, all EtherTypes are reflected. Use an EtherType value that matches the EtherType value set on the customer premises equipment (CPE) to which your router connects. The EtherType value appears in the Ethernet type field of the packet. It specifies the protocol being transported in the Ethernet frame.

```
user@host# set reflect-etype ethertype-value
```

13. (Optional) Specify the reflection mode for the benchmarking test.

```
user@host# set reflect-mode (mac-swap | no-mac-swap)
```

Start and Stop the RFC 2544-Based Benchmarking Test

To start an RFC 2544-based benchmarking test:

- For Junos OS, issue the test services rpm rfc2544-benchmarking test *test-name* start CLI command.
- For Junos OS Evolved, issue the test services monitoring rfc2544 test *test-name* start CLI command.

To stop an RFC 2544-based benchmarking test:

- For Junos OS, issue the test services rpm rfc2544-benchmarking test *test-name* stop CLI command.
- For Junos OS Evolved, issue the test services monitoring rfc2544 test *test-name* stop CLI command.

To start an RFC 2544 benchmarking inet tests on Layer 3 VPN or virtual router:

- For Junos OS, issue the test services rpm rfc2544-benchmarking test *test-name* routing-instance *routing-instance-name* start CLI command.
- For Junos OS Evolved, issue the test services monitoring rfc2544 test *test-name* routing-instance *routing-instance-name* start CLI command.

To stop an RFC 2544 benchmarking inet tests on Layer 3 VPN or virtual router:

- For Junos OS, issue the test services rpm rfc2544-benchmarking test *test-name* routing-instance *routing-instance-name* stop CLI command.
- For Junos OS Evolved, issue the test services monitoring rfc2544 test *test-name* routing-instance *routing-instance-name* stop CLI command.

Copying an RFC 2544-Based Benchmarking Test Result

You can copy the RFC 2544-based benchmarking test results for a particular test ID to a local or a remote file.

- To copy test results to a local file:
 - For Junos OS, issue the show services rpm rfc2544-benchmarking test-id *number* detail | save rfc-2544-test-result-session-id-*number* CLI command.
 - For Junos OS Evolved, issue the show services monitoring rfc2544 test-id *number* detail | save rfc-2544-test-result-session-id-*number* CLI command.

- To copy test results to a remote file:
 - For Junos OS, issue the `show services rpm rfc2544-benchmarking test-id number detail | save ftp://username:password@sftpchannel.example.com/rfc-2544-test-result-session-id-number.`
 - For Junos OS Evolved, issue the `show services monitoring rfc2544 test-id number detail | save ftp://username:password@sftpchannel.example.com/rfc-2544-test-result-session-id-number.`

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Layer 2 and Layer 3 RFC 2544-Based Benchmarking Test Overview | 859](#)

[show services rpm rfc2544-benchmarking | 1791](#)

[show services rpm rfc2544-benchmarking test-id | 1800](#)

Configuring Ethernet Loopback for RFC 2544-Based Benchmarking Tests

Ethernet loopback is a feature that you can use for verifying the connectivity and identifying or isolating faults in a network.

On ACX Series routers, Ethernet loopback is supported on the egress user-to-network interfaces (UNIs) direction for a bridge family configuration. In ACX Series routers, Ethernet loopback is configured on the logical interfaces. The Ethernet loopback feature can be used in performance measurements where packets are looped back to the measuring device for testing various services.

Figure 69: Testing End-to-End Service in Ethernet Loopback Mode



Figure 69 on page 882 shows a scenario where UNI-B interface is configured in Ethernet loopback mode in the egress direction. The packets received on the network-to-network interface (NNI) of the ACX Series router are forwarded to the UNI-B interface and looped back at the UNI-B interface after the source and destination MAC addresses are swapped. This is a use case for testing an end-to-end service.

You can use the following optional parameters to identify an egress traffic flow for Ethernet loopback:

- Source MAC address
- Destination MAC address
- Source IPv4 address
- Destination IPv4 address
- VLAN
- VLAN .1p priority
- EtherType
- Test iterator duration

While performing RFC2544 benchmarking tests, configure Ethernet loopback as the test mode on a logical interface by including the Ethernet-loopback CLI statement at the [edit services rpm rfc2544-benchmarking] hierarchy level.

If you configure Ethernet loopback on logical interfaces without configuring any of the optional parameters, then any unknown unicast traffic in the same bridge domain also gets looped back and does not get forwarded to other logical interfaces while the test is being performed.

When an RFC2544 benchmarking test is being performed, if the **test-iterator-duration** parameter is not configured, then Ethernet loopback continues until the test is completed or terminated.

NOTE: When performing RFC2544 benchmarking tests, you can configure the test in initiator, reflector, or loopback mode. You cannot perform the RFC2544 benchmarking tests in a combination of these test modes.

The following is a sample Ethernet loopback configuration:

```
[edit services rpm rfc2544-benchmarking]
  tests {
    test-name test1{
      source-mac-address 00:bb:cc:dd:ee:ff;
      destination-mac-address 00:11:22:33:44:55;
      vlan-id 100;
      vlan-priority 2;
      vlan-cfi 1;
      ip-swap;
      udp-tcp-port-swap;
      forwarding-class network-control;
```

```

        packet-loss-priority medium-high;
        mode ethernet-loopback;
        family bridge;
        reflect-etype 2048;
        direction egress;
        source-udp-port 2020;
        destination-udp-port 3030;
        test-iterator-duration 50;
        test-interface ge-0/1/6.0;
    }
}
[edit interfaces]
ge-0/1/4 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 1000;
        family bridge {
            filter {
                input ft1;
            }
        }
    }
}
ge-0/1/6 {
    flexible-vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 0 {
        encapsulation vlan-bridge;
        vlan-id 100;
        input-vlan-map {
            push;
            vlan-id 1000;
        }
        output-vlan-map pop;
    }
}
[edit routing-options]
ppm {
    traceoptions {
        file ppmd size 100m;
        flag packet;
        flag event;
    }
}

```

```

        flag distribute;
        flag pipe;
        flag all;
    }
}
[edit firewall]
  family bridge {
    filter ft1 {
      term t1 {
        from {
          user-vlan-id 100;
        }
        then count loopback;
      }
    }
  }
}
[edit bridge-domains]
  bd1 {
    interface ge-0/1/4.0;
    interface ge-0/1/6.0;
  }
}

```

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Layer 2 and Layer 3 RFC 2544-Based Benchmarking Test Overview | 859](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[RFC 2544-Based Benchmarking Test States | 885](#)

[Example: Configure an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services | 887](#)

[Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires | 901](#)

[Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires | 913](#)

RFC 2544-Based Benchmarking Test States

When you trigger an RFC 2544-based benchmarking test, it passes through a series of states. These states are displayed in the `Test state` field in the brief or detailed output of the `show services rpm rfc2544-`

benchmarking command. The following are the names of the states through which the test progresses after it is initiated:

1. **RFC2544_TEST_STATE_START_REQUEST**—This is the first state that all the triggered tests enter. When a test enters this state, the state denotes that a request has been sent to a Packet Forwarding Engine to start the test.
2. **RFC2544_TEST_STATE_START_FAILED**—This state indicates that the test failed to start. This state occurs when the Packet Forwarding Engine responds to the **START_REQUEST** message. The Status field of the brief or detailed output of the `show` command displays a reason for the failure. When a test enters this state, it is categorized as an terminated test.
3. **RFC2544_TEST_STATE_RUNNING**—This state occurs if the Packet Forwarding Engine is able to successfully start the test. This state indicates that the test is in progress. You can use the output of the `show` command to learn additional information about the test progress.
4. **RFC2544_TEST_STATE_STOP_REQUEST**—A test enters this state when you use the `test services rpm rfc2544-benchmarking test-id stop` command. A request is sent to the Packet Forwarding Engine to stop the test.
5. **RFC2544_TEST_STATE_STOP_FAILED**—This state is entered when the Packet Forwarding Engine failed to stop a test after it received the **STOP_REQUEST** message. The Status field displays further information regarding the exact reason for failure.
6. **RFC2544_TEST_STATE_STOPPED**—This state is entered when the Packet Forwarding Engine successfully managed to stop a test when it received the **STOP_REQUEST** message.
7. **RFC2544_TEST_STATE_COMPLETED**—This state is entered when the test successfully completes all necessary test steps.
8. **RFC2544_TEST_STATE_ABORTED_TIMEOUT**—When a request is sent to the Packet Forwarding Engine for any test, a 10-second timer control is started. If a response is not received from the Packet Forwarding Engine and the timer elapses, the test is transitioned to the **ABORTED_TIMEOUT** state. This state is introduced to prevent a test from indefinitely waiting to receive a reply from the Packet Forwarding Engine.
9. **RFC2544_TEST_STATE_RUNTIME_ERROR**—This state is entered if the Packet Forwarding Engine encounters an error when the test is running. The Status field of the brief or detailed output specifies the reason for the failure. Tests that encounter the **RUNTIME_ERROR** state are added to the count of the terminated-tests category, which can be viewed from the output of the `show services rpm rfc2544-benchmarking` command.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Layer 2 and Layer 3 RFC 2544-Based Benchmarking Test Overview | 859](#)[Configuring RFC 2544-Based Benchmarking Tests | 864](#)[show services rpm rfc2544-benchmarking | 1791](#)[show services rpm rfc2544-benchmarking test-id | 1800](#)

Example: Configure an RFC 2544-Based Benchmarking Test for Layer 3 IPv4 Services

IN THIS SECTION

- [Requirements | 887](#)
- [Overview | 887](#)
- [Configuration | 888](#)
- [Verify the Results of the Benchmarking Test for Layer 3 IPv4 Services | 900](#)

This example shows how to configure the benchmarking test for a Layer 3 IPv4 service.

NOTE: This example is not applicable for ACX5448, ACX5048, and ACX5096 routers.

Requirements

This example uses the following hardware and software components:

- An ACX Series router
- Junos OS Release 12.3X53 or later


Overview

Consider a sample topology in which a router, Router A, functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A is connected over a Layer 3 network to another router, Router B, which functions as a reflector to reflect back the test frames it receives from Router A. IPv4 is used for transmission of test frames over the Layer 3 network. This benchmarking test is used to compute the IPv4 service parameters between Router A and Router B. Logical interfaces on

both the routers are configured with IPv4 addresses to measure the performance attributes, such as throughput, latency, frame loss, and bursty frames, of network devices for the IPv4 service.

Figure 70 on page 888 shows the sample topology to perform an RFC 2544 test for a Layer 3 IPv4 service.

Figure 70: RFC 2544-Based Benchmarking Test for a Layer 3 IPv4 Service



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 889](#)
- [Configure Benchmarking Test Parameters on Router B | 889](#)
- [Configure Benchmarking Test Parameters on Router A | 893](#)
- [Results | 899](#)

In this example, you configure the benchmarking test for a Layer 3 IPv4 service that is between interface ge-0/0/0 on Router A and interface ge-0/0/4 on Router B to detect and analyze the performance of the interconnecting routers. You do not configure a test profile on Router B, because it operates as a reflector. You must configure the reflector (Router B) before you configure the initiator (Router A), because the reflector needs to be already configured and the tests running before you start tests on the initiator. If you start the tests on the initiator first, then all the packets sent are lost until you start the tests on the reflector.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configure Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 unit 0 family inet address 192.0.2.2/24
set services rpm rfc2544-benchmarking tests test-name test1 test-interface ge-0/0/4.0
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 destination-ipv4-address 192.0.2.2
set services rpm rfc2544-benchmarking tests test-name test1 destination-udp-port 4001
set rfc2544-benchmarking tests test-name test1 source-ipv4-address 192.0.2.1
```

Configure Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 unit 0 family inet address 192.0.2.1/24
set rfc2544-benchmarking profiles test-profile throughput test-type throughput
set rfc2544-benchmarking profiles test-profile throughput packet-size 64
set rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 1000
set rfc2544-benchmarking tests test-name test1 test-profile throughput
set rfc2544-benchmarking tests test-name test1 test-interface ge-0/0/0.0
set rfc2544-benchmarking tests test-name test1 mode initiate-and-terminate
set rfc2544-benchmarking tests test-name test1 family inet
set rfc2544-benchmarking tests test-name test1 destination-ipv4-address 192.0.2.2
set rfc2544-benchmarking tests test-name test1 destination-udp-port 4001
set rfc2544-benchmarking tests test-name test1 source-ipv4-address 192.0.2.1
```

Configure Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the test parameters on Router B:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

```
[edit]
user@RouterB# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@RouterB# edit ge-0/0/4
```

3. Configure a logical unit and specify the protocol family as inet.

```
[edit interfaces ge-0/0/4]
user@RouterB# edit unit 0 family inet
```

4. Specify the address for the logical interface.

```
[edit interfaces ge-0/0/4 unit 0 family inet]
user@RouterB# set address 192.0.2.2/24
```

5. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit interfaces ge-0/0/4 unit 0 family inet]
user@RouterB# up
```

6. Go to the top level of the configuration mode.

```
[edit interfaces ge-0/0/4 unit 0]
user@RouterB# top
```

7. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]
user@RouterB# edit services
```

8. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@RouterB# edit rpm
```

9. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@RouterB# edit rfc2544-benchmarking
```

10. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@RouterB# edit tests test-name test1
```

11. Specify the logical interface, ge-0/0/4.0, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set test-interface ge-0/0/4.0
```

12. Specify reflect as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set mode reflect
```

13. Configure the address type family, inet, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set family inet
```

14. Configure the destination IPv4 address for the test packets as 192.0.2.2. The destination IPv4 address configured on the reflector must match the destination IPv4 address configured on the

initiator. If you configure 192.0.2.1 instead, you get this error message: error: test test1 - Could not determine local interface for address 192.0.2.1.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set destination-ipv4-address 192.0.2.2
```

15. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set destination-udp-port 4001
```

16. Configure the source IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# set source-ipv4-address 192.0.2.1
```

17. Go to the top level of the configuration mode.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterB# top
```

18. Commit the configuration.

```
[edit]
user@RouterB# commit
```

19. Confirm the configuration. If the output does not contain the configuration below, repeat the configuration instructions in this example to correct it.

```
[edit interfaces]
ge-0/0/4 {
  unit 0 {
    family inet {
      address 192.0.2.2/24;
    }
  }
}
```

```

}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name test1 {
      test-interface ge-0/0/4.0;
      mode reflect;
      family inet;
      destination-ipv4-address 192.0.2.2;
      destination-udp-port 4001;
      source-ipv4-address 192.0.2.1
    }
  }
}

```

20. Exit to operational mode.

```

[edit]
user@RouterB# exit
user@RouterB>

```

21. Start the benchmarking test on the reflector.

```

user@RouterB> test services rpm rfc2544-benchmarking test test1 start

```

Once you configure the initiator (Router A), you can start the test on the initiator, and the initiator starts sending packets to the reflector. Once the test is successfully completed at the initiator, you can stop the test at the reflector by entering the `test services rpm rfc2544-benchmarking test test1 stop` command in operational mode.

Configure Benchmarking Test Parameters on Router A

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the test parameters on Router A:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

```
[edit]
user@RouterA# edit interfaces
```

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@RouterA# edit ge-0/0/0
```

3. Configure a logical unit and specify the protocol family.

```
[edit interfaces ge-0/0/0]
user@RouterA# edit unit 0 family inet
```

4. Specify the address for the logical interface.

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@RouterA# set address 192.0.2.1/24
```

5. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@RouterA# up
```

6. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/0 unit 0]
user@RouterA# top
```

7. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]
user@RouterA# edit services
```

8. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@RouterA# edit rpm
```

9. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@RouterA# edit rfc2544-benchmarking
```

10. Define a name for a test profile—for example, throughput.

```
[edit services rpm rfc2544-benchmarking]
user@RouterA# edit profiles test-profile throughput
```

11. Configure the type of test to be performed as throughput.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# set test-type throughput
```

12. Specify the size of the test packet as 64 bytes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# set test-type packet-size 64
```

13. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1,000 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# set test-type bandwidth-kbps 1000
```

14. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@RouterA# up
```

15. Enter the `up` command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@RouterA# up
```

16. Define a name for the test—for example, `test1`. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@RouterA# edit tests test-name test1
```

17. Specify the name of the test profile—for example, `throughput`—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set test-profile throughput
```

18. Specify the logical interface, `ge-0/0/0.0`, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set test-interface ge-0/0/0.0
```

19. Specify the test mode for the packets that are sent during the benchmarking test as `initiate` and `terminate`.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set mode initiate-and-terminate
```

20. Configure the address type family, `inet`, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set family inet
```

21. Configure the destination IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set destination-ipv4-address 192.0.2.2
```

22. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set destination-udp-port 4001
```

23. Configure the source IPv4 address for the test packets.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# set source-ipv4-address 192.0.2.1
```

24. Go to the top level of the configuration mode.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@RouterA# top
```

25. Commit the configuration.

```
[edit]
user@RouterA# commit
```

26. Confirm the configuration. If the output does not contain the configuration below, repeat the configuration instructions in this example to correct it.

```
[edit]
user@RouterA# show
[edit interfaces]
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 192.0.2.1/24;
      }
    }
  }
```

```

    }
}

[edit services rpm]
rfc2544-benchmarking {
  profiles {
    test-profile throughput {
      test-type throughput
      packet-size 64;
      bandwidth-kbps 1000;
    }
  }

  tests {
    test-name test1 {
      test-profile throughput;
      interface ge-0/0/0.0;
      mode initiate-and-terminate;
      family inet;
      destination-ipv4-address 192.0.2.2
      destination-udp-port 4001;
      source-ipv4-address 192.0.2.1
    }
  }
}

```

27. Exit to operational mode.

```

[edit]
user@RouterA# exit
user@RouterA>

```

28. Start the benchmarking test on the initiator.

```

user@RouterA> test services rpm rfc2544-benchmarking test test1 start

```

After the test successfully completes, it automatically stops at the initiator. Once the test is successfully completed at the initiator, you can stop the test at the reflector by entering the test services rpm rfc2544-benchmarking test test1 stop command on Router B in operational mode.

Results

If you have not done so already, confirm your configuration on Router A and Router B by entering the `show` command in configuration mode at the [edit interfaces] and [edit services rpm] hierarchy levels. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Configuration for Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family inet {
      address 192.0.2.1/24;
    }
  }
}

[edit services rpm]
rfc2544-benchmarking {
  profiles {
    test-profile throughput {
      test-type throughput
      packet-size 64;
      bandwidth-kbps 1000;
    }
  }

  tests {
    test-name test1 {
      test-profile throughput;
      interface ge-0/0/0.0;
      mode initiate-and-terminate;
      family inet;
      destination-ipv4-address 192.0.2.2
      destination-udp-port 4001;
      source-ipv4-address 192.0.2.1
    }
  }
}
```

Configuration for Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
  unit 0 {
    family inet {
      address 192.0.2.2/24;
    }
  }
}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name test1 {
      test-interface ge-0/0/4.0;
      mode reflect;
      family inet;
      destination-ipv4-address 192.0.2.2;
      destination-udp-port 4001;
      source-ipv4-address 192.0.2.1
    }
  }
}
```

Verify the Results of the Benchmarking Test for Layer 3 IPv4 Services

IN THIS SECTION

- [Verify the Benchmarking Test Results | 901](#)

Examine the results of the benchmarking test performed on the configured service between Router A and Router B.

Verify the Benchmarking Test Results

Purpose

Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.

Action

In operational mode, enter the `show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests | summary)` command, on either the initiator or the reflector, to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed.

RELATED DOCUMENTATION

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

[profiles \(RFC 2544 Benchmarking\) | 1324](#)

[tests | 1459](#)

[show services rpm rfc2544-benchmarking | 1791](#)

[show services rpm rfc2544-benchmarking test-id | 1800](#)

Example: Configuring an RFC 2544-Based Benchmarking Test for NNI Direction of Ethernet Pseudowires

IN THIS SECTION

- [Requirements | 902](#)
- [Overview | 902](#)

● Configuration | 903

● Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service | 912

This example shows how to configure the benchmarking test for a network-to-network interface (NNI) direction of an Ethernet pseudowire service.

Requirements

This example uses the following hardware and software components:

- An ACX Series router
- Junos OS Release 12.3X52 or later


Overview

Consider a sample topology in which a router, Router A, functions as an initiator and terminator of the test frames for an RFC 2544-based benchmarking test. Router A operates as a provider edge device, PE1, which is connected to a customer edge device, CE1, on one side and over an Ethernet pseudowire to another router, Router B, which functions as a reflector to reflect back the test frames it receives from Router A. Router B operates as a provider edge device, PE2, which is the remote router located at the other side of the service provider core. The UNI direction of CE1 is connected to the NNI direction of PE1. An MPLS tunnel connects PE1 and PE2 over the Ethernet pseudowire or the Ethernet line (E-LINE).

This benchmarking test is used to compute the performance attributes in the network-to-network interface (NNI) direction of an Ethernet pseudowire service between Router A and Router B. The logical interface under test on Router A is the CE1 interface with UNI as the direction, and the logical interface under test on Router B is the CE2 interface with NNI as the direction. Data traffic arriving from UNI towards NNI is ignored while the test is in progress. Packets from NNI are not sent toward the customer edge because all packets are assumed to be test frames. The CCC family and NNI direction are configured on routers A and B.

Figure 71 on page 903 shows the sample topology to perform an RFC 2544 test for the NNI direction of an Ethernet pseudowire service.

Figure 71: RFC 2544-Based Benchmarking Test for NNI Direction of an Ethernet Pseudowire



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 903](#)
- [Configuring Benchmarking Test Parameters on Router B | 904](#)
- [Configuring Benchmarking Test Parameters on Router B | 908](#)
- [Results | 911](#)

In this example, you configure the benchmarking test for the NNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/0 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration 20
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name test1 family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction nni
```

Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 reflector-port 25
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction uni
```

Configuring Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the test parameters on Router A:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/0
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```

8. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]  
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]  
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]  
user@host# edit rfc2544-benchmarking
```

11. Define a name for a test profile—for example, throughput.

```
[edit services rpm rfc2544-benchmarking]  
user@host# edit profiles test-profile throughput
```

12. Configure the type of test to be performed as throughput.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# set test-type throughput
```

13. Specify the size of the test packet as 64 bytes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# set test-type packet-size 64
```

14. Specify the period—for example, 20 minutes—for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds).

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```

15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```

16. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# up
```

17. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```

18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-profile throughput
```

20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, ccc, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

23. Specify the direction of the interface on which the test must be run, which is NNI in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction nni
```

Configuring Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the test parameters on Router B:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# top
```

8. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]
user@host# edit services
```


9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

11. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

12. Specify the logical interface, ge-0/0/4.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/4.1
```

13. Specify reflect as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

14. Configure the address type family, ccc, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

15. Specify the direction of the interface on which the test must be run, which is NNI in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction nni
```

Results

In configuration mode, confirm your configuration on Router A and Router B by entering the `show` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Configuring Benchmarking Test Parameters on Router A:

```
[edit interfaces]
  ge-0/0/0 {
    vlan-tagging;
    unit 0 {
      encapsulation vlan-ccc;
      vlan-id 101;
    }
  }

[edit services rpm]
  rfc2544-benchmarking {
    profiles {
      test-profile throughput {
        test-type throughput
        packet-size 64;
        test-duration 20m;
        bandwidth-kbps 500;
      }
    }

    tests {
      test-name test1 {
        interface ge-0/0/0.1;
        test-profile throughput;
        mode initiate-and-terminate;
        family ccc;
        direction nni;
      }
    }
  }
}
```

Configuring Benchmarking Test Parameters on Router B:

```
[edit interfaces]
ge-0/0/4 {
  vlan-tagging;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 101;
  }
}

[edit services rpm]
rfc2544-benchmarking {
  # Note, When in reflector mode, test profile is not needed
  tests {
    test-name test1 {
      interface ge-0/0/4.1;
      mode reflect;
      family ccc;
      direction nni;
    }
  }
}
```

After you have configured the device, enter the `commit` command in configuration mode.

Verifying the Results of the Benchmarking Test for NNI Direction of an Ethernet Pseudowire Service

IN THIS SECTION

- [Verifying the Benchmarking Test Results | 913](#)

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

Verifying the Benchmarking Test Results

Purpose

Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.

Action

In operational mode, enter the `run show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the `run show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` in the [CLI Explorer](#).

RELATED DOCUMENTATION

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

[profiles \(RFC 2544 Benchmarking\) | 1324](#)

[tests | 1459](#)

[show services rpm rfc2544-benchmarking | 1791](#)

[show services rpm rfc2544-benchmarking test-id | 1800](#)

Example: Configuring an RFC 2544-Based Benchmarking Test for UNI Direction of Ethernet Pseudowires

IN THIS SECTION

● [Requirements | 914](#)

- [Overview | 914](#)
- [Configuration | 915](#)
- [Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service | 925](#)

This example shows how to configure the benchmarking test for the user-to-network interface (UNI) direction of an Ethernet pseudowire service.

Requirements

This example uses the following hardware and software components:

- An ACX Series router
- Junos OS Release 12.3X53 or later

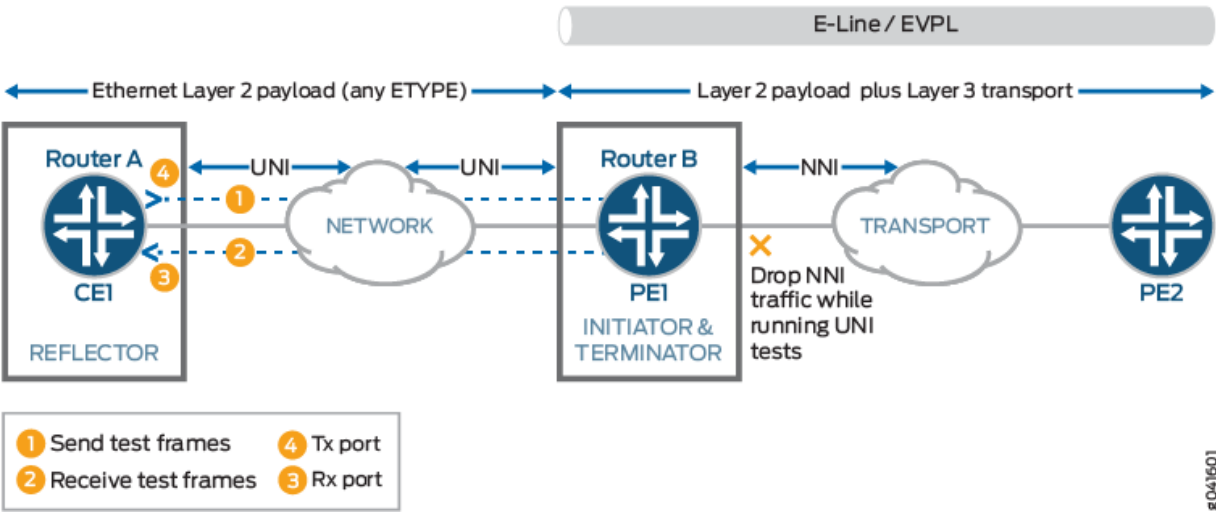
Overview

Consider a sample topology in which a router, Router A, functions as a reflector of the test frames for an RFC 2544-based benchmarking test. The logical customer edge (CE)-facing interface and inet family are configured on Router A. Router A is not part of a pseudowire and therefore, a Layer 3 family configuration is required on it. Router A, which is a customer edge device, CE1, is connected to Router B, which functions as a provider edge device, PE1, over an Ethernet pseudowire in the UNI direction with EtherType or Layer 2 Ethernet payload. The logical interface, CCC family, and UNI direction are configured on Router B. Router B or PE1 is connected over an Ethernet pseudowire in the NNI direction to a provider edge device at the remote site, PE2. The link between CE1 and PE1 is an Ethernet Layer 2 network and it can be configured with any EtherType value. The link between PE1 and PE2 is an Ethernet line (E-LINE) or an Ethernet Private Line (EPL) that has Layer 2 payload and Layer 3 transport sent over it. Router B or PE1 functions as an initiator and terminator of the test frames that are sent to Router A and reflected back from it.

This benchmarking test is used to compute the performance attributes in the user-to-network interface (UNI) direction of an Ethernet pseudowire service between Router A and Router B. Data traffic arriving from a network-to-network interface (NNI) toward the customer edge is ignored while the test is in progress. Packets from the CE are not sent toward the NNI because all packets are assumed to be test probes.

Figure 72 on page 915 shows the sample topology to perform an RFC 2544 test for the UNI direction of an Ethernet pseudowire service.

Figure 72: RFC 2544-Based Benchmarking Test for UNI Direction of an Ethernet Pseudowire



Configuration

IN THIS SECTION

- [CLI Quick Configuration | 915](#)
- [Configuring Benchmarking Test Parameters on Router A | 916](#)
- [Configuring Benchmarking Test Parameters on Router B | 920](#)
- [Results | 923](#)

In this example, you configure the benchmarking test for the UNI direction of an Ethernet pseudowire service that is enabled between two routers to detect and analyze the performance of the interconnecting routers.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them in a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

Configuring Benchmarking Test Parameters on Router A

```
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 vlan-id 101
set interfaces ge-0/0/0 unit 0 family inet address 200.0.0.1/24
set services rpm rfc2544-benchmarking profiles test-profile throughput test-type throughput
set services rpm rfc2544-benchmarking profiles test-profile throughput packet-size 64
set services rpm rfc2544-benchmarking profiles test-profile throughput test-duration 20m
set services rpm rfc2544-benchmarking profiles test-profile throughput bandwidth-kbps 500
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/0.1
set services rpm rfc2544-benchmarking tests test-name test1 test-profile throughput
set services rpm rfc2544-benchmarking tests test-name test1 mode initiate-and-terminate
set services rpm rfc2544-benchmarking tests test-name test1 family inet
set services rpm rfc2544-benchmarking tests test-name test1 dest-address 10.200.0.2
set services rpm rfc2544-benchmarking tests test-name test1 udp-port 4001
```

Configuring Benchmarking Test Parameters on Router B

```
set interfaces ge-0/0/4 vlan-tagging
set interfaces ge-0/0/4 unit 0 encapsulation vlan-ccc
set interfaces ge-0/0/4 unit 0 vlan-id 101
set services rpm rfc2544-benchmarking tests test-name test1 interface ge-0/0/4.1
set services rpm rfc2544-benchmarking tests test-name test1 mode reflect
set services rpm rfc2544-benchmarking tests test-name test1 mode family ccc
set services rpm rfc2544-benchmarking tests test-name test1 direction uni
```

Configuring Benchmarking Test Parameters on Router A

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the test parameters on Router A:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/0
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/0]
user@host# set vlan-tagging
```

4. Configure a logical unit and specify the protocol family as inet.

```
[edit interfaces ge-0/0/0]
user@host# edit unit 0 family inet
```

5. Specify the address for the logical interface.

```
[edit interfaces ge-0/0/0 unit 0 family inet]
user@host# set address 10.200.0.1/24
```

6. Configure the VLAN ID on the logical interface as 101.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/0 unit 0]
user@host# top
```


8. In configuration mode, go to the [edit services] hierarchy level.

```
[edit]  
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]  
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]  
user@host# edit rfc2544-benchmarking
```

11. Define a name for a test profile—for example, throughput.

```
[edit services rpm rfc2544-benchmarking]  
user@host# edit profiles test-profile throughput
```

12. Configure the type of test to be performed as throughput.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# set test-type throughput
```

13. Specify the size of the test packet as 64 bytes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]  
user@host# set test-type packet-size 64
```

14. Specify the period for which the test is to be performed in hours, minutes, or seconds by specifying a number followed by the letter h (for hours), m (for minutes), or s (for seconds). In this example, you configure the period as 20 minutes.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type test-duration 20m
```

15. Define the theoretical maximum bandwidth for the test in kilobits per second, with a value from 1 Kbps through 1,000,000 Kbps.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# set test-type bandwidth-kbps 500
```

16. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles test-profile throughput]
user@host# up
```

17. Enter the up command to go the previous level in the configuration hierarchy.

```
[edit services rpm rfc2544-benchmarking profiles]
user@host# up
```

18. Define a name for the test—for example, test1. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

19. Specify the name of the test profile—for example, throughput—to be associated with a particular test name.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-profile throughput
```

20. Specify the logical interface, ge-0/0/0.1, on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/0.1
```

21. Specify the test mode for the packets that are sent during the benchmarking test as initiation and termination.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode initiate-and-terminate
```

22. Configure the address type family, inet, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family inet
```

23. Configure the destination IPv4 address for the test packets as 200.0.0.2.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set dest-address 200.0.0.2
```

24. Specify the UDP port of the destination to be used in the UDP header for the generated frames as 4001.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set udp-port 4001
```

Configuring Benchmarking Test Parameters on Router B

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode*.

To configure the test parameters on Router B:

1. In configuration mode, go to the [edit interfaces] hierarchy level:

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

2. Configure the interface on which the test must be run.

```
[edit interfaces]
user@host# edit ge-0/0/4
```

3. Configure VLAN tagging for transmission and reception of 802.1Q VLAN-tagged frames.

```
[edit interfaces ge-0/0/4]
user@host# set vlan-tagging
```

4. Configure a logical unit for the interface.

```
[edit interfaces ge-0/0/4]
user@host# edit unit 0
```

5. Specify the encapsulation for Ethernet VLAN circuits.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set encapsulation vlan-ccc
```

6. Configure the VLAN ID as 101 on the logical interface.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# set vlan-id 101
```

7. Go to the top level of the configuration command mode.

```
[edit interfaces ge-0/0/4 unit 0]
user@host# top
```

8. In configuration mode, go to the `[edit services]` hierarchy level.

```
[edit]
user@host# edit services
```

9. Configure a real-time performance monitoring service (RPM) instance.

```
[edit services]
user@host# edit rpm
```

10. Configure an RFC 2544-based benchmarking test for the RPM instance.

```
[edit services rpm]
user@host# edit rfc2544-benchmarking
```

11. Define a name for the test—for example, `test1`. The test name identifier can be up to 32 characters in length.

```
[edit services rpm rfc2544-benchmarking]
user@host# edit tests test-name test1
```

12. Specify the logical interface on which the RFC 2544-based benchmarking test is run.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set test-interface ge-0/0/4.1
```

13. Specify `reflect` as the test mode for the packets that are sent during the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set mode reflect
```

14. Configure the address type family, `ccc`, for the benchmarking test.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set family ccc
```

15. Specify the direction of the interface on which the test must be run, which is UNI in this example.

```
[edit services rpm rfc2544-benchmarking tests test-name test1]
user@host# set direction uni
```

Results

In configuration mode, confirm your configuration on Router A and Router B by entering the `show` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

Configuring Benchmarking Test Parameters on Router A:

```
[edit interfaces]
ge-0/0/0 {
  vlan-tagging;
  unit 0 {
    vlan-id 101;
    family inet {
      address 10.200.0.1/24;
    }
  }
}

[edit services rpm]
rfc2544-benchmarking {
  profiles {
    test-profile throughput {
      test-type throughput
      packet-size 64;
      test-duration 20m;
      bandwidth-kbps 500;
    }
  }

  tests {
    test-name test1 {
      interface ge-0/0/0.1;
      test-profile throughput;
      mode initiate-and-terminate;
```

```

        family inet;
        dest-address 10.200.0.2
        udp-port 4001;
    }
}
}

```

Configuring Benchmarking Test Parameters on Router B:

```

[edit interfaces]
ge-0/0/4 {
    vlan-tagging;
    unit 0 {
        encapsulation vlan-ccc;
        vlan-id 101;
    }
}

[edit services rpm]
rfc2544-benchmarking {
    # Note, When in reflector mode, test profile is not needed
    tests {
        test-name test1 {
            interface ge-0/0/4.1;
            mode reflect;
            family ccc;
            direction uni;
        }
    }
}

```

After you have configured the device, enter the `commit` command in configuration mode.

Verifying the Results of the Benchmarking Test for UNI Direction of an Ethernet Pseudowire Service

IN THIS SECTION

- [Verifying the Benchmarking Test Results | 925](#)

Examine the results of the benchmarking test that is performed on the configured service between Router A and Router B.

Verifying the Benchmarking Test Results

Purpose

Verify that the necessary and desired statistical values are displayed for the benchmarking test that is run on the configured service between Router A and Router B.

Action

In operational mode, enter the `run show services rpm rfc2544-benchmarking (aborted-tests | active-tests | completed-tests | summary)` command to display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance.

Meaning

The output displays the details of the benchmarking test that was performed. For more information about the `run show services rpm rfc2544-benchmarking operational` command, see `show services rpm rfc2544-benchmarking` in the [CLI Explorer](#).

RELATED DOCUMENTATION

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

[profiles \(RFC 2544 Benchmarking\) | 1324](#)

[tests | 1459](#)

[show services rpm rfc2544-benchmarking | 1791](#)

[show services rpm rfc2544-benchmarking test-id | 1800](#)

Configuring a Service Package to be Used in Conjunction with PTP

On ACX1100 routers, you can configure a service package on the router for the RFC 2544-based benchmarking test, or for NAT and IPsec applications. When you configure the service package for the RFC 2544-based benchmarking test or for the NAT and IPsec applications, a reboot of the Forwarding Engine Board (FEB) occurs to apply the service package selection. By default, the service package for RFC 2544 benchmarking test is selected. The selection of a service package is needed on ACX1100 routers when you configure such routers for IEEE 1588v2 Precision Time Protocol (PTP) because both RFC 2544-based benchmarking tests and a combination of NAT and IPsec protocols are not supported simultaneously; you can configure only PTP and RFC 2544-based tests, or PTP and the combination of NAT and IPsec at a point in time.

You need to specify the service package to be RFC 2544-based or NAT and IPsec-based only for ACX1100-AC routers. The selection of a service package is not needed on ACX Series routers other than the ACX1100-AC and ACX500 routers because on such routers, only the RFC 2544-based benchmarking tests are supported; NAT and IPsec applications are not supported on those routers.

To configure the RFC 2544-based service package on a particular FPC, include the `service-package bundle-rfc2544` statement at the `[edit chassis fpc slot-number]` hierarchy level.

```
[edit chassis]
fpc slot-number {
    service-package bundle-rfc2544;
}
```

To configure the NAT and IPsec applications service package on a particular FPC, include the `service-package bundle-nat-ipsec` statement at the `[edit chassis fpc slot-number]` hierarchy level.

```
[edit chassis]
fpc slot-number {
    service-package bundle-nat-ipsec;
}
```

RELATED DOCUMENTATION

[service-package](#)

Tracking Streaming Media Traffic Using Inline Video Monitoring

IN THIS CHAPTER

- [Understanding Inline Video Monitoring on MX Series Routers | 927](#)
- [Configuring Inline Video Monitoring on MX Series Routers | 934](#)
- [Inline Video Monitoring Syslog Messages on MX Series Routers | 946](#)
- [Generation of SNMP Traps and Alarms for Inline Video Monitoring on MX Series Routers | 947](#)
- [SNMP Traps for Inline Video Monitoring Statistics on MX Series Routers | 951](#)
- [Processing SNMP GET Requests for MDI Metrics on MX Series Routers | 952](#)

Understanding Inline Video Monitoring on MX Series Routers

Junos OS supports inline video monitoring using media delivery index (MDI) metrics.

Before you use the inline video monitoring feature, ensure that you understand the following terms:

- **media delivery index**—MDI metrics facilitate identification of buffering needs for streaming media. Buffering must be adequate to compensate for packet jitter, measured by the MDI delay factor, and quality problems indicated by lost packets, measured by the MDI media loss rate (MLR). By performing measurements under varying load conditions, you can identify sources of significant jitter or packet loss and take appropriate action.
- **delay factor**—Delay factor is the maximum observed time difference between the arrival of media data and the drain of media data. The expected drain rate is the nominal, constant traffic rate for constant bit rate streams or the computed traffic rate of variable rate media stream packet data.

For typical stream rates of 1 megabit per second and higher, an interval of one second provides an adequate sample time. The delay factor indicates how long a data stream must be buffered (delayed) at its nominal bit rate to prevent packet loss.

The delay factor suggests the minimum size of the buffer required at the next downstream node. As a stream progresses, the variation of the delay factor indicates packet bunching or packet gaps

(jitter). Greater delay factor values also indicate that more network latency is needed to deliver a stream due to the need to pre-fill a receive buffer before beginning the drain to guarantee no underflow.

When the nominal drain bit rate at a receiving node is known, the delay factor's maximum indicates the size of buffer required to accommodate packet jitter.

- **Media rate variation (MRV)**—This value is the difference between the expected packet rate and actual packet rate, expressed as a percentage of the expected packet rate.
- **Media loss rate (MLR)**—This value is the number of media packets lost over a configurable time interval (*interval-duration*), where the flow packets carry streaming application information. A single IP packet can contain one or more streaming packets. For example, an IP packet typically contains seven 188-byte MPEG transport stream packets. In this case, a single IP packet loss results in seven lost packets counted (if those seven lost packets did not include null packets). Including out-of-order packets is important, because many consumer-type streaming devices do not attempt to reorder packets that are received out of order.

To configure the monitoring process, define criteria templates and apply them to the interfaces and flows you want to monitor. Monitoring templates include the following criteria:

- Duration of each measurement cycle
- Flow rate information used to establish expected flow rates
- Threshold levels for delay factor, media rate variation, and media loss rate that trigger desired system log alerts

For each interface you want to monitor, you can define one or more filters to select IPv4 flows for monitoring. Flows are designated as input or output flows. Starting in Junos OS Release 17.2R1, you can identify IPv4-over-MPLS flows. Starting in Junos OS Release 17.4R1, you can identify IPv6 flows and IPv6-over-MPLS flows. Starting in Junos OS Release 19.1R1, you can configure MX Series routers for inline video monitoring of uncompressed HD or 4K stream video (Payload Type 98 and 99). MDI functionality has been extended to video flows such as ST 2000-5 (RTP PT 98) and ST 2000-6 (RTP PT 99). These are non-MPEG video flows over IP/UDP/RTP and are constant bit-rate flows. The operator would specify proper IP addresses and UDP ports so that non-video flows over RTP will not go through MDI processing.

MPLS flows with more than three labels cannot be monitored.

IPv4 flows are uniquely identified by:

- Destination IP address
- Destination port
- Source IP address

- Source port
- Direction
- Interface index
- Media type (RTP or MPEG)

IPv4-over-MPLS flows are uniquely identified by:

- The top three MPLS labels
- Destination IP address
- Destination port
- Source IP address
- Source port
- Direction
- Interface index
- Media type (RTP or MPEG)

IPv6 flows are uniquely identified by:

- Destination IP address
- Destination port
- Source IP address
- Source port
- Direction
- Interface index
- Media type (RTP or MPEG)

IPv6-over-MPLS flows are uniquely identified by:

- The top three MPLS labels
- Destination IP address
- Destination port
- Source IP address

- Source port
- Direction
- Interface index
- Media type (RTP or MPEG)

Junos OS supports the definition of filters for up to 256 flows on an interface, which can consist of input flows, output flows, or a combination of input and output flows. These filters provide criteria for selecting flows for monitoring. If the selection criteria consist of lists of IP addresses or ports, you can exceed the maximum number of match conditions for flows. Video monitoring selects a widely variable number of flows based on flow filters.

The total number of destination IP addresses configured in a flow for an interface cannot exceed 32, and the total number of source IP addresses configured in a flow for an interface cannot exceed 32.

Inline video monitoring is not supported when you enable Next Gen Services on an MX Series router.

Inline video monitoring is available on MX Series 5G Universal Routing Platforms using only the following MPCs:

- MPC1
- MPC1E
- MPC2
- MPC2E
- MPC2E-NG
- MPC3E-NG
- MPC-16XGE
- MPC5E
- MPC6E
- MPC7E
- MPC8E
- MPC9E
- MPC10E
- MPC11E

NOTE: Traffic throughput is reduced below the interface bandwidth when video monitoring is used with an MPC2E-NG or MPC3E-NG in the following scenario:

- The input and output ports are on the same slot.
- The input-flows is configured as inet and the output-flows is configured as mpls.
- At least one flow has a traffic rate greater than 2 Gbps.

To avoid this reduced throughput, use input and output ports on different slots.

Starting in Junos OS Release 16.1R1, you can configure the number of flows that can be measured per Packet Forwarding Engine at a time, up to a value of 8192. The maximum configured number of flows that can be measured for each MPC model is shown in the second column of [Table 132 on page 931](#). The default number of flows that can be measured for each MPC model is shown in the third column of [Table 132 on page 931](#). In Junos OS Release 15.1 and earlier, you cannot configure the number of flows that can be measured.

When you do not define input or output flow filters for a monitored interface, all flows on the interface are subject to monitoring.

Table 132: MPC Flow Monitoring Capacity by Model

MPC Model	Maximum Configurable Number of Flows Monitored Simultaneously (Starting in Junos OS Release 16.1)	Default Number of Flows Monitored Simultaneously
MPC1	8000	1000
MPC1E	8000	1000
MPC2	16,000	2000
MPC2E	16,000	2000
MPC2E-NG	8000	1000
MPC3E-NG	8000	1000

Table 132: MPC Flow Monitoring Capacity by Model (Continued)

MPC Model	Maximum Configurable Number of Flows Monitored Simultaneously (Starting in Junos OS Release 16.1)	Default Number of Flows Monitored Simultaneously
MPC-16XGE	32,000	4000
MPC5E	40,000	5000
MPC6E	40,000	5000
MPC7E	40,000	5000
MPC8E	40,000	5000
MPC9E	40,000	5000
MPC10E	24,000 (starting in Junos OS Release 20.3R1)	3000
MPC11E	64,000 (starting in Junos OS Release 20.3R1)	8000

NOTE: Junos OS measures both UDP flows (the default) and RTP flows. Junos OS differentiates media traffic over UDP or RTP by inspecting the first byte in the UDP payload. If the first byte of the UDP payload is 0x47 (MPEG2-TS sync byte), the traffic is treated as media traffic over UDP. Traffic is treated as media traffic over RTP if the version field is 2 and the payload type is 33 in the RTP header. When neither of these criteria are met, the packet is not considered for video monitoring.

Starting in Junos OS Release 15.1R1, MX Series routers support the inline video monitoring to measure media delivery index (MDI) metrics that can be accessed using the SNMP GET operation. Currently, inline MDI can generate only a system log when the computed value is not within the configured range. SNMP is primarily used to monitor alarms raised by the inline video monitoring feature. The alarms are monitored in the network management systems either to troubleshoot the problem or to diagnose degradation in video quality.

You use the video-monitoring statement at the [edit services] hierarchy level to specify monitoring criteria for two key indicators of video traffic problems: delay factor and media loss rate (MLR), and to apply these metrics to flows on designated interfaces.

Release History Table

Release	Description
19.3R2	Inline video monitoring is not supported when you enable Next Gen Services on an MX Series router.
19.1R1	Starting in Junos OS Release 19.1R1, you can configure MX Series routers for inline video monitoring of uncompressed HD or 4K stream video (Payload Type 98 and 99). MDI functionality has been extended to video flows such as ST 2000-5 (RTP PT 98) and ST 2000-6 (RTP PT 99). These are non-MPEG video flows over IP/UDP/RTP and are constant bit-rate flows. The operator would specify proper IP addresses and UDP ports so that non-video flows over RTP will not go through MDI processing.
17.4R1	Starting in Junos OS Release 17.4R1, you can identify IPv6 flows and IPv6-over MPLS flows.
17.2R1	Starting in Junos OS Release 17.2R1, you can identify IPv4-over-MPLS flows.
16.1R1	Starting in Junos OS Release 16.1R1, you can configure the number of flows that can be measured per Packet Forwarding Engine at a time, up to a value of 8192.
15.1R1	Starting in Junos OS Release 15.1R1, MX Series routers support the inline video monitoring to measure media delivery index (MDI) metrics that can be accessed using the SNMP GET operation. Currently, inline MDI can generate only a system log when the computed value is not within the configured range.

RELATED DOCUMENTATION

[Configuring Inline Video Monitoring on MX Series Routers | 934](#)

[show services video-monitoring mdi stats fpc-slot | 1872](#)

[show services video-monitoring mdi errors fpc-slot | 1862](#)

[show services video-monitoring mdi flows fpc-slot | 1865](#)

[alarms | 973](#)

Configuring Inline Video Monitoring on MX Series Routers

IN THIS SECTION

- [Configuring Media Delivery Indexing Criteria | 934](#)
- [Configuring Interface Flow Criteria | 937](#)
- [Configuring the Number of Flows That Can Be Measured | 945](#)

Configuring Media Delivery Indexing Criteria

To configure media delivery indexing criteria:

1. In edit mode, create a named template for video monitoring.

```
user@host# edit services video-monitoring templates template-name
```

For example,

```
user@host# edit services video-monitoring templates t1
```

2. Set the duration for sampling in seconds. Flow media delivery indexing statistics are updated at the end of this interval.

```
[edit services video-monitoring templates template-name]  
user@host# set interval-duration interval-duration
```

For example,

```
[edit services video-monitoring templates t1]  
user@host# set interval-duration 1
```

BEST PRACTICE: If you change the interval duration when a template is being used, you cause a change in the calculated number of expected packets in an measurement interval

for the template. We recommend that you do not change the interval duration for a template that is in use.

3. Set the inactivity timeout.

```
[edit services video-monitoring templates template-name]
user@host# set inactive-timeout inactive-timeout
```

For example,

```
[edit services video-monitoring templates t1]
user@host# set inactive-timeout 30
```

4. Configure either the media rate or layer 3 packet rate to establish expected flow rates used to compare to monitored flow rates.

NOTE: The media rate is the configured media bit rate for the stream. The media rate is used to establish *expected packets per second* (pps).

The Layer 3 packet rate in packets per second (pps) is used to establish *expected bits per second* (bps).

```
[edit services video-monitoring templates template-name]
user@host# set rate media media-bits-per-second
```

For example,

```
[edit services video-monitoring templates t1]
user@host# set rate media 2972400
```

5. Set delay factor thresholds for syslog message levels.

```
[edit services video-monitoring templates template-name]
user@host# set delay-factor threshold info delay-factor-threshold
user@host# set delay-factor threshold warning delay-factor-threshold
user@host# set delay-factor threshold critical delay-factor-threshold
```

For example,

```
[edit services video-monitoring templates t1]
user@host# set delay-factor threshold info 100
user@host# set delay-factor threshold warning 200
user@host# set delay-factor threshold critical 300
```

6. Set media loss rate thresholds for syslog message levels. You can set the threshold based on number of packets lost, or percentage of packets lost.

```
[edit services video-monitoring templates template-name]
user@host# set media-loss-rate threshold info percentage mlr-percentage
user@host# set media-loss-rate threshold warning percentage mlr-percentage
user@host# set media-loss-rate threshold critical percentage mlr-percentage
```

For example,

```
[edit services video-monitoring templates t1]
user@host# set media-loss-rate threshold info percentage 5
user@host# set media-loss-rate threshold warning percentage 10
user@host# set media-loss-rate threshold critical percentage 20
```

7. Set the media rate variation thresholds for syslog message levels. The threshold is based on the ratio of the *difference* between the configured media rate and the monitored media rate to the configured media rate, expressed as a percentage.

```
[edit services video-monitoring templates template-name]
user@host# set media-rate-variation threshold info mrsv-variation
user@host# set media-rate-variation threshold warning mrsv-variation
user@host# set media-rate-variation threshold critical mrsv-variation
```

For example,

```
[edit services video-monitoring templates t1]
user@host# set media-rate-variation threshold info 10
user@host# set media-rate-variation threshold warning 15
user@host# set media-rate-variation threshold critical 20
```

Configuring Interface Flow Criteria

You can identify the input and output flows that you want to monitor. If you do not specify any identifiers, all flows on the interface are monitored. Starting in Junos OS Release 17.2R1, you can identify IPv4-over-MPLS flows. Starting in Junos OS Release 17.4R1, you can identify IPv6 flows and IPv6-over-MPLS flows. MPLS flows with more than three labels cannot be monitored.

NOTE: You can configure a maximum of 256 flow definitions for an interface. If your flow definitions contain lists of addresses and ports, you can exceed the number of match conditions. When you exceed the limits for flows or match conditions, you receive the following constraint message when you commit:

```
'interfaces xe-0/2/2.0'
  Number of flows or Number of match condition under flows exceeded limit
error: configuration check-out failed
```

To configure monitoring of flows for interfaces:

1. In edit mode, identify an interface for monitoring.

```
user@host# edit services video-monitoring interfaces interface-name
```

2. Identify IPv4 input flows for monitoring.

- a. Assign a name to the input flow.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set input-flows input-flow-name
```

- b. Identify the source IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set input-flows input-flow-name source-address [ address ]
```

- c. Identify the destination IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set input-flows input-flow-name destination-address [ address ]
```

- d. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set input-flows input-flow-name source-port [ port ]
```

- e. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set input-flows input-flow-name destination-port [ port ]
```

- f. Identify the template used to monitor the input flow on the interface.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set input-flows input-flow-name template template-name
```

3. Identify IPv4 output flows for monitoring.

- a. Assign a name to the output flow.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set output-flows output-flow-name
```

- b. Identify the source IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set output-flows output-flow-name source-address [ address ]
```

- c. Identify the destination IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set output-flows output-flow-name destination-address [ address ]
```

- d. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set output-flows output-flow-name source-port [ port ]
```

- e. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set output-flows output-flow-name destination-port [ port ]
```

- f. Identify the template used to monitor the output flow on the interface.

```
[edit services video-monitoring interfaces interface-name family inet]
user@host# set output-flows output-flow-name template template-name
```

4. Identify IPv4-over-MPLS input flows for monitoring:

- a. Assign a name to the input flow.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name
```

- b. Identify the payload type as IPv4 over MPLS.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name payload-type ipv4
```

- c. Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use up to 32 destination addresses and up to 32 source addresses.

Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use up to 32 destination addresses and up to 32 source addresses.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name destination-address [address]
user@host# set input-flows input-flow-name source-address [address]
```

If you configure multiple addresses for both the destination and source, then either all the destination or all the source values must have the same prefix length. For example, the following is allowed, because all the destination addresses have the same prefix length.

```
[edit services video-monitoring interfaces ge-0/2/2.0 family mpls]
user@host# set input-flows input-flow-name destination-address [203.0.13.0/24]
```

```
198.51.100.0/24]
```

```
user@host# set input-flows input-flow-name source-address [172.16.0.0/12 192.0.2.11/32]
```

- d. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
```

```
user@host# set input-flows input-flow-name destination-port [ port ]
```

- e. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
```

```
user@host# set input-flows input-flow-name source-port [ port ]
```

- f. Identify the template used to monitor the input flow on the interface.

```
[edit services video-monitoring interfaces interface-name family mpls]
```

```
user@host# set input-flows input-flow-name template template-name
```

5. Identify IPv4-over-MPLS output flows for monitoring:

- a. Assign a name to the output flow.

```
[edit services video-monitoring interfaces interface-name family mpls]
```

```
user@host# set output-flows output-flow-name
```

- b. Identify the payload type as IPv4 over MPLS.

```
[edit services video-monitoring interfaces interface-name family mpls]
```

```
user@host# set output-flows output-flow-name payload-type ipv4
```

- c. Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use up to 32 destination addresses and up to 32 source addresses.

Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use up to 32 destination addresses and up to 32 source addresses.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name destination-address [address]
user@host# set output-flows output-flow-name source-address [address]
```

If you configure multiple addresses for both the destination and source, then either all the destination or all the source values must have the same prefix length. For example, the following is allowed, because all the destination addresses have the same prefix length.

```
[edit services video-monitoring interfaces ge-0/2/2.0 family mpls]
user@host# set output-flows output-flow-name destination-address [203.0.13.0/24
198.51.100.0/24]
user@host# set output-flows output-flow-name source-address [172.16.0.0/12 192.0.2.11/32]
```

- d. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name source-port [port]
```

- e. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name destination-port [ port ]
```

- f. Identify the template used to monitor the output flow on the interface.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name template template-name
```

6. Identify IPv6 input flows for monitoring.

- a. Assign a name to the input flow.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set input-flows input-flow-name
```


- b. Identify the source IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set input-flows input-flow-name source-address [ address ]
```

- c. Identify the destination IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set input-flows input-flow-name destination-address [ address ]
```

- d. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set input-flows input-flow-name source-port [ port ]
```

- e. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set input-flows input-flow-name destination-port [ port ]
```

- f. Identify the template used to monitor the input flow on the interface.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set input-flows input-flow-name template template-name
```

7. Identify IPv6 output flows for monitoring.

- a. Assign a name to the output flow.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set output-flows output-flow-name
```

- b. Identify the source IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set output-flows output-flow-name source-address [ address ]
```

- c. Identify the destination IP address or prefix value for the flow. You can use up to 32 addresses.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set output-flows output-flow-name destination-address [ address ]
```

- d. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set output-flows output-flow-name source-port [ port ]
```

- e. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set output-flows output-flow-name destination-port [ port ]
```

- f. Identify the template used to monitor the output flow on the interface.

```
[edit services video-monitoring interfaces interface-name family inet6]
user@host# set output-flows output-flow-name template template-name
```

8. Identify IPv6-over-MPLS input flows for monitoring:

- a. Assign a name to the input flow.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name
```

- b. Identify the payload type as IPv6 over MPLS.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name payload-type ipv6
```

- c. Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use multiple addresses (up to 32) for either the destination or the source IP address, but not for both.

Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use multiple addresses (up to 32) for either the destination or the source IP address, but not for both.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name destination-address [address]
user@host# set input-flows input-flow-name source-address [address]
```

- d. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name destination-port [ port ]
```

- e. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name source-port [ port ]
```

- f. Identify the template used to monitor the input flow on the interface.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set input-flows input-flow-name template template-name
```

9. Identify IPv6-over-MPLS output flows for monitoring:

- a. Assign a name to the output flow.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name
```

- b. Identify the payload type as IPv6 over MPLS.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name payload-type ipv6
```

- c. Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use multiple addresses (up to 32) for either the destination or the source IP address, but not for both.

Identify the destination IP address or prefix value, the source IP address or prefix value, or both for the flow. You can use multiple addresses (up to 32) for either the destination or the source IP address, but not for both.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name destination-address [address]
user@host# set output-flows output-flow-name source-address [address]
```

- d. Identify the source port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name source-port [port]
```

- e. Identify the destination port for the flow. You can use multiple port numbers and port ranges.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name destination-port [ port ]
```

- f. Identify the template used to monitor the output flow on the interface.

```
[edit services video-monitoring interfaces interface-name family mpls]
user@host# set output-flows output-flow-name template template-name
```

Configuring the Number of Flows That Can Be Measured

Starting in Junos OS Release 16.1R1, you can configure the number of flows that can be measured per Packet Forwarding Engine at a given time by an MPC. This value takes effect the next time the MPC is rebooted. If you do not configure this value, the default maximum value for an MPC is given in ["Understanding Inline Video Monitoring on MX Series Routers" on page 927](#).

To configure the number of flows that can be measured per Packet Forwarding Engine by an MPC at a given time:

- Configure the flow table size. The range is 16 through 8192.

```
[edit chassis fpc slot inline-video-monitoring]
user@host# set flow-table-size size
```

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, you can identify IPv6 flows and IPv6-over MPLS flows.
17.2R1	Starting in Junos OS Release 17.2R1, you can identify IPv4-over-MPLS flows.
16.1R1	Starting in Junos OS Release 16.1R1, you can configure the number of flows that can be measured per Packet Forwarding Engine at a given time by an MPC.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers | 927](#)

[templates | 1452](#)

[interfaces \(Video Monitoring\) | 1181](#)

Inline Video Monitoring Syslog Messages on MX Series Routers

The following examples show the syslog messages produced when configured video monitoring thresholds are exceeded.

/var/log/messages

```
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] DF: 56.71 ms, exceeded threshold for flow(src:192.0.2.2
dst:198.51.100.2 sport:1024 dport:2048) ingressing at interface xe-2/2/1.0 with template t1.
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] MLR : 112, exceeded threshold for flow (src:192.0.2.2
dst:198.51.100.2 sport:1024 dport:2048) ingressing at interface xe-2/2/1.0 with template t1.
Mar 11 18:36:25 tstrtr01 fpc2 [MDI] MRV : -5.67, exceeded threshold for flow (src:192.0.2.2
dst:198.51.100.2 sport:1024 dport:2048) ingressing at interface xe-2/2/1.0 with template t1.
```

Console Messages

```
NPC2(tstrtr01 vty)# [Mar 12 01:40:58.411 LOG: Critical] [MDI] MLR : 420, exceeded threshold for
flow (src:192.0.2.2 dst:198.51.100.2 sport:1024 dport:2048) ingressing at interface xe-2/2/1.0
with template t1.
[Mar 12 01:40:58.411 LOG: Critical] [MDI] MRV : -14.89, exceeded threshold for flow
(src:192.0.2.2 dst:198.51.100.2 sport:1024 dport:2048) ingressing at interface xe-2/2/1.0 with
```

```
template t1.
[Mar 12 01:40:59.412 LOG: Critical] [MDI] DF: 141.74 ms, exceeded threshold for
flow(src:192.0.2.2 dst:198.51.100.2 sport:1024 dport:2048) ingressing at interface xe-2/2/1.0
with template t1.
```

RELATED DOCUMENTATION

[Configuring Inline Video Monitoring on MX Series Routers](#) | 934

Generation of SNMP Traps and Alarms for Inline Video Monitoring on MX Series Routers

IN THIS SECTION

- [Collection of MDI Statistics Associated with an FPC Slot](#) | 948
- [Collection of MDI Errors Associated with an FPC Slot](#) | 949
- [Collection of MDI Flows Associated with an FPC Slot](#) | 950
- [Collection of MDI Record-Level Metrics](#) | 951

Starting in Junos OS Release 15.1, SNMP support is introduced for the Media Delivery Index (MDI) metrics of inline video monitoring. Inline video monitoring is available on MX Series routers using only MPCE1, MPC2, MPC2E, MPC2E-NG, MPC5E, MPC6E, MPC7E, MPC8E, and MPC- 16XGE. Starting in Junos OS Release 20.3R1, inline video monitoring is available on MX Series routers using MPC10E and MPC11E.

Until Junos OS Release 14.2, inline MDI generated only system log messages when the computed MDI metric value was not within the configured range. SNMP support is now added to enable SNMP traps to be triggered when the computed delay factor (DF), media rate variation (MRV), and media loss rate (MLR) value is not within the configured range. You can retrieve the MDI statistics, flow levels, error details, and MDI record-level information using SNMP Get and Get Next requests. The SNMP traps and alarms that are generated when the MDI metrics exceed the configured ranges can be cleared as necessary. Also, you can control the flooding of SNMP traps on the system.

The following sections describe the statistical counters and parameters that are collected for MDI records and for generation of SNMP traps and alarms when the DF, MRV, and MLR values are not within the specified ranges.

Collection of MDI Statistics Associated with an FPC Slot

The FPC-level statistics include the following parameters that are displayed in the output of the `show services video-monitoring mdi stats fpc-slot fpc-slot` command. All of these attributes can be obtained using the SNMP Get request.

Table 133: show services video-monitoring mdi stats fpc-slot Output Fields

Field Name	Field Description
FPC Slot	Slot number of the monitored FPC
Active Flows	Number of active flows currently monitored. active flows = inserted flows - deleted flows.
Total Inserted Flows	Number of flows initiated under video monitoring.
Total Deleted Flows	Number of flows deleted due to inactivity timeout.
Total Packets Count	Number of total packets monitored.
Total Bytes Count	Number of total bytes monitored.
DF Alarm Count	Number of delay factor alarms at each of the following levels: <ul style="list-style-type: none"> • Info level • Warning level • Critical level

Table 133: show services video-monitoring mdi stats fpc-slot Output Fields (Continued)

Field Name	Field Description
MLR Alarm Count	<p>Number of media loss rate (MLR) alarms at each of the following levels:</p> <ul style="list-style-type: none"> • Info level • Warning level • Critical level
MRV alarm count	<p>Number of media rate variation (MRV) alarms at each of the following levels:</p> <ul style="list-style-type: none"> • Info level • Warning level • Critical level

Collection of MDI Errors Associated with an FPC Slot

The FPC-level statistics include the following parameters that are displayed in the output of the `show services video-monitoring mdi errors fpc-slot fpc-slot` command. All of these attributes can be obtained using the SNMP Get request.

Table 134: show services video-monitoring mdi errors fpc-slot Output Fields

Field Name	Field Description
FPC slot	Slot number of the monitored FPC.
Flow Insert Error	Number of errors during new flow insert operations.
Flow Policer Drops	<p>Number of packets dropped by flow policer process.</p> <p>NOTE: New flows usually arrive within a very short time interval (1.5 microseconds). These errors do not represent the loss of entire flows, because subsequent packets in the flow can establish the flow. All packets are monitored after a flow has been established. Packet forwarding occurs independently of the video monitoring, and packets are not dropped due to video monitoring errors.</p>

Table 134: show services video-monitoring mdi errors fpc-slot Output Fields (Continued)

Field Name	Field Description
Unsupported Media Packets Count	Number of packets dropped because they are not media packets or they are unsupported media packets.
PID Limit Exceeded	<p>Number of packets unmonitored because the process identifier (PID) limit exceeded has been exceeded.</p> <p>NOTE: The current PID limit is 6.</p>

Collection of MDI Flows Associated with an FPC Slot

The FPC-level statistics include the following parameters that are displayed in the output of the `show services video-monitoring mdi flows fpc-slot fpc-slot` command. All of these attributes can be obtained using the SNMP Get request.

Table 135: show services mdi flows Output Fields

Field Name	Field Description
SIP	Source IP address
DIP	Destination IP address
SP	Source port
DP	Destination port
Di	Direction (I=Input, O=Output)
Ty	Type of flow
Last DF:MLR	Delay factor and media loss rate value of last media delivery index record

Table 135: show services mdi flows Output Fields (Continued)

Field Name	Field Description
Avg DF:MLR	Average value of delay factor and media loss rate
Last MRV	Media rate variation value of last media delivery index record
Avg MRV	Average value of media rate variation
IFL	Interface name on which flow is receiving
Template Name	Name of template associated with flow

Collection of MDI Record-Level Metrics

The computed DF, MLR, and MRV counters of all valid MDI records of a flow that you can view by using the output of the `show services video-monitoring mdi flow fpc-slot fpc-slot detail` command can be obtained by using the SNMP Get request.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers](#) | 927

SNMP Traps for Inline Video Monitoring Statistics on MX Series Routers

SNMP is primarily used to monitor alarms raised by the inline video monitoring feature. The alarms sent to a network management system (NMS) either to troubleshoot the problem quickly or to proactively diagnose degradation in video quality. The following SNMP traps or alarms are implemented with the Cleared, Info, Warning, and Critical severity levels. The Cleared severity level is used to indicate a normal condition and to clear a particular alarm. Whenever a change in the alarm level occurs, the corresponding alarm is generated.

All the alarms include the following information pertaining to the MDI flows:

- Source IP address

- Destination IP address
- Source Port Destination
- Port Traffic type (UDP or RTP)
- Computed DF, MLR, and MRV values

The following traps are generated for MDI metrics:

- **mdiMLRAAlarm**—This trap is generated when the computed MLR value is not within the configured range.
- **mdiDFAAlarm**—This trap is generated when the computed DF value is not within the configured range.
- **mdiMRVAAlarm**—This trap is generated when the computed MRV value is not within the configured range.

To enable the generation of SNMP traps or alarms for inline video monitoring or MDI metrics, include the `alarms` statement and its substatements at the `[edit services video-monitoring]` hierarchy level.

Processing SNMP GET Requests for MDI Metrics on MX Series Routers

A query on-demand mechanism without caching facility is used to process the SNMP Get requests. The Routing Engine queries the Packet Forwarding Engine to obtain the computed metrics on every Get request. The Routing Engine does not maintain computed metrics locally. No additional memory is required to cache queried metrics. The network management system (NMS) server can receive latest information on every Get request, especially regarding the MDI records because MDI records are updated very frequently. However, querying the Packet Forwarding Engine PFE on each GET request is resource-consuming if the volume of metrics is large. The response to a Get request might be relatively delayed as the Routing Engine has to poll the Packet Forwarding Engine to obtain the metrics.

Inline MDI metrics are real-time data where cached information might not be valid. Reporting cached or invalid metrics is not beneficial because it a real-time monitoring feature. An increase in the number of flows and number of MDI records per flow causes a proportional increase in the volume of memory required in the Routing Engine to store flows and MDI records for all flows. Because asynchronous traps are generated for threshold with enough contents, frequent Get request from NMS are not highly expected, reducing the periodicity of polling to the Packet Forwarding Engine. SNMP traps are triggered with the severity level of Info, Warning, Critical, or Cleared. A trap with the cleared severity level is used to clear an alarm.

Whenever a change in the alarm level occurs, the designated trap is triggered. For example, if the delay factor (DF) alarm changes from informational level to warning level, or from warning to critical, the **mdiDFAAlarm** trap is triggered. Alarm can be immediate or average. If the immediate alarm is configured,

an immediate trap is raised at the end of interval duration if the metric value exceeds the configured range. If the average alarm is configured, a trap is generated, based on the average value for specified number of interval duration.

Storm control is applied for SNMP traps at the flow level and not at the FPC level. The NMS system can obtain SNMP trap from all the flows even if multiple flows are generating traps at approximately the same time. If multiple flows are generating traps at nearly the same time, NMS is flooded by many traps at the same time. For example, no traffic received on a logical interface owing to any reason can trigger all alarms and cause an avalanche of alarms on the NMS server.

RELATED DOCUMENTATION

| [Understanding Inline Video Monitoring on MX Series Routers](#) | 927



Configuration Statements and Operational Commands

[Configuration Statements](#) | 955

[Operational Commands](#) | 1551

Configuration Statements

IN THIS CHAPTER

- [accounting](#) | 965
- [address \(Interfaces\)](#) | 967
- [address \(Services Dynamic Flow Capture\)](#) | 968
- [aggregate-export-interval](#) | 970
- [aggregation](#) | 971
- [alarms](#) | 973
- [alarm-mode](#) | 975
- [allowed-destinations](#) | 977
- [analyzer-address](#) | 978
- [analyzer-id](#) | 979
- [archive-sites](#) | 980
- [authentication-mode](#) | 982
- [authentication-key-chain \(TWAMP\)](#) | 983
- [autonomous-system-type](#) | 985
- [bandwidth-kbps](#) | 987
- [bgp](#) | 988
- [bridge-template](#) | 990
- [capture-group](#) | 991
- [category](#) | 993
- [cflowd \(Discard Accounting\)](#) | 995
- [cflowd \(Flow Monitoring\)](#) | 997
- [client](#) | 998
- [client-delegate-probes](#) | 1002
- [client-list](#) | 1003
- [collector](#) | 1005
- [collector \(Inline Monitoring\)](#) | 1006

- collector (Flow Monitoring Logs for NAT) | 1009
- collector (Flow Template Profiles for NAT) | 1010
- collector-group (Flow Template Profiles for NAT) | 1012
- collector-group (Flow Monitoring Logs for NAT) | 1014
- content-destination | 1016
- control-connection (Junos OS) | 1017
- control-connection (Junos OS Evolved) | 1019
- control-source | 1024
- controller | 1025
- core-dump | 1028
- data-fill | 1029
- data-fill-with zeros | 1031
- data-format | 1032
- data-record-fields | 1033
- data-size | 1037
- delay-factor | 1039
- delegate-probes | 1041
- destination (Interfaces) | 1043
- destination-address (Flow Monitoring Logs for NAT) | 1044
- destination-interface | 1046
- destination-ipv4-address | 1048
- destination-mac-address | 1050
- destination-port | 1051
- destination-port (Flow Monitoring Logs for NAT) | 1054
- destination-udp-port | 1055
- destinations | 1057
- direction | 1058
- disable (Forwarding Options) | 1060
- disable-signature-check | 1062
- dscp (flow-server) | 1063
- dscp-code-points (RPM) | 1065
- dscp-code-points (RFC 2544 Benchmarking) | 1067

- [dump-on-flow-control | 1070](#)
- [duplicates-dropped-periodicity | 1071](#)
- [dynamic-flow-capture | 1072](#)
- [em-hw-profile | 1074](#)
- [engine-id \(Forwarding Options\) | 1076](#)
- [engine-type | 1077](#)
- [exception-reporting | 1079](#)
- [exceptions | 1080](#)
- [export-format | 1082](#)
- [family \(Monitoring\) | 1083](#)
- [family | 1085](#)
- [family \(Sampling\) | 1087](#)
- [features | 1090](#)
- [file \(Sampling\) | 1092](#)
- [file \(Trace Options\) | 1094](#)
- [file-specification \(File Format\) | 1095](#)
- [file-specification \(Interface Mapping\) | 1096](#)
- [filename | 1098](#)
- [filename-prefix | 1099](#)
- [files | 1100](#)
- [filter | 1102](#)
- [flex-flow-sizing | 1103](#)
- [flow-active-timeout | 1105](#)
- [flow-collector | 1107](#)
- [flow-control-options | 1109](#)
- [flow-export-destination | 1111](#)
- [flow-export-rate | 1113](#)
- [flow-export-timer | 1114](#)
- [flow-inactive-timeout | 1116](#)
- [flow-key \(Flow Monitoring\) | 1117](#)
- [flow-monitoring | 1119](#)
- [flow-monitoring \(Inline Monitoring Services\) | 1122](#)

- [flow-server](#) | 1125
- [flow-table-size](#) | 1127
- [flow-table-size \(Chassis\)](#) | 1129
- [flow-tap](#) | 1130
- [forwarding-class \(RFC 2544 Benchmarking\)](#) | 1132
- [forwarding-class \(Sampling\)](#) | 1134
- [ftp \(Flow Collector Files\)](#) | 1135
- [ftp \(Transfer Log Files\)](#) | 1138
- [g-duplicates-dropped-periodicity](#) | 1139
- [g-max-duplicates](#) | 1140
- [generate-snmp-traps](#) | 1142
- [halt-on-prefix-down \(RFC 2544 Benchmarking\)](#) | 1143
- [hard-limit](#) | 1145
- [hard-limit-target](#) | 1146
- [hardware-timestamp](#) | 1147
- [history-size](#) | 1148
- [host-outbound media-interface](#) | 1150
- [icmp](#) | 1151
- [in-service](#) | 1153
- [inactivity-timeout \(Services RPM\)](#) | 1155
- [inet6-options \(Services\)](#) | 1156
- [inband-flow-telemetry](#) | 1157
- [inline-jflow](#) | 1162
- [inline-monitoring](#) | 1163
- [instance](#) | 1165
- [input \(Sampling\)](#) | 1168
- [input-interface-index](#) | 1169
- [input-packet-rate-threshold](#) | 1170
- [instance \(Sampling\)](#) | 1171
- [interface \(Accounting or Sampling\)](#) | 1174
- [interfaces](#) | 1175
- [interface \(Services Flow Tap\)](#) | 1177

- interface-map | 1178
- interfaces (Services Dynamic Flow Capture) | 1179
- interfaces (Video Monitoring) | 1181
- inet6-options (Services) | 1185
- ipfix-sw-mode | 1186
- ip-swap | 1187
- ipv4-flow-table-size | 1189
- ipv4-template | 1191
- ipv6-flow-table-size | 1192
- ipv6-extended-attrib | 1194
- ipv6-template | 1195
- ivlan-cfi (RFC 2544 Benchmarking) | 1196
- ivlan-id (RFC 2544 Benchmarking) | 1198
- ivlan-priority (RFC 2544 Benchmarking) | 1199
- jflow-log (Interfaces) | 1200
- jflow-log (Services) | 1202
- label-position | 1204
- license-server | 1205
- light | 1207
- local-dump | 1209
- logical-system | 1210
- managed | 1211
- match | 1214
- max-connection-duration | 1215
- max-duplicates | 1216
- max-packets-per-second | 1218
- maximum-age | 1219
- maximum-connections | 1221
- maximum-connections-per-client | 1222
- maximum-packet-length | 1224
- maximum-sessions | 1226
- maximum-sessions-per-connection | 1227

- [media-loss-rate](#) | 1229
- [media-rate-variation](#) | 1230
- [message-rate-limit \(Flow Monitoring Logs for NAT\)](#) | 1232
- [minimum-priority](#) | 1233
- [mode](#) | 1235
- [monitoring \(Forwarding Options\)](#) | 1236
- [monitoring \(Services\)](#) | 1238
- [moving-average-size](#) | 1241
- [mpls-flow-table-size](#) | 1243
- [mpls-ipv4-template](#) | 1245
- [mpls-ipvx-template](#) | 1246
- [mpls-template](#) | 1248
- [multiservice-options](#) | 1250
- [name-format](#) | 1251
- [next-hop \(Forwarding Options\)](#) | 1253
- [next-hop \(RPM\)](#) | 1255
- [next-hop-group \(Forwarding Options\)](#) | 1256
- [nexthop-learning](#) | 1258
- [no-remote-trace \(Trace Options\)](#) | 1260
- [no-syslog](#) | 1261
- [no-syslog-generation](#) | 1262
- [notification-targets](#) | 1264
- [observation-domain-id](#) | 1265
- [offload-type](#) | 1267
- [one-way-hardware-timestamp](#) | 1268
- [option-refresh-rate](#) | 1270
- [options-template-id](#) | 1271
- [outer-tag-protocol-id \(RFC 2544 Benchmarking\)](#) | 1273
- [output \(Accounting\)](#) | 1275
- [output \(Monitoring\)](#) | 1276
- [output \(Sampling\)](#) | 1278
- [output-interface-index](#) | 1280

- [ovlan-cfi \(RFC 2544 Benchmarking\) | 1282](#)
- [ovlan-id \(RFC 2544 Benchmarking\) | 1283](#)
- [ovlan-priority \(RFC 2544 Benchmarking\) | 1284](#)
- [owner | 1286](#)
- [packet-loss-priority \(RFC 2544 Benchmarking\) | 1288](#)
- [packet-size \(RFC 2544 Benchmarking\) | 1289](#)
- [passive-monitor-mode | 1291](#)
- [password \(Flow Collector File Servers\) | 1292](#)
- [password \(Transfer Log File Servers\) | 1293](#)
- [peer-as-billing-template | 1294](#)
- [persistent-results | 1296](#)
- [pfe | 1297](#)
- [pic-memory-threshold | 1300](#)
- [pop-all-labels | 1301](#)
- [port \(Flow Monitoring\) | 1303](#)
- [port \(RPM\) | 1304](#)
- [port \(TWAMP\) | 1306](#)
- [post-cli-implicit-firewall | 1307](#)
- [pre-rewrite-tos | 1309](#)
- [primary-data-record-fields | 1310](#)
- [probe | 1313](#)
- [probe-count | 1316](#)
- [probe-interval | 1317](#)
- [probe-limit | 1319](#)
- [probe-server | 1320](#)
- [probe-type | 1323](#)
- [profiles \(RFC 2544 Benchmarking\) | 1324](#)
- [rate \(Interface Services\) | 1326](#)
- [rate \(Forwarding Options\) | 1327](#)
- [receive-failure-threshold \(RFC 2544 Benchmarking\) | 1329](#)
- [receive-options-packets | 1330](#)
- [receive-ttl-exceeded | 1331](#)

- [reflect-etype](#) | 1333
- [reflect-mode](#) | 1334
- [refresh-rate \(Flow Monitoring Logs for NAT\)](#) | 1337
- [required-depth](#) | 1338
- [resiliency](#) | 1340
- [retry \(Services Flow Collector\)](#) | 1343
- [retry-delay](#) | 1344
- [rfc2544](#) | 1345
- [rfc2544-benchmarking](#) | 1347
- [routing-instance \(RPM\)](#) | 1350
- [routing-instance \(cflowd\)](#) | 1352
- [routing-instance-list \(TWAMP\)](#) | 1354
- [routing-instances](#) | 1355
- [rpm \(Interfaces\)](#) | 1357
- [rpm \(Services\)](#) | 1358
- [rpm-scale](#) | 1372
- [rpm-tracking](#) | 1375
- [run-length](#) | 1377
- [sample-once](#) | 1379
- [sampling \(Forwarding Options\)](#) | 1380
- [sampling \(Interfaces\)](#) | 1384
- [sampling-instance](#) | 1385
- [server \(Junos OS\)](#) | 1387
- [server \(Junos OS Evolved\)](#) | 1388
- [server-inactivity-timeout](#) | 1391
- [service-port](#) | 1392
- [service-type](#) | 1393
- [services-options](#) | 1395
- [shared-key](#) | 1397
- [size](#) | 1398
- [skip-arp-iteration \(RFC 2544 Benchmarking\)](#) | 1400
- [slamon-services](#) | 1401

- [soft-limit](#) | [1403](#)
- [soft-limit-clear](#) | [1404](#)
- [source-address \(Forwarding Options\)](#) | [1405](#)
- [source-address \(RPM\)](#) | [1407](#)
- [source-address \(TWAMP\)](#) | [1408](#)
- [source-addresses](#) | [1410](#)
- [source-id](#) | [1412](#)
- [source-ip \(Flow Monitoring Logs for NAT\)](#) | [1413](#)
- [source-ipv4-address \(RFC 2544 Benchmarking\)](#) | [1415](#)
- [source-mac-address](#) | [1416](#)
- [source-udp-port \(RFC 2544 Benchmarking\)](#) | [1418](#)
- [stamp](#) | [1420](#)
- [step-percent \(RFC 2544 Benchmarking\)](#) | [1421](#)
- [store](#) | [1423](#)
- [storm-control](#) | [1426](#)
- [syslog](#) | [1427](#)
- [target-address](#) | [1428](#)
- [tcp](#) | [1431](#)
- [tcp-keepcnt](#) | [1432](#)
- [tcp-keepidle](#) | [1434](#)
- [tcp-keepintvl](#) | [1435](#)
- [template \(Flow Monitoring IPFIX Version\)](#) | [1436](#)
- [template \(Flow Monitoring Version 9\)](#) | [1438](#)
- [template \(Forwarding Options\)](#) | [1440](#)
- [template \(Forwarding Options Version IPFIX\)](#) | [1441](#)
- [template \(Inline Monitoring\)](#) | [1442](#)
- [template-id](#) | [1446](#)
- [template-profile \(Flow Monitoring Logs for NAT\)](#) | [1448](#)
- [template-refresh-rate](#) | [1449](#)
- [template-type \(Flow Monitoring Logs for NAT\)](#) | [1451](#)
- [templates](#) | [1452](#)
- [test](#) | [1456](#)

- tests | **1459**
- test-count | **1462**
- test-finish-wait-duration (RFC 2544 Benchmarking) | **1463**
- test-interface (RFC 2544 Benchmarking) | **1465**
- test-interval | **1467**
- test-iterator-duration (RFC 2544 Benchmarking) | **1469**
- test-iterator-pass-threshold (RFC 2544 Benchmarking) | **1470**
- test-name | **1471**
- test-profile (RFC 2544 Benchmarking) | **1474**
- test-session (Junos OS) | **1476**
- test-session (Junos OS Evolved) | **1477**
- test-type (RFC 2544 Benchmarking) | **1483**
- thresholds (Junos OS) | **1485**
- thresholds (Junos OS Evolved) | **1487**
- timestamp-format (RFC 2544 Benchmarking) | **1489**
- traceoptions (Dynamic Flow Capture) | **1490**
- traceoptions (Forwarding Options) | **1492**
- traceoptions (Inline Monitoring) | **1493**
- traceoptions (Resiliency) | **1496**
- traceoptions (RPM) | **1499**
- transfer | **1502**
- transfer-log-archive | **1503**
- transmit-failure-threshold (RFC 2544 Benchmarking) | **1504**
- traps | **1506**
- ttl | **1509**
- ttl (RPM probe) | **1511**
- tunnel-observation | **1513**
- twamp | **1515**
- twamp-server | **1521**
- trio-flow-offload | **1522**
- udp | **1524**
- udp-tcp-port-swap | **1526**

- [unit](#) | 1527
- [use-extended-flow-memory](#) | 1529
- [username \(Services\)](#) | 1530
- [variant](#) | 1532
- [version](#) | 1533
- [version \(Flow Monitoring Logs for NAT\)](#) | 1534
- [version9 \(Forwarding Options\)](#) | 1536
- [version9 \(Flow Monitoring\)](#) | 1537
- [version-ipfix \(Forwarding Options\)](#) | 1539
- [version-ipfix \(Services\)](#) | 1540
- [video-monitoring](#) | 1542
- [vpls-flow-table-size](#) | 1546
- [vpls-template](#) | 1548
- [world-readable](#) | 1549

accounting

IN THIS SECTION

- [Syntax](#) | 965
- [Hierarchy Level](#) | 966
- [Description](#) | 966
- [Required Privilege Level](#) | 966
- [Release Information](#) | 967

Syntax

```
accounting name {
  output {
    aggregate-export-interval seconds;
```



```

cflowd hostname {
    aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
            caida-compliant;
        }
        source-prefix;
    }
    autonomous-system-type (origin | peer);
    port port-number;
    version format;
}
flow-active-timeout seconds;
flow-inactive-timeout seconds;
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
}
}

```

Hierarchy Level

[edit forwarding-options]

Description

Specify the discard accounting instance name and options.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Discard Accounting](#) | 435

address (Interfaces)

IN THIS SECTION

- [Syntax](#) | 967
- [Hierarchy Level](#) | 967
- [Description](#) | 968
- [Options](#) | 968
- [Required Privilege Level](#) | 968
- [Release Information](#) | 968

Syntax

```
address address {  
    destination address;  
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-numberfamily family]
```

Description

Configure the interface address.

Options

address—Address of the interface.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Junos OS Network Interfaces Library for Routing Devices](#)

[Configuring Flow Monitoring | 5](#)

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

address (Services Dynamic Flow Capture)

IN THIS SECTION

- [Syntax | 969](#)
- [Hierarchy Level | 969](#)
- [Description | 969](#)
- [Options | 969](#)
- [Required Privilege Level | 969](#)

Syntax

```
address address;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Configure an IP address for the flow capture destination.

Options

address—IP address for the content destination.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

aggregate-export-interval

IN THIS SECTION

- [Syntax | 970](#)
- [Hierarchy Level | 970](#)
- [Description | 970](#)
- [Options | 970](#)
- [Required Privilege Level | 970](#)
- [Release Information | 971](#)

Syntax

```
aggregate-export-interval seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting name output],  
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output],  
[edit forwarding-options sampling family (inet |inet6 |mpls) output]
```

Description

Specify the duration, in seconds, of the interval for exporting aggregate accounting information.

Options

seconds—Duration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Discard Accounting](#) | 435

aggregation

IN THIS SECTION

- [Syntax](#) | 971
- [Hierarchy Level](#) | 972
- [Description](#) | 972
- [Options](#) | 972
- [Required Privilege Level](#) | 972
- [Release Information](#) | 972

Syntax

```
aggregation {  
    autonomous-system;  
    destination-prefix;  
    protocol-port;  
    source-destination-prefix {  
        caida-compliant;  
    }  
    source-prefix;  
}
```

Hierarchy Level

```
[edit forwarding-options accounting output cflowd hostname],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server
hostname],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```

Description

For cflowd version 8 only, specify the type of data to be aggregated; cflowd records and sends only those flows that match the specified criteria.

Options

`autonomous-system`—Aggregate by autonomous system (AS) number.

`caida-compliant`—Record source and destination mask-length values in compliance with the Version 2.1b1 release of CAIDA's cflowd application. If this statement is not configured, the Junos OS records source and destination mask length values in compliance with the *cflowd Configuration Guide*, dated August 30, 1999.

`destination-prefix`—Aggregate by destination prefix.

`protocol-port`—Aggregate by protocol and port number.

`source-destination-prefix`—Aggregate by source and destination prefix.

`source-prefix`—Aggregate by source prefix.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Flow Aggregation](#) | 577

alarms

IN THIS SECTION

- [Syntax](#) | 973
- [Hierarchy Level](#) | 974
- [Description](#) | 974
- [Required Privilege Level](#) | 974
- [Release Information](#) | 974

Syntax

```
alarms {  
    delay-factor {  
        no-syslog-generation;  
        generate-snmp-traps;  
        storm-control {  
            count number;  
            interval number;  
        }  
        alarm-mode {  
            mdi-records-count number;  
            average;  
        }  
    }  
    media-rate-variation {  
        no-syslog-generation;  
        generate-snmp-traps;  
        storm-control {  
            count number;  
            interval number;  
        }  
    }  
}
```



```

alarm-mode {
    mdi-records-count number;
    average;
}
}
media-loss-rate {
    no-syslog-generation;
    generate-snmp-traps;
    storm-control {
        count number;
        interval number;
    }
    alarm-mode {
        immediate;
    }
}
}

```

Hierarchy Level

[edit services]

Description

Configure the alarm to monitor and report active alarms. SNMP is used to monitor alarms raised by the inline video monitoring feature. The alarms are monitored in the network management system either to troubleshoot the problem or to diagnose degradation in video quality.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers](#) | 927

[delay-factor](#) | 1039

[no-syslog-generation](#) | 1262

[generate-snmp-traps](#) | 1142

[storm-control](#) | 1426

[alarm-mode](#) | 975

[media-rate-variation](#) | 1230

[media-loss-rate](#) | 1229

alarm-mode

IN THIS SECTION

- [Syntax](#) | 975
- [Hierarchy Level](#) | 976
- [Description](#) | 976
- [Default](#) | 976
- [Options](#) | 976
- [Required Privilege Level](#) | 976
- [Release Information](#) | 976

Syntax

```
alarm-mode {  
    mdi-records-count number;  
    average;  
}
```

Hierarchy Level

[edit services]

Description

If this statement is configured you can set the alarm as immediate or average mode. If immediate alarm is configured, an immediate trap is raised at the end of interval duration when the metric value exceeds the configured range. If average alarm is configured, a trap is generated based on average value for the specified number of interval duration.

Default

The default alarm mode is immediate mode.

Options

mdi-records-count number	Use the specified media delivery index record count number for immediate alarm mode.
average	Generate traps for average values that are not within the configured range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers | 927](#)
alarms | [973](#)

allowed-destinations

IN THIS SECTION

- [Syntax | 977](#)
- [Hierarchy Level | 977](#)
- [Description | 977](#)
- [Options | 977](#)
- [Required Privilege Level | 977](#)
- [Release Information | 978](#)

Syntax

```
allowed-destinations [ destinations ];
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Identify flow capture destinations that are allowed in messages sent from this control source.

Options

destinations—Allowed content destination name.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

analyzer-address

IN THIS SECTION

- [Syntax](#) | 978
- [Hierarchy Level](#) | 978
- [Description](#) | 978
- [Options](#) | 979
- [Required Privilege Level](#) | 979
- [Release Information](#) | 979

Syntax

```
analyzer-address address;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure an IP address for the packet analyzer that overrides the default value.

Options

address—IP address for packet analyzer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

analyzer-id

IN THIS SECTION

- [Syntax](#) | [979](#)
- [Hierarchy Level](#) | [980](#)
- [Description](#) | [980](#)
- [Options](#) | [980](#)
- [Required Privilege Level](#) | [980](#)
- [Release Information](#) | [980](#)

Syntax

```
analyzer-id name;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure an identifier for the packet analyzer that overrides the default value.

Options

name—Identifier for packet analyzer.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

archive-sites

IN THIS SECTION

- [Syntax](#) | 981
- [Hierarchy Level](#) | 981
- [Description](#) | 981
- [Required Privilege Level](#) | 981

Syntax

```
archive-sites {  
  ftp:url {  
    password "password";  
    username username;  
  }  
}
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive]
```

Description

Specify the destination for transfer logs.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

authentication-mode

IN THIS SECTION

- [Syntax | 982](#)
- [Hierarchy Level | 982](#)
- [Description | 982](#)
- [Options | 982](#)
- [Required Privilege Level | 983](#)
- [Release Information | 983](#)

Syntax

```
authentication-mode (authenticated | encrypted | none);
```

Hierarchy Level

```
[edit services rpm twamp server],  
[edit services rpm twamp client control-connection control-client-name]
```

Description

Specify the authentication or encryption mode support for the TWAMP test protocol. This statement is required in the configuration; if no authentication or encryption is specified, you must set the value to `none`.

Options

authenticated Authenticate all TWAMP packets.

NOTE: This mode is supported only on TWAMP servers.

encrypted Encrypt all TWAMP packets.

NOTE: This mode is supported only on TWAMP servers.

none Do not authenticate or encrypt packets.

NOTE: This mode is supported on both TWAMP servers and clients.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches | 694](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

authentication-key-chain (TWAMP)

IN THIS SECTION

● [Syntax | 984](#)

● [Hierarchy Level | 984](#)

- [Description | 984](#)
- [Options | 984](#)
- [Required Privilege Level | 984](#)
- [Release Information | 985](#)

Syntax

```
authentication-key-chain identifier {
    key-id identifier {
        secret password-string;
    }
}
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Apply and enable an authentication key chain to the routing device. Note that the referenced key chain must be defined. When configuring the authentication key update mechanism for TWAMP, you cannot commit the `0.0.0.0/allow` statement with authentication keys or key chains. The CLI issues a warning and fails to commit such configurations.

Options

identifier—Authentication key chain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").

password-string—Authentication key, consisting of 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

autonomous-system-type

IN THIS SECTION

- [Syntax](#) | 985
- [Hierarchy Level](#) | 986
- [Description](#) | 986
- [Default](#) | 986
- [Options](#) | 986
- [Required Privilege Level](#) | 986
- [Release Information](#) | 986

Syntax

```
autonomous-system-type (origin | peer);
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server hostname],  
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```

Description

Specify the type of AS numbers that cflowd exports.

Default

origin

Options

origin—Export origin AS numbers of the packet source address in the Source Autonomous System cflowd field.

peer—Export peer AS numbers through which the packet passed in the Source Autonomous System cflowd field.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Flow Aggregation](#) | 577

bandwidth-kbps

IN THIS SECTION

- [Syntax | 987](#)
- [Hierarchy Level | 987](#)
- [Description | 987](#)
- [Options | 987](#)
- [Required Privilege Level | 988](#)
- [Release Information | 988](#)

Syntax

```
bandwidth-kbps kbps;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profiletest-profile profile-name]
```

Description

Define the theoretical maximum bandwidth, in kilobits per second, for the test. The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested. The range is 1,000 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value used for this test.

Options

- kbps* Bandwidth limit, in kilobits per second (kpbs).
- **Range:** 1,000 kbps through 1,000,000 Kbps.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

bgp

IN THIS SECTION

- [Syntax | 988](#)
- [Hierarchy Level | 989](#)
- [Description | 989](#)
- [Options | 989](#)
- [Required Privilege Level | 989](#)
- [Release Information | 990](#)

Syntax

```
bgp {
  data-fill data;
  data-size size;
  destination-port port;
```

```

history-size size;
logical-system logical-system-name <routing-instances routing-instance-name>;
moving-average-size size;
probe-count count;
probe-interval seconds;
probe-type type;
routing-instances instance-name;
rfc6514-compliant-safi129;
test-interval interval;
}

```

Hierarchy Level

```

[edit services rpm bgp],
[edit protocols bgp group group-name],
[edit routing-instances instance-name protocols bgp group group-name],
[edit logical-system logical-system-name protocols bgp group group-name],
[edit logical-system logical-system-name routing-instances instance-name protocols bgp group group-name]

```

Description

Configure BGP neighbor discovery through Real-Time Performance Monitoring (RPM).

Options

bgp—Define properties for configuring BGP neighbor discovery.

The remaining statements are explained separately. See [CLI Explorer](#).

NOTE: On MX Series routers, you can configure all the statements. On M Series and T Series routers, you can configure only the logical-system and routing-instances statements.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring BGP Neighbor Discovery Through RPM](#) | 671

bridge-template

IN THIS SECTION

- [Syntax](#) | 990
- [Hierarchy Level](#) | 990
- [Description](#) | 990
- [Required Privilege Level](#) | 991
- [Release Information](#) | 991

Syntax

```
bridge-template;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix |version9 template template-name]
```

Description

Specify that the template is used for bridge records or for VPLS records.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

capture-group

IN THIS SECTION

- [Syntax | 991](#)
- [Hierarchy Level | 992](#)
- [Description | 992](#)
- [Required Privilege Level | 992](#)
- [Release Information | 992](#)

Syntax

```
capture-group client-name {
  content-destination identifier {
    address address;
    hard-limit bandwidth;
    hard-limit-target bandwidth;
    soft-limit bandwidth;
    soft-limit-clear bandwidth;
    ttl hops;
```

```

}
control-source identifier {
    allowed-destinations [ destinations ];
    minimum-priority value;
    no-syslog;
    notification-targets address port port-number;
    service-port port-number;
    shared-key value;
    source-addresses [ addresses ];
}
duplicates-dropped-periodicity seconds;
input-packet-rate-threshold rate;
interfaces interface-name;
max-duplicates number;
pic-memory-threshold percentage percentage;
}

```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Define the capture group values.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Junos Capture Vision](#) | 289

category

IN THIS SECTION

- [Syntax](#) | 993
- [Hierarchy Level](#) | 993
- [Description](#) | 993
- [Options](#) | 994
- [Required Privilege Level](#) | 994
- [Release Information](#) | 994

Syntax

```
category category-name {  
    inline-monitoring-instance inline-monitoring-instance;  
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pfe identifier exception-reporting]
```

Description

(Required) Configure the PFE exception category type and assign a particular inline-monitoring instance to it.

Options

category-name Name of the category. Include more than one *category* statement in your configuration if you want reporting for more than one exception type (Junos OS only). You can specify one of the following types on each *category* statement:

- *all*—(Junos OS Evolved only) All exceptions
- *firewall*—(Junos OS only) Firewall exceptions
- *forwarding-state*—(Junos OS only) Forwarding-state-related exceptions
- *layer2*—(Junos OS only) Layer 2 exceptions
- *layer3*—(Junos OS only) Layer 3 exceptions
- *packet-errors*—(Junos OS only) Packet-format-error-related exceptions

inline-monitoring-instance inline-monitoring-instance

Assign a particular inline-monitoring instance to the exception reporting category.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

[Juniper Resiliency Interface](#) | 409

cflowd (Discard Accounting)

IN THIS SECTION

- [Syntax | 995](#)
- [Hierarchy Level | 996](#)
- [Description | 996](#)
- [Options | 996](#)
- [Required Privilege Level | 996](#)
- [Release Information | 996](#)

Syntax

```
cflowd hostname {  
    aggregation {  
        autonomous-system;  
        destination-prefix;  
        protocol-port;  
        source-destination-prefix {  
            caida-compliant;  
        }  
        source-prefix;  
    }  
    autonomous-system-type (origin | peer);  
    label-position {  
        template template-name;  
    }  
    (local-dump | no-local-dump);  
    port port-number;  
    source-address address;  
    version format;  
}
```

Hierarchy Level

```
[edit forwarding-options accounting name output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility `cfddcollect`.

You can configure up to one version 5 and one version 8 flow format at the `[edit forwarding-options accounting name output]` hierarchy level.

Options

hostname—IP address or identifier of the host system (the workstation running the `cflowd` utility).

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Flow Aggregation](#) | 577

cflowd (Flow Monitoring)

IN THIS SECTION

- [Syntax | 997](#)
- [Hierarchy Level | 997](#)
- [Description | 997](#)
- [Options | 997](#)
- [Required Privilege Level | 998](#)
- [Release Information | 998](#)

Syntax

```
cflowd hostname {
    port port-number;
}
```

Hierarchy Level

```
[edit forwarding-options monitoring (Forwarding Options) name inet output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility `cfcollect`.

You can configure up to eight version 5 flow formats at the `[edit forwarding-options monitoring (Forwarding Options) name output]` hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.

Options

hostname—IP address or identifier of the host system (the workstation running the `cflowd` utility).

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Flow Aggregation](#) | [577](#)

client

IN THIS SECTION

- [Syntax \(Junos OS\) | 998](#)
- [Syntax \(Junos OS Evolved\) | 1000](#)
- [Junos OS Hierarchy Level | 1001](#)
- [Junos OS Evolved Hierarchy Level | 1001](#)
- [Description | 1001](#)
- [Required Privilege Level | 1001](#)
- [Release Information | 1001](#)

Syntax (Junos OS)

```
client {
  control-connection control-client- name {
    authentication-mode;
    control-type (managed | light);
```

```

destination-interface interface-name;
destination-port port;
history-size size;
moving-average-size number;
persistent-results;
routing-instance instance-name;
target (url url | address address);
tcp-keepcnt;
tcp-keepidle;
tcp-keepintvl;
test-interval interval;
traps {
    control-connection-closed;
    test-iteration-done;
}
data-fill-with zeros;
data-size size;
dscp-code-points (RPM) dscp-bits;
probe-count count;
probe-interval seconds;
thresholds thresholds;
test-session session-name {
    data-fill-with zeros data;
    data-size size;
    dscp-code-points (RPM) dscp-bits;
    probe-count count;
    probe-interval seconds;
    source-address source-address;
    target-address target-address local-link IPv6-link-local-interface-name;
    traps {
        egress-jitter-exceeded;
        egress-std-dev-exceeded;
        egress-time-exceeded;
        ingress-jitter-exceeded;
        ingress-std-dev-exceeded;
        ingress-time-exceeded;
        jitter-exceeded;
        max-rtt-exceeded;
        probe-failure;
        rtt-exceeded;
        std-dev-exceeded;
        test-completion;
        test-failure;
    }
}

```

```

    }
  }
}

```

Syntax (Junos OS Evolved)

```

client {
  control-connection control-client-name {
    control-type (managed | light);
    destination-port destination-port;
    routing-instance routing-instance-name;
    source-address source-address;
    target target-address;
    test-start (auto | manual);
    test-interval seconds;
    traps {
      control-connection-closed;
      test-iteration-done;
    }
    test-session name {
      data-size data-size;
      destination-port destination-port;
      dscp-code-points dscp-code-points;
      history-size history-size;
      moving-average-size moving-average-size;
      offload-type (none | inline-timestamping | pfe-timestamp);
      probe-count probe-count;
      probe-interval seconds;
      source-address source-address;
      target target-address local-link IPv6-link-local-interface-name;
      thresholds {
        control-failure (on | off);
        successive-loss number;
        total-loss number;
        threshold-type (microseconds | average);
      }
      traps {
        egress-jitter-exceeded;
        egress-time-exceeded;
        ingress-jitter-exceeded;

```

```

        ingress-time-exceeded;
        jitter-exceeded;
        probe-failure;
        rtt-exceeded;
        test-completion;
        test-failure;
    }
    ttl hop-count;
}
}
}

```

Junos OS Hierarchy Level

```
[edit services rpm twamp]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp]
```

Description

Specify the TWAMP client configuration settings.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Statement introduced in Junos OS Evolved Release 20.3R1.

control-type option introduced in Junos OS Release 21.1R1 and Junos OS Evolved Release 20.3R1.

traps option introduced in Junos OS Evolved Release 21.3R1.

source-address option and the local-link sub-option of the target-address option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

The local-link sub-option of the target option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The inline-timestamping suboption of the offload-type option introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

[Understand Two-Way Active Measurement Protocol](#) | 686

client-delegate-probes

IN THIS SECTION

- [Syntax](#) | 1002
- [Hierarchy Level](#) | 1002
- [Description](#) | 1003
- [Required Privilege Level](#) | 1003
- [Release Information](#) | 1003

Syntax

```
rpm client-delegate-probes;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Generate real-time performance monitoring (RPM) probes on an MS-MPC or MS-MIC services interface, which increases the number of RPM probes that can run at the same time.

The destination-interface statement must be configured at the [edit services rpm probe *owner* test *test-name*] hierarchy level to point to the interface and logical unit number and for which you configure client-delegate-probes. Configure the delegate-probes statement at the [edit services rpm probe *owner*] hierarchy level to complete the configuration.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

| [Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | 651

client-list

IN THIS SECTION

- [Syntax](#) | 1004
- [Junos OS Hierarchy Level](#) | 1004
- [Junos OS Evolved Hierarchy Level](#) | 1004
- [Description](#) | 1004
- [Options](#) | 1004
- [Required Privilege Level](#) | 1004
- [Release Information](#) | 1004

Syntax

```
client-list list-name {  
    address address;  
}
```

Junos OS Hierarchy Level

```
[edit services rpm twamp server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

Specify the list of allowed control client hosts that can connect to this server. Each entry is a Classless Interdomain Routing (CIDR) address (IP address plus mask) that represents a network of allowed hosts. You can configure more than one list, but you must configure at least one client address to enable TWAMP. Each list can contain up to 64 entries.

Options

list-name—Name of client address list.

address—Address and mask for an allowed client.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Evolved 20.3R1.

IPv6 address support introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

collector

IN THIS SECTION

- [Syntax](#) | 1005
- [Hierarchy Level](#) | 1005
- [Description](#) | 1005
- [Options](#) | 1006
- [Required Privilege Level](#) | 1006
- [Release Information](#) | 1006

Syntax

```
collector interface-name;
```

Hierarchy Level

```
[edit services flow-collector interface-map]
```

Description

Configure the default flow collector interface for interface mapping.

Options

interface-name—Default flow collector interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

collector (Inline Monitoring)

IN THIS SECTION

- [Syntax \(Junos OS\) | 1007](#)
- [Syntax \(Junos OS Evolved\) | 1007](#)
- [Hierarchy Level | 1007](#)
- [Description | 1007](#)
- [Options | 1007](#)
- [Required Privilege Level | 1008](#)
- [Release Information | 1008](#)

Syntax (Junos OS)

```
collector collector-name {
    destination-address destination-IP-address;
    destination-port destination-port-number;
    dscp dscp;
    forwarding-class forwarding-class;
    routing-instance routing-instance-name;
    sampling-rate sampling-rate;
    source-address source-IP-address;
}
```

Syntax (Junos OS Evolved)

```
collector collector-name {
    destination-address destination-IP-address;
    destination-port destination-port;
    source-address source-IP-address;
}
```

Hierarchy Level

```
[edit services inline-monitoring instance instance-name]
```

Description

Configure an collector for inline monitoring. The monitored packets are exported to the collector in an IPFIX format. The actual packet is exported in an IPFIX format up to the configured clip length. By default, Junos OS supports a maximum packet length of 126 bytes, starting with the Ethernet header. The IPFIX format exports information on the original packet size, and the incoming or outgoing interface for further processing on the collector.

Options

<i>collector-name</i>	Name of the collector.
<i>destination-address</i> <i>destination-IP-address</i>	IPv4 destination IP address.

destination-port <i>destination-port-number</i>	Destination port value. <ul style="list-style-type: none">• Range: 1 through 65535
dscp <i>dscp</i>	(Junos OS only) DSCP value. <ul style="list-style-type: none">• Default: 0• Range: 0 through 63
forwarding-class <i>forwarding-class</i>	(Junos OS only) Forwarding class for exported frames. <ul style="list-style-type: none">• Default: best-effort
routing-instance <i>routing-instance-name</i>	(Junos OS only) Name of the routing instance. <ul style="list-style-type: none">• Default: default.inet
sampling-rate <i>sampling-rate</i>	(Junos OS only) Rate at which the packets are sampled. In Junos OS Evolved, you do not specify the sampling rate here; you specify the sampling rate at the [edit services inline-monitoring instance <i>instance-name</i>] hierarchy level. <p>For example, if you specify 1000, then 1 packet out of every 1000 packets is sampled.</p> <ul style="list-style-type: none">• Range: 1 through 16000000• Default: 1
source-address <i>source-IP-address</i>	IPv4 source IP address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 19.4R1.

Statement introduced in Junos Evolved OS Release 22.2R1.

RELATED DOCUMENTATION

[Understanding Inline Monitoring Services](#) | 334

collector (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax](#) | 1009
- [Hierarchy Level](#) | 1009
- [Description](#) | 1009
- [Options](#) | 1010
- [Required Privilege Level](#) | 1010
- [Release Information](#) | 1010

Syntax

```
collector collector-name {  
    source-ip address;  
    destination-address address;  
    destination-port port-number;  
}
```

Hierarchy Level

```
[edit services jflow-log]
```

Description

Specify the name of the collector to which flow monitoring log messages in IPFIX or version 9 flow template format for NAT events must be sent. The generated flow monitoring logs for NAT events in flow template format are sent to the specified host or external device that functions as the NetFlow

collector. You must associate a collector with a template profile for the template characteristics, such as refresh rate of messages and the template format, to be used for generated flow monitoring logs.

Options

collector-name—Name of the collector to which flow monitoring log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

The remaining statements are described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

collector (Flow Template Profiles for NAT)

IN THIS SECTION

 [Syntax | 1011](#)

- Hierarchy Level | 1011
- Description | 1011
- Options | 1011
- Required Privilege Level | 1011
- Release Information | 1012

Syntax

```
collector collector-name;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the name of the collector to be associated with a template profile. The generated flow monitoring logs for NAT events in flow template format are sent to the specified collector. You must have previously configured the collector by using the `collector collector-name` statement at the `[edit services jflow-log]` hierarchy level before you associate a collector with a template profile.

Options

collector-name—Name of the collector to which flow monitoring log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are `[a-zA-Z0-9_]`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

collector-group (Flow Template Profiles for NAT)

IN THIS SECTION

- [Syntax | 1012](#)
- [Hierarchy Level | 1013](#)
- [Description | 1013](#)
- [Options | 1013](#)
- [Required Privilege Level | 1013](#)
- [Release Information | 1013](#)

Syntax

```
collector-group collector-group-name;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the name of the collector group to be associated with a template profile. The generated flow monitoring logs for NAT events in flow template format are sent to the specified collector group. By using a collector group, you can effectively and optimally transmit flow monitoring logs to a cluster of collectors in a single, one-step operation. A maximum of up to eight collectors can be aggregated into a collector group. You must have previously configured the collector group by using the `collector-group collector-group-name` statement at the `[edit services jflow-log]` hierarchy level before you associate a collector-group with a template profile.

Options

collector-group-name—Name of the collector group to which log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are `[a-zA-Z0-9_]`

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275

collector-group (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax | 1014](#)
- [Hierarchy Level | 1014](#)
- [Description | 1014](#)
- [Options | 1015](#)
- [Required Privilege Level | 1015](#)
- [Release Information | 1015](#)

Syntax

```
collector-group collector-group-name {  
    [collector-name1 collector-name2];  
}
```

Hierarchy Level

```
[edit services jflow-log]
```

Description

Specify the name of the collector group that contains a set of NetFlow collectors to which flow monitoring log messages in IPFIX or version 9 flow template format for NAT events must be sent. You must define at least one collector in the group. A maximum of up to eight collectors can be aggregated into a collector group.

The generated flow monitoring logs for NAT events in flow template format are sent to the specified collector group. By using a collector group, you can effectively and optimally transmit flow monitoring logs to a cluster of collectors in a single, one-step operation.

Options

collector-group-name—Name of the collector group to which flow monitoring log messages for NAT events in flow monitoring format (IPFIX or version 9 flow template format) must be sent. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

collector-name—Name of the collector to be assigned to the group of collectors. You must have previously defined the collector by including the collector *collector-name* statement at the [edit services jflow-log] hierarchy level. You can specify a list of valid collector names. Specify the names individually by using a space to separate each collector name. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_]

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

content-destination

IN THIS SECTION

- [Syntax | 1016](#)
- [Hierarchy Level | 1016](#)
- [Description | 1016](#)
- [Options | 1016](#)
- [Required Privilege Level | 1017](#)
- [Release Information | 1017](#)

Syntax

```
content-destination identifier {  
    address address;  
    hard-limit bandwidth;  
    hard-limit-target bandwidth;  
    soft-limit bandwidth;  
    soft-limit-clear bandwidth;  
    ttl hops;  
}
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Identify the destination for captured packets.

Options

identifier—Name of the destination.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

control-connection (Junos OS)

IN THIS SECTION

- [Syntax](#) | 1017
- [Hierarchy Level](#) | 1018
- [Description](#) | 1018
- [Options](#) | 1019
- [Required Privilege Level](#) | 1019
- [Release Information](#) | 1019

Syntax

```
control-connection control-client-name {  
    authentication-mode  
    control-type (managed | light);  
    destination-interface interface-name;  
    destination-port port;
```

```

history-size size;
moving-average-size number;
persistent-results
routing-instance instance-name;
target-address (url url | address);
tcp-keepcnt
tcp-keepidle
tcp-keepintvl
test-interval interval;
traps traps;
data-fill-with zeros data;
data-size size;
dscp-code-points dscp-bits;
probe-count count;
probe-interval seconds;
thresholds thresholds;
test-session session-name{
    data-fill-with zeros data;
    data-size size;
    dscp-code-points dscp-bits;
    destination-port destination-port;
    probe-count count;
    probe-interval seconds;
    source-address source-address;
    target-address (url url | address) local-link IPv6-link-local-interface-name;
}
}

```

Hierarchy Level

```
[edit services rpm twamp client]
```

Description

List all the TWAMP control clients that can connect to this server. You must configure at least one client to enable TWAMP.

Options

control-client-name Name of the control 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

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

Statement introduced in Junos OS Evolved Release 20.3R1.

control-type option introduced in Junos OS Release 21.1R1 and Junos OS Evolved Release 20.3R1.

source-address option and the local-link sub-option of the target-address option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

[Understand Two-Way Active Measurement Protocol | 686](#)

[control-connection \(Junos OS Evolved\) | 1019](#)

control-connection (Junos OS Evolved)

IN THIS SECTION

- [Syntax | 1020](#)
- [Hierarchy Level | 1021](#)
- [Description | 1021](#)
- [Options | 1021](#)

- Required Privilege Level | 1023
- Release Information | 1023

Syntax

```
control-connection control-client-name {
    control-type (managed | light);
    destination-port destination-port;
    routing-instance routing-instance-name;
    source-address source-address;
    target target-address;
    test-start (auto | manual);
    test-interval seconds;
    traps {
        control-connection-closed;
        test-iteration-done;
    }
    test-session name {
        data-size data-size;
        destination-port destination-port;
        dscp-code-points dscp-code-points;
        history-size history-size;
        moving-average-size moving-average-size;
        offload-type (none | inline-timestamping | pfe-timestamp);
        probe-count probe-count;
        probe-interval seconds;
        source-address source-address;
        target target-address local-link IPv6-link-local-interface-name;
        thresholds {
            control-failure (on | off);
            successive-loss number;
            total-loss number;
            threshold-type (microseconds | average);
        }
        traps {
            egress-jitter-exceeded
            egress-time-exceeded;
            ingress-jitter-exceeded;
```

```

        ingress-time-exceeded;
        jitter-exceeded;
        probe-failure;
        rtt-exceeded;
        test-completion;
        test-failure;
    }
    ttl hop-count;
}
}

```

Hierarchy Level

```
[edit services monitoring twamp client]
```

Description

For Junos OS Evolved, configure all the Two-Way Active Measurement Protocol (TWAMP) control clients that can connect to the server. A TWAMP control connection is responsible for initiating, starting, and ending the test sessions between a TWAMP client and a TWAMP server for performance measurement. However, when the control-type option is set to `light`, the test session parameters are predefined and not negotiated, and no control session is created. Therefore, with TWAMP `light`, UDP probe generation and reception, as well as reflection, is not part of a control session.

Options

control-client-name Name of the control client.

control-type (managed | light) Specify how you want to manage the control connection.

- **Values:** Configure one of the following:
 - **managed**—Specify that you want to have a stateful version of TWAMP where test parameters are negotiated. You can specify particular source or target addresses, or non-default destination ports for the control connections. The ports are negotiated and agreed upon between the client and server. All test packets received by the server on a destination port are reflected back and forgotten right away.

- **light**—Specify that you want a stateless version of TWAMP where test parameters are predefined instead of negotiated. All test packets received by the server on a destination port are reflected back and forgotten right away.
- **Default:** managed

destination-port *port*

Specify the Transmission Control Protocol (TCP) port number for the endpoint of the control connection. You must set the control-type option to `managed` to set the destination port for the entire control connection. Otherwise, if the control-type option is set to `light`, you will have to configure the destination-port option on each test session.

- **Range:** You can specify port 862, or any port from 49152 through 65535.
- **Default:** 862 (IANA port for TWAMP)

routing-instance *instance-name*

Specify the routing instance used by the probes. This routing instance is configured at the [edit routing-instance] hierarchy level.

NOTE: The media interface from where the TWAMP control and test or data packets arrive and exit must be part of the same routing instance.

- **Default:** No routing instance is used. The Internet (IPv4) routing table `inet.0` is used instead.

source-address *address*

Specify the IPv4 or IPv6 source address used for control connections. If the source address is not one of the device's assigned addresses, the control connection uses the outgoing interface's address as its source.

You cannot use IPv6 link-local addresses or the following addresses as the source address for control connections:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

target *target-address*

Specify the IPv4 or IPv6 address for the target destination for the control connection. You cannot use IPv6 link-local addresses as the target destination for the control connection.

NOTE: You can configure this option only when the control-type option is set or defaulted to managed. The target option is required for a managed control connection.

**test-start (auto
| manual)**

Specify whether or not you want to manually start TWAMP tests.

- **Values:** Configure one of the following:
 - auto—Start test sessions as soon as you commit the configuration.
 - manual—Wait to start test sessions until you issue the request services monitoring twamp client start operational mode command.
- **Default:** auto

**test-interval
seconds**

Specify the time to wait between tests.

- **Range:** 1 through 86400 seconds
- **Default:** 1

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.3R1.

traps option introduced in Junos OS Evolved Release 21.3R1.

IPv6 address support introduced in Junos OS Evolved Release 21.4R1.

The local-link sub-option of the target option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The inline-timestamping suboption of the offload-type option introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

[Understand Two-Way Active Measurement Protocol](#) | 686

[control-connection \(Junos OS\)](#) | 1017

[request services monitoring twamp client](#) | 1577

control-source

IN THIS SECTION

- [Syntax](#) | 1024
- [Hierarchy Level](#) | 1024
- [Description](#) | 1025
- [Options](#) | 1025
- [Required Privilege Level](#) | 1025
- [Release Information](#) | 1025

Syntax

```
control-source identifier {  
    allowed-destinations [ destinations ];  
    minimum-priority value;  
    no-syslog;  
    notification-targets address port port-number;  
    service-port port-number;  
    shared-key value;  
    source-addresses [ addresses ];  
}
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Identify the source of the dynamic flow capture request.

Options

identifier—Name of control source.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

controller

IN THIS SECTION

- [Syntax](#) | 1026
- [Hierarchy Level](#) | 1026
- [Description](#) | 1026
- [Options](#) | 1027
- [Required Privilege Level](#) | 1028
- [Release Information](#) | 1028

Syntax

```
controller(p4 | re);
```

Hierarchy Level

```
[edit services inline-monitoring instance]
```

Description

PTX10004, PTX10008, and PTX10016 devices running Junos OS Evolved Release 22.1R1 can provide an SDN-based backbone (data center interconnect) for target networks. It works by using a firewall filter to redirect matching packets, via a P4Runtime agent running in Junos OS Evolved on the Routing Engine, to a P4 controller that is also running on the PTX router. (The P4Runtime agent registers with the Juniper Extension Toolkit (JET) services daemon (JSD) to open the gRPC connections and listen for P4 requests from clients.)

You can match:

- IPv4, IPv6, UDP, and TCP protocol packets according to the destination IP address
- Google Discovery protocol packets (matched by specifying VLAN ID: 4000, EtherType: 0x6007)
- traceroute redirect packets (matched by specifying TTL=0 for IPv4 and TTL=1 for IPv6)

The P4Runtime agent supports packet I/O from network devices to the SDN controller, as well as [OpenConfig](#) for switch configuration, and gRIBI for route programming.

To configure inline-monitoring services for packet redirects to the P4 controller, you need to configure inline-monitoring, create an instance, and set the instance type to controller P4, as shown here:

```
{master}
[edit services inline-monitoring]
instance {
  Instance-1 {
    controller p4;
```

```

    }
}

```

And configure a firewall filter **action** to redirect matching packets to your *instance*, as shown here:

```

{master}
[edit firewall family (any | inet | inet6) filter f1]
term t1 {
    then redirect Instance-1;
}

```

Optionally, you may also want to modify the DDoS protocol parameters:

```

{master}
[edit system ddos-protection protocols custom]
aggregate {
    bandwidth value;
    burst value;
}

```

or disable DDoS, as the case may be:

```

{master}
[edit system ddos-protection protocols custom]
aggregate {
    disable-fpc;
}

```

Controller type:

Options

- p4** Select this option to have the P4Runtime agent send packets to the P4 controller (instead of the default Routing Engine).
- re** Select this option to send packets to the default Routing Engine.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 22.1R1.

RELATED DOCUMENTATION

[Firewall Filter Match Conditions and Actions \(PTX Series Routers\)](#)

core-dump

IN THIS SECTION

- [Syntax | 1028](#)
- [Hierarchy Level | 1028](#)
- [Description | 1029](#)
- [Required Privilege Level | 1029](#)
- [Release Information | 1029](#)

Syntax

```
(core-dump | no-core-dump);
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port multiservice-options]
```

Description

A useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory `/var/tmp` contains core files. Junos OS saves the current core file (0) and the four previous core files, which are numbered from 1 through 4 (from newest to oldest):

NOTE: By default, all members of a configured user group (with read-only permissions) can access the core dump files and attach them to cases associated with JTAC.

- `core-dump`—Enable the core dumping operation.
- `no-core-dump`—Disable the core dumping operation.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Flow Monitoring | 5](#)

data-fill

IN THIS SECTION

- [Syntax | 1030](#)
- [Junos OS Hierarchy Levels | 1030](#)
- [Junos OS Evolved Hierarchy Level | 1030](#)
- [Description | 1030](#)

- Options | 1030
- Required Privilege Level | 1030
- Release Information | 1031

Syntax

```
data-fill data;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],  
[edit servicesrpm probe owner test test-name],
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the contents of the data portion of Internet Control Message Protocol (ICMP) probes. The data-fill statement is not valid with the http-get or http-metadata-get probe types.

Options

data—A hexadecimal value; for example, 0-9, A-F.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

data-fill-with zeros

IN THIS SECTION

- [Syntax | 1031](#)
- [Hierarchy Level | 1031](#)
- [Description | 1032](#)
- [Required Privilege Level | 1032](#)
- [Release Information | 1032](#)

Syntax

```
data-fill-with-zeros;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Description

If this statement is configured, then the contents of the test packet are zeros, if the statement is not configured, then the data content is a pseudo-random number.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

| [Understand Two-Way Active Measurement Protocol](#) | 686

data-format

IN THIS SECTION

- [Syntax](#) | 1032
- [Hierarchy Level](#) | 1033
- [Description](#) | 1033
- [Options](#) | 1033
- [Required Privilege Level](#) | 1033
- [Release Information](#) | 1033

Syntax

```
data-format format;
```

Hierarchy Level

```
[edit services flow-collector file-specification variant variant-number]
```

Description

Specify the data format for a specific file format variant.

Options

format—Data format. Specify **flow-compressed** as the data format.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

data-record-fields

IN THIS SECTION

- [Syntax](#) | 1034
- [Hierarchy Level](#) | 1034
- [Description](#) | 1034
- [Default](#) | 1035

- [Options | 1035](#)
- [Additional Information | 1036](#)
- [Required Privilege Level | 1037](#)
- [Release Information | 1037](#)

Syntax

```
data-record-fields {  
    source-prefix-as-path count;  
    destination-prefix-as-path count;  
    bgp-source-standard-community count;  
    bgp-destination-standard-community count;  
    bgp-source-extended-community count;  
    bgp-destination-extended-community count;  
    bgp-source-large-community count;  
    bgp-destination-large-community count;  
}
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Export the BGP community and AS path information in the data record using IPFIX, per RFC 8549. The IPFIX collector should be RFC 6313 compliant to decode data records carrying fields with the basicList format. Any new or changed configuration of this statement is effective immediately and the exported packets contain the new information.

count defines the size of the list that will be exported for that Information Element in the data record. If the total number of items exceeds the configured list size, the exported list contains the first and subsequent items up to and including the configured list size. For example, if the AS path list for the source prefix contains 8 AS numbers, but *count* is set to 4, only the first four AS numbers are exported in the data record.

Each Information Element is independent of the others and can be exported separately in the collector record.

Default

When this statement is not configured, the routing protocol process (rpd) only sends the peer AS number, origin AS number, and the BGP next-hop information. When this statement is configured with at least one Information Element, rpd sends all of the community and AS-path-related information. You do not have to configure all of the Information Elements.

Options

source-prefix-as-path count	Export the source prefix AS path (Information Element 16, per RFC 6313). <ul style="list-style-type: none">• Default: 8
destination-prefix-as-path count	Export the destination prefix AS path (Information Element 17, per RFC 6313). <ul style="list-style-type: none">• Default: 8
bgp-source-standard-community count	Export the standard BGP community list for the source prefix, corresponding to a specific flow's source IP address (Information Element 484, per Appendix A of RFC 8549). <ul style="list-style-type: none">• Default: 8
bgp-destination-standard-community count	Export the standard BGP community list for the destination prefix, corresponding to a specific flow's destination IP address (Information Element 485, per Appendix A of RFC 8549). <ul style="list-style-type: none">• Default: 8
bgp-source-extended-community count	Export the extended BGP community list corresponding to a specific flow's source IP address (Information Element 487, per Appendix A of RFC 8549). <ul style="list-style-type: none">• Default: 8
bgp-destination-extended-community count	Export the extended BGP community list corresponding to a specific flow's destination IP address (Information Element 488, per Appendix A of RFC 8549). <ul style="list-style-type: none">• Default: 8
bgp-source-large-community count	Export the large BGP community list corresponding to a specific flow's source IP address (Information Element 490, per Appendix A of RFC 8549).

- Default: 8

bgp-destination-large-community *count*

Export the large BGP community list corresponding to a specific flow's destination IP address (Information Element 491, per Appendix A of RFC 8549).

- Default: 8

Additional Information

AS path and BGP community information are exported using the following Information Elements, as shown in [Table 136 on page 1036](#):

Table 136: AS Path and BGP Community Information Elements

IE Name	IE ID	Export Data Type Format
BGP Community	483	Unsigned 32-bit value
BGP Source Community List, which is a list of BGP community numbers (IE 483)	484	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Destination Community List, which is a list of BGP community numbers (IE 483)	485	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Extended Community	486	8-byte octetArray
BGP Source Extended Community List, which is a list of BGP extended community numbers (IE 486)	487	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Destination Extended Community List, which is a list of BGP extended community numbers (IE 486)	488	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Large Community	489	12-byte octetArray

Table 136: AS Path and BGP Community Information Elements *(Continued)*

IE Name	IE ID	Export Data Type Format
BGP Source Large Community List, which is a list of BGP large community numbers (IE 489)	490	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
BGP Destination Large Community List, which is a list of BGP large community numbers (IE 489)	491	basicList: data is exported in the format defined in the example in Appendix A of RFC 8549.
Source AS Path List	16	basicList: data is exported in the format defined in RFC 6313.
Destination AS Path List	17	basicList: data is exported in the format defined in RFC 6313.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.4R1.

data-size

IN THIS SECTION

- [Syntax | 1038](#)
- [Junos OS Hierarchy Levels | 1038](#)
- [Junos OS Evolved Hierarchy Level | 1038](#)

- [Description | 1038](#)
- [Options | 1038](#)
- [Required Privilege Level | 1039](#)
- [Release Information | 1039](#)

Syntax

```
data-size size;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the size of the data portion of ICMP probes. The data-size statement is not valid with the http-get or http-metadata-get probe type.

Options

size—0 through 65400 for RPM, for TWAMP the value is from 60 through 1400.

- **Default:** 0 for RPM and 60 for TWAMP.

NOTE: If you configure the hardware timestamp feature (see ["Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches"](#) on page 663):

- The default value of data-size is 32 bytes and 32 is the minimum value for explicit configuration. The UDP timestamp probe type is an exception; it requires a minimum data size of 52 bytes.
- The data size must be at least 100 bytes smaller than the default MTU of the interface of the RPM client interface.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

delay-factor

IN THIS SECTION

● [Syntax | 1040](#)

- Hierarchy Level | 1040
- Description | 1040
- Required Privilege Level | 1041
- Release Information | 1041

Syntax

```
delay-factor {  
    no-syslog-generation;  
    generate-snmp-traps;  
    storm-control {  
        count number;  
        interval number;  
    }  
    alarm-mode {  
        mdi-records-count number;  
        average;  
    }  
}
```

Hierarchy Level

[edit services]

Description

Configure the maximum observed time difference between the arrival of media data and the drain of media data. The delay factor suggests the minimum size of the buffer required at the next downstream node. As a stream progresses, the variation of the delay factor indicates packet bunching or packet gaps (jitter). Greater delay factor values also indicate that more network latency is needed to deliver a stream because of the need to pre-fill a receive buffer before beginning the drain to guarantee no underflow.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers | 927](#)
[alarms | 973](#)

delegate-probes

IN THIS SECTION

- [Syntax | 1041](#)
- [Hierarchy Level | 1041](#)
- [Description | 1042](#)
- [Required Privilege Level | 1042](#)
- [Release Information | 1042](#)

Syntax

```
delegate-probes;
```

Hierarchy Level

```
[edit services rpm probe owner]
```

Description

Generate real-time performance monitoring (RPM) probes on an MS-MPC or MS-MIC card, which increases the number of RPM probes that can run at the same time.

To use the `delegate-probes` statement, you must first configure the `destination-interface` statement at the `[edit services rpm probe owner test test-name]` hierarchy level to point to a valid logical unit number of a multiservices interface. Then configure the same unit and multiservice interface with the `rpm client-delegate-probes` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

The probe-type `type` at the `[edit services rpm probe owner test test-name]` hierarchy level can be `icmp-ping` or `icmp-ping-timestamp` starting in Junos OS Release 17.3R1, and `icmp6-ping` starting in Junos OS Release 18.1R1.

To avoid packet bursts in the network due to RPM, probes will be distributed in a better way.

The chances of multiple tests starting and ending at the same time are smaller. This way RPM syslog bursts and a potential performance bottleneck in event-processing are avoided.. This does not exclude potential syslog drops on the RE if more than 12000 RPM tests are running simultaneously. For scaled configurations (with more than 12000 RPM tests) we recommend you to configure syslogs to sent to an external hosts for offloaded processing.

NOTE: You cannot configure the `routing-instance` statement at the `[edit services rpm probe owner test test-name]` hierarchy level for RMP probes that are generated on an MS-MPC or MS-MIC card.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.3R1.

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[client-delegate-probes | 1002](#)

destination (Interfaces)

IN THIS SECTION

- [Syntax | 1043](#)
- [Hierarchy Level | 1043](#)
- [Description | 1043](#)
- [Options | 1044](#)
- [Required Privilege Level | 1044](#)
- [Release Information | 1044](#)

Syntax

```
destination address;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number tunnel],  
[edit interfaces interface-name unit logical-unit-number family inet address address],  
[edit interfaces interface-name unit logical-unit-number tunnel],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number  
family inet address address]
```

Description

For CoS on ATM interfaces, specify the remote address of the connection.

For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.

For tunnel and encryption interfaces, specify the remote address of the tunnel.

Options

address—Address of the remote side of the connection.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Linear RED Profiles on ATM Interfaces

Configuring Link and Multilink Services Logical Interfaces

Configuring Encryption Interfaces

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

[Configuring Flow Monitoring | 5](#)

Configuring Unicast Tunnels

destination-address (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax | 1045](#)
- [Hierarchy Level | 1045](#)
- [Description | 1045](#)
- [Options | 1045](#)
- [Required Privilege Level | 1045](#)
- [Release Information | 1045](#)

Syntax

```
destination-address address;
```

Hierarchy Level

```
[edit services jflow-log collector collector-name]
```

Description

Specify the destination IP address or identifier of the host or external device that functions as the collector for receiving the generated flow monitoring logs that are sent from the exporter. You can configure an IPv4 address, or an identifier of the host system (the workstation either running the Jflow utility or collecting traffic flows using version 9 or IPFIX format). For external NetFlow collectors or servers, the hostname must be reachable from the same routing instance to which the initial data packet (that triggered session establishment) is delivered. You can specify a maximum of eight collectors per profile.

Options

address—Destination hostname, or IPv4 or IPv6 address of the collector.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

destination-interface

IN THIS SECTION

- [Syntax | 1046](#)
- [Hierarchy Level | 1046](#)
- [Description | 1046](#)
- [Options | 1047](#)
- [Required Privilege Level | 1047](#)
- [Release Information | 1047](#)

Syntax

```
destination-interface interface-name;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name],  
[edit services rpm probe-server (tcp | udp)],  
[edit services rpm twamp client control-connection control-client-name]
```

Description

On M Series and T Series routers, specify a services (sp-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the sp- interface and include the unit 0 family inet statement with a /32 address.

On M Series, MX Series, and T Series routers, specify a multiservices (ms-) interface that adds a timestamp to RPM probe messages. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types. You must also configure the rpm statement on the ms- interface and include the unit 0 family inet statement with a /32 address.

The inline service interface (si- interface) is a virtual physical service interface that resides on the Packet Forwarding Engine to provide L2TP services without a special services PIC. The inline service interface is supported only by MPCs on MX Series routers. Four inline service interfaces are configurable per MPC-occupied chassis slot. Specify a multiservices (si-) interface that adds a timestamp to TWAMP probe messages. You must also configure the rpm twamp-client or twamp-server statement on the si- interface and include the unit 0 family inet statement with a /32 address.

To enable RPM for the extension-provider packages on the adaptive services interface, configure the object-cache-size, policy-db-size, and package statements at the [edit chassis fpc *slot-number* pic *pic-number* adaptive-services service-package extension-provider] hierarchy level. For the extension-provider package, *package-name* in the package *package-name* statement is jservices-rpm.

Starting in Junos OS Release 17.3R1, you can use destination-interface *interface-name.logical-unit-number* at the [edit services rpm probe *owner* test *test-name*] hierarchy level to configure the generation of probes on an MS-MPC or MS-MIC. You must also include the delegate-probes statement at the [edit services rpm probe *owner*] hierarchy level and the rpm client-delegate-probes and the family (inet | inet6) address *address* statements at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

Options

interface-name—Name of the adaptive services interface.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.5.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

RELATED DOCUMENTATION

[Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches](#) | 663

[Configuring RPM Receiver Servers](#) | 662

[Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches](#) | 663

[hardware-timestamp](#) | 1147

[rpm \(Interfaces\)](#) | 1357

[Enabling RPM on MX, M and T Series Routers and SRX Firewalls for the Services SDK](#) | 685

destination-ipv4-address

IN THIS SECTION

- [Syntax](#) | 1048
- [Junos OS Hierarchy Level](#) | 1048
- [Junos OS Evolved Hierarchy Level](#) | 1049
- [Description](#) | 1049
- [Options](#) | 1049
- [Required Privilege Level](#) | 1049
- [Release Information](#) | 1049

Syntax

```
destination-ipv4-address address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the destination IPv4 address to be used in generated test frames. You must configure this option if you specify `inet` as the family. This option is not required if you specify `ccc` as the family.

Options

address (Required if you specify `inet` as the family.) Valid IPv4 address.

- **Default:** When you specify `ccc` as the family, if you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test](#) | 739

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices](#) | 728

[rfc2544-benchmarking](#) | 1347

destination-mac-address

IN THIS SECTION

- [Syntax | 1050](#)
- [Junos OS Hierarchy Level | 1050](#)
- [Junos OS Evolved Hierarchy Level | 1050](#)
- [Description | 1050](#)
- [Options | 1051](#)
- [Required Privilege Level | 1051](#)
- [Release Information | 1051](#)

Syntax

```
destination-mac-address mac-address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the destination MAC address used in the generated test frames. This is a mandatory parameter for family bridge and vpls and optional for family ccc.

Options

mac-address Specify the MAC address as six hexadecimal bytes in one of the following formats:
nnnn.nnnn.nnnn or *nn:nn:nn:nn:nn:nn*—for example, 0000:5e00:5355 or 00:00:5e:00:53:55.

Default: 0x00:0x11:0xAE:0x92:0x2F:0x28

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[rfc2544-benchmarking](#) | [1347](#)

[rfc2544](#) | [1345](#)

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | [855](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices](#) | [728](#)

[Configuring an RFC 2544-Based Benchmarking Test](#) | [739](#)

destination-port

IN THIS SECTION

- [Syntax](#) | [1052](#)
- [Junos OS Hierarchy Levels](#) | [1052](#)
- [Junos OS Evolved Hierarchy Level](#) | [1052](#)
- [Description](#) | [1052](#)

- Options | 1053
- Required Privilege Level | 1053
- Release Information | 1053

Syntax

```
destination-port port;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
[edit services rpm twamp client control-connection control-client-name test-session test-session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) port to which a probe is sent. This statement is used only for TCP or UDP probe types.

The value for the `destination-port` can be only 7 when you configure the destination port along with hardware timestamping. A constraint check prevents you for configuring any other value for the destination port in this case.

This constraint does not apply when you are using one-way hardware timestamping along with `destination-port` and either `probe-type udp-ping` or `probe-type udp-ping-timestamp`.

For a managed TWAMP control connection (the default or configured with the `control-type managed` statement), the `destination-port` statement is required at the `[edit services rpm twamp control-connection`

control-connection-name] hierarchy level (Junos OS) or the [edit services monitoring twamp control-connection *control-connection-name*] hierarchy level (Junos OS Evolved) and is not allowed at the individual test-session level, at [edit services rpm twamp control-connection *control-connection-name* test-session *test-session-name*] (Junos OS) or [edit services monitoring twamp control-connection *control-connection-name* test-session *test-session-name*] (Junos OS Evolved).

For a TWAMP control connection configuration that includes the control-type light statement, the destination-port statement is not allowed at the [edit services rpm twamp control-connection *control-connection-name*] hierarchy level (Junos OS) or the [edit services monitoring twamp control-connection *control-connection-name*] hierarchy level (Junos OS Evolved), but is required for each test session, at [edit services rpm twamp control-connection *control-connection-name* test-session *test-session-name*] (Junos OS) or [edit services monitoring twamp control-connection *control-connection-name* test-session *test-session-name*] (Junos OS Evolved).

Options

port—Port number 7 or from 862 through 65,535.

- **Default:** The default value for the port is 862 to which the TWAMP client establishes the control connection.

NOTE: The specified port numbers are recommended for RPM only.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport Routers.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

Support at the [edit services rpm twamp client control-connection *control-client-name* test-session *test-session-name*] hierarchy level for Junos OS Release 21.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#) | **646**

[Configuring BGP Neighbor Discovery Through RPM](#) | **671**

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | **651**

destination-port (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax](#) | **1054**
- [Hierarchy Level](#) | **1054**
- [Description](#) | **1054**
- [Options](#) | **1055**
- [Required Privilege Level](#) | **1055**
- [Release Information](#) | **1055**

Syntax

```
destination-port port-number;
```

Hierarchy Level

```
[edit services jflow-log collector collector-name]
```

Description

Specify the UDP port of the destination to be used in the UDP header for the generated flow monitoring logs. This is a required setting.

Options

port-number—UDP port number for the test frames.

- **Default:** 4041

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

destination-udp-port

IN THIS SECTION

- [Syntax | 1056](#)
- [Junos OS Hierarchy Level | 1056](#)
- [Junos OS Evolved Hierarchy Level | 1056](#)
- [Description | 1056](#)
- [Options | 1056](#)

- Required Privilege Level | 1056
- Release Information | 1057

Syntax

```
destination-udp-port port-number;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

(Required if you specify `inet` as the family.) Specify the UDP port of the destination to be used in the UDP header for the generated frames. For other families, if you do not specify the UDP port, the default value of 4041 is used.

Options

- port-number* UDP port number for the test frames
- **Default:** 4041

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)

destinations

IN THIS SECTION

- [Syntax | 1057](#)
- [Hierarchy Level | 1058](#)
- [Description | 1058](#)
- [Required Privilege Level | 1058](#)
- [Release Information | 1058](#)

Syntax

```
destinations {  
  ftp:url {  
    password "password";  
  }  
}
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Specify the primary and secondary destination FTP servers.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

direction

IN THIS SECTION

- [Syntax](#) | [1059](#)
- [Junos OS Hierarchy Level](#) | [1059](#)
- [Junos OS Evolved Hierarchy Level](#) | [1059](#)
- [Description](#) | [1059](#)
- [Options](#) | [1059](#)
- [Required Privilege Level](#) | [1059](#)

Syntax

```
direction (egress | ingress);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the direction of the interface on which the test must be run. This parameter is valid only for a ccc , a ethernet-switching or a bridge family. RFC2544 tests are supported only in the egress direction or the user-to-network interface (UNI) direction of an E-line or E-LAN service parameters in a bridge domain between two routers for unicast traffic. You cannot compute the NNI direction of Ethernet services between two routers for multicast or broadcast traffic.

Options

- egress** Run the test in the egress direction of the interface (network-to-network interface (NNI)). This option is applicable for a ccc, bridge, or ethernet-switching family.
- ingress** Run the test in the ingress direction of the interface (user-to-network interface (UNI)). You cannot configure this option for a bridge or ethernet-switching family.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[rfc2544-benchmarking | 1347](#)

[rfc2544 | 1345](#)

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

disable (Forwarding Options)

IN THIS SECTION

- [Syntax | 1060](#)
- [Hierarchy Level | 1061](#)
- [Description | 1061](#)
- [Required Privilege Level | 1061](#)
- [Release Information | 1061](#)

Syntax

```
disable;
```

Hierarchy Level

```
[edit forwarding-options port-mirror],
[edit forwarding-options port-mirror instance instance-name],
[edit forwarding-options sampling],
[edit forwarding-options sampling instance instance-name],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls) ],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls) output file]
```

Description

Disable traffic accounting, port mirroring, or sampling.

NOTE: The disable statement at the [edit forwarding-options sampling] hierarchy level disables only Routing Engine-based sampling. To disable PIC-based sampling and inline sampling, include the disable statement at the [edit forwarding-options sampling instance *instance-name*] hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement added to port-mirror hierarchy in Junos OS Release 9.6.

NOTE: Beginning in Junos OS Release 15.1F5 and later 15.1 releases and Junos OS Release 16.1 and later, the disable option has been deprecated at the forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) hierarchy level on PTX3000 Series routers. When configured, the option does not take effect, so packets continue to be sampled. Instead of the disable option, use the deactivate forwarding-options sampling instance *instance-name* family (inet | inet6 | mpls) command to prevent sampling.

RELATED DOCUMENTATION

Disabling Traffic Sampling

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

[Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

disable-signature-check

IN THIS SECTION

- [Syntax](#) | 1062
- [Junos OS Hierarchy Level](#) | 1062
- [Junos OS Evolved Hierarchy Level](#) | 1062
- [Description](#) | 1063
- [Required Privilege Level](#) | 1063
- [Release Information](#) | 1063

Syntax

```
disable-signature-check;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Disable signature verification on the received test frames. This statement is valid only if you configure the test mode to be a reflector. The configuration is useful when the test traffic is generated using a third-party vendor tool, instead of an ACX Series router.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | 737](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

dscp (flow-server)

IN THIS SECTION

- [Syntax | 1064](#)
- [Hierarchy Level | 1064](#)
- [Description | 1064](#)
- [Options | 1064](#)
- [Required Privilege Level | 1064](#)
- [Release Information | 1064](#)

Syntax

```
dscp dscp-value
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet | inet6) output flow-  
server hostname]
```

Description

Specify the Differentiated Services Code Point (DSCP) mapping that is applied to exported packets for inline active flow monitoring. This allows different levels of service to be assigned to sampled traffic.

Options

dscp
dscp-value Can be a value between 0 and 63 (the default is 0). When the same flow-server is configured under both the inet and inet6 families in a sampling instance, use the same dscp value for both flow-server appearances.

The *dscp-value* is overwritten by the CoS DSCP value if you configure dscp under the [edit class-of-service] hierarchy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F4.

Statement introduced in Junos OS Release 16.1 for the MX Series.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring on PTX Series Routers](#) | 540

[Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers | 615](#)

[Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches | 433](#)

dscp-code-points (RPM)

IN THIS SECTION

- [Syntax | 1065](#)
- [Junos OS Hierarchy Levels | 1065](#)
- [Junos OS Evolved Hierarchy Level | 1065](#)
- [Description | 1066](#)
- [Options | 1066](#)
- [Required Privilege Level | 1067](#)
- [Release Information | 1067](#)

Syntax

```
dscp-code-points dscp-bits;
```

Junos OS Hierarchy Levels

```
[edit services rpm probe owner test test-name],  
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.

Options

dscp-bits—A valid 6-bit pattern; for example, 001111, or one of the following configured DSCP aliases:

- af11—Default: 001010
- af12—Default: 001100
- af13—Default: 001110
- af21—Default: 010010
- af22—Default: 010100
- af23 —Default: 010110
- af31 —Default: 011010
- af32 —Default: 011100
- af33 —Default: 011110
- af41 —Default: 100010
- af42 —Default:100100
- af43 —Default:100110
- be—Default: 000000
- cs1—Default: 001000
- cs2—Default: 010000
- cs3—Default: 011000
- cs4—Default: 100000
- cs5—Default: 101000
- cs6—Default: 110000
- cs7—Default: 111000

- ef—Default: 101110
- nc1—Default: 110000
- nc2—Default: 111000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release for PTX Series Packet Transport routers.

Support at the `[edit services rpm twamp client control-connection control-client-name]` hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

dscp-code-points (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1068](#)
- [Hierarchy Level | 1068](#)

- [Description | 1068](#)
- [Options | 1068](#)
- [Required Privilege Level | 1069](#)
- [Release Information | 1069](#)

Syntax

```
dscp-code-points dscp-bits;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the value of the Differentiated Services (DiffServ) field within the IP header of the test frames. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern. If you do not specify this value, 0 is used in the DSCP fields in the IP header.

Options

dscp-bits—A valid 6-bit pattern; for example, 001111, or one of the following configured DSCP aliases:

- af11—Default: 001010
- af12—Default: 001100
- af13—Default: 001110
- af21—Default: 010010
- af22—Default: 010100
- af23 —Default: 010110
- af31 —Default: 011010

- af32 —Default: 011100
- af33 —Default: 011110
- af41 —Default: 100010
- af42 —Default:100100
- af43 —Default:100110
- be—Default: 000000
- cs1—Default: 001000
- cs2—Default: 010000
- cs3—Default: 011000
- cs4—Default: 100000
- cs5—Default: 101000
- cs6—Default: 110000
- cs7—Default: 111000
- ef—Default: 101110
- nc1—Default: 110000
- nc2—Default: 111000

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

dump-on-flow-control

IN THIS SECTION

- [Syntax](#) | [1070](#)
- [Hierarchy Level](#) | [1070](#)
- [Description](#) | [1070](#)
- [Required Privilege Level](#) | [1070](#)
- [Release Information](#) | [1071](#)

Syntax

```
dump-on-flow-control;
```

Hierarchy Level

```
[edit interfaces interface-name multiservice-options]
```

Description

This option supports high availability functionality and can be used with various service interfaces, including `rsp`, `rms`, `lsq`, and `rlsq`.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.5.

RELATED DOCUMENTATION

[Configuring Multiservice Physical Interface Properties](#)

[Junos OS Services Interfaces Library for Routing Devices](#)

passive-monitor-mode

duplicates-dropped-periodicity

IN THIS SECTION

- [Syntax | 1071](#)
- [Hierarchy Level | 1071](#)
- [Description | 1072](#)
- [Options | 1072](#)
- [Required Privilege Level | 1072](#)
- [Release Information | 1072](#)

Syntax

```
duplicates-dropped-periodicity seconds;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the `max-duplicates` threshold has been reached.

Options

seconds—Period for sending `DuplicatesDropped` notifications.

- **Default:** 30 seconds

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[g-duplicates-dropped-periodicity](#) | [1139](#)

[Configuring Junos Capture Vision](#) | [289](#)

[max-duplicates](#) | [1216](#)

dynamic-flow-capture

IN THIS SECTION

- [Syntax](#) | [1073](#)
- [Hierarchy Level](#) | [1073](#)
- [Description](#) | [1073](#)
- [Required Privilege Level](#) | [1074](#)
- [Release Information](#) | [1074](#)

Syntax

```
dynamic-flow-capture {
  capture-group client-name {
    content-destination identifier {
      address address;
      hard-limit bandwidth;
      hard-limit-target bandwidth;
      soft-limit bandwidth;
      soft-limit-clear bandwidth;
      ttl hops;
    }
    control-source identifier {
      allowed-destinations [ destinations ];
      minimum-priority value;
      no-syslog;
      notification-targets address port port-number;
      service-port port-number;
      shared-key value;
      source-addresses [ addresses ];
    }
    duplicates-dropped-periodicity seconds;
    input-packet-rate-threshold rate;
    interfaces interface-name;
    max-duplicates number;
    pic-memory-threshold percentage percentage;
  }
  g-duplicates-dropped-periodicity seconds;
  g-max-duplicates number;
}
```

Hierarchy Level

[edit services]

Description

Define the dynamic flow capture properties to be applied to traffic.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

[Understanding Junos Capture Vision](#) | 286

em-hw-profile

IN THIS SECTION

- [Syntax](#) | 1074
- [Hierarchy Level](#) | 1074
- [Description](#) | 1075
- [Default](#) | 1075
- [Required Privilege Level](#) | 1075
- [Release Information](#) | 1075

Syntax

```
em-hw-profile;
```

Hierarchy Level

```
[edit chassis forwarding-options]
```

Description

(QFX5120 only) Configure a unified forwarding table (UFT) profile to allocate the amount of MAC address, layer 3 host, longest prefix match (LPM), and exact-match (EM) memory available for software-based flow-based telemetry (FBT) for VXLANs. You must reboot the switch after committing this statement to the configuration to create the UFT profile. The software allocates these amounts of memory in the UFT for each of the following types, in bytes:

- MAC address: 48K
- Layer 3 host address: 48K
- LPM memory: 16K
- EM memory: 32K

Default

The UFT uses the default profile.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

[Flow-Based Telemetry for VXLANs \(QFX5120\)](#) | 361

engine-id (Forwarding Options)

IN THIS SECTION

- [Syntax | 1076](#)
- [Hierarchy Level | 1076](#)
- [Description | 1076](#)
- [Options | 1076](#)
- [Required Privilege Level | 1077](#)
- [Release Information | 1077](#)

Syntax

```
engine-id number;
```

Hierarchy Level

```
[edit forwarding-options accounting name output interface interface-name],  
[edit forwarding-options monitoring name output interface interface-name],  
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output  
interface interface-name],  
[edit forwarding-options sampling family (inet | inet6 | mpls) output interface interface-name]
```

Description

Specify the engine ID number for flow monitoring and accounting services.

Options

number—Identity of accounting interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

[Configuring Flow Monitoring | 5](#)

[Configuring Discard Accounting | 435](#)

engine-type

IN THIS SECTION

- [Syntax | 1077](#)
- [Hierarchy Level | 1078](#)
- [Description | 1078](#)
- [Options | 1078](#)
- [Required Privilege Level | 1078](#)
- [Release Information | 1078](#)

Syntax

```
engine-type number;
```


Hierarchy Level

```
[edit forwarding-options accounting name output interface interface-name],
[edit forwarding-options monitoring name output interface interface-name],
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output
interface interface-name],
[edit forwarding-options sampling family (inet | inet6 | mpls) output interface interface-name]
```

Description

Specify the engine type number for flow monitoring and accounting services. The engine type attribute refers to the type of the flow switching engine, such as the route processor or a line module. The configured engine type is inserted in output cflowd packets. The Source ID, a 32-bit value to ensure uniqueness for all flows exported from a particular device, is the equivalent of the engine type and the engine ID fields.

NOTE: You must configure a source address in the output interface statements. The interface-level statement of engine-type is added automatically but you can override this value with manually configured statements to track different flows with a single cflowd collector.

Options

number—Platform-specific accounting interface type.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

[Configuring Flow Monitoring | 5](#)[Configuring Discard Accounting | 435](#)

exception-reporting

IN THIS SECTION

- [Syntax | 1079](#)
- [Hierarchy Level | 1079](#)
- [Description | 1079](#)
- [Options | 1080](#)
- [Required Privilege Level | 1080](#)
- [Release Information | 1080](#)

Syntax

```
exception-reporting {  
  category category-name {  
    inline-monitoring-instance inline-monitoring-instance;  
  }  
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pfe identifier]
```

Description

Enable reporting of exceptions in the forwarding path.

Options

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

[Juniper Resiliency Interface](#) | 409

exceptions

IN THIS SECTION

- [Syntax \(Junos OS\)](#) | 1081
- [Syntax \(Junos OS Evolved\)](#) | 1081
- [Hierarchy Level](#) | 1081
- [Description](#) | 1081
- [Default](#) | 1081
- [Options](#) | 1081
- [Required Privilege Level](#) | 1082
- [Release Information](#) | 1082

Syntax (Junos OS)

```
exceptions {  
  forwarding;  
  os;  
  routing;  
}
```

Syntax (Junos OS Evolved)

```
exceptions {  
  forwarding;  
  routing;  
}
```

Hierarchy Level

```
[edit system resiliency]
```

Description

Subscribe to forwarding, operating system, and routing exceptions for the on-box collector.

Default

The on-box collector is disabled.

Options

- | | |
|-------------------|---|
| forwarding | Subscribe to forwarding exceptions. |
| os | (Junos OS only) Subscribe to operating system exceptions. |
| routing | Subscribe to routing exceptions. |

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

[Juniper Resiliency Interface](#) | 409

[inline-monitoring](#) | 1163

[primary-data-record-fields](#) | 1310

[exception-reporting](#) | 1079

export-format

IN THIS SECTION

- [Syntax](#) | 1082
- [Hierarchy Level](#) | 1083
- [Description](#) | 1083
- [Options](#) | 1083
- [Required Privilege Level](#) | 1083
- [Release Information](#) | 1083

Syntax

```
export-format format;
```

Hierarchy Level

```
[edit forwarding-options monitoring name output]
```

Description

Flow monitoring export format.

Options

format—Format of the flows.

- **Values:** 5 or 8
- **Default:** 5

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[version | 1533](#)

[Configuring Flow Monitoring | 5](#)

family (Monitoring)

IN THIS SECTION

[Syntax | 1084](#)

- Hierarchy Level | 1085
- Description | 1085
- Required Privilege Level | 1085
- Release Information | 1085

Syntax

```
family inet {
  output {
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    export-format format;
    cflowd hostname {
      aggregation {
        autonomous-system;
        destination-prefix;
        protocol-port;
        source-destination-prefix {
          caida-compliant;
        }
        source-prefix;
      }
      port port-number;
    }
    interface interface-name {
      engine-id number;
      engine-type number;
      input-interface-index number;
      output-interface-index number;
      source-address address;
    }
  }
}
```

Hierarchy Level

```
[edit forwarding-options monitoring name]
```

Description

Specify input and output interfaces and properties for flow monitoring. Only IPv4 (*inet*) is supported.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

family

IN THIS SECTION

- [Syntax](#) | 1086
- [Junos OS Hierarchy Level](#) | 1086
- [Junos OS Evolved Hierarchy Level](#) | 1086
- [Description](#) | 1086
- [Options](#) | 1086
- [Required Privilege Level](#) | 1087

Syntax

```
family (bridge | ccc | ethernet-switching | inet | vpls);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Configure the address type family for the benchmarking test.

Options

bridge	(Junos OS) Run the test on a Layer 2 Ethernet line (E- Line) or an Ethernet LAN (E-LAN) service configured in a bridge domain. You can run the RFC2544-based benchmarking test only in the egress direction or the user-to-network interface (UNI) direction of an Ethernet line.
ccc	Run the test on a layer 2 VPN, circuit cross-connect (CCC), or Ethernet pseudowire service (such as EVPN-VPWS). You can run the RFC2544-based benchmarking test either in the egress or ingress direction.
ethernet-switching	(Junos OS Evolved) Run the test over a Layer 2 bridge service, EVPN-MPLS network, or a Virtual Private LAN Service (VPLS). You can run the RFC2544-based benchmarking test only in the egress direction or the user-to-network interface (UNI) direction of an Ethernet line.
inet	Run the test on an IPv4 service.

vp1s (Junos OS) Run the test over a Virtual Private LAN Service (VPLS).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

bridge option introduced in Junos OS Release 12.3X53 for ACX Series routers.

bridge option introduced in Junos OS Release 14.2 for MX104 Universal Routing Platforms.

Statement introduced in Junos OS Evolved Release 21.1R1 for the **inet** option only.

ccc and **ethernet-switching** options introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[rfc2544-benchmarking](#) | [1347](#)

[rfc2544](#) | [1345](#)

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | [855](#)

[Configuring an RFC 2544-Based Benchmarking Test](#) | [739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices](#) | [728](#)

family (Sampling)

IN THIS SECTION

- [Syntax](#) | [1088](#)
- [Hierarchy Level](#) | [1089](#)
- [Description](#) | [1089](#)
- [Required Privilege Level](#) | [1090](#)

Syntax

```

family (inet | inet6 | mpls | vpls | bridge) {
    disable;
    input {
        max-packets-per-second max-packets-per-second;
        maximum-packet-length maximum-packet-length;
        rate rate;
        run-length run-length
    }

    output {
        aggregate-export-interval seconds;
        flow-active-timeout seconds;
        flow-inactive-timeout seconds;
        extension-service service-name;
        flow-server hostname {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            dscp dscp-value;
            forwarding-class class-name;
            (local-dump | no-local-dump);
            port port-number;
            source-address address;
            version format;
            version9 {
                template template-name;
            }
            version-ipfix {

```

```

        template template-name;
    }
}
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
file {
    disable;
    filename filename;
    files number;
    size bytes;
    (stamp | no-stamp);
    (world-readable | no-world-readable);
}
inline-jflow {
    source-address address;
    flow-export-rate rate;
}
}
}

```

Hierarchy Level

```

[edit forwarding-options sampling],
[edit forwarding-options sampling instance instance-name]
[edit forwarding-options sampling instance instance-name family (inet | inet6 | bridge)

```

Description

Configure the protocol family to be sampled. IPv4 (**inet**) is supported for most purposes, but you can configure **family mpls** to collect and export MPLS label information, **family inet6** to collect and export IPv6 traffic using flow aggregation version 9, and **vpls** to collect and export VPLS information, and **bridge** to collect and export bridge information.

The remaining statements are explained separately. See [CLI Explorer](#).

NOTE: The `inline-jflow` statement is valid only under the `[edit forwarding-options sampling instance instance-name family inet output]` hierarchy level. The `file` statement is valid only under the `[edit forwarding-options sampling family inet output]` hierarchy level.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

mpls option introduced in Release 8.3.

inet6 option introduced in Release 9.4.

vpls option added in Junos OS Release 13.2 for MX Series routers.

bridge option introduced in Release 18.2R1 for MX Series routers.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

features

IN THIS SECTION

- [Syntax](#) | 1091
- [Hierarchy Level](#) | 1091
- [Description](#) | 1091
- [Options](#) | 1091
- [Required Privilege Level](#) | 1092

Syntax

```
features {
  aggregate-intf-member-id;
  chip-delay;
  egress-drop-reason;
  flow-start-end-time;
  ingress-drop-reason;
  inter-arrival-time;
  inter-departure-time;
  queue-congestion-level;
  security-enable;
  shared-pool-congestion;
}
```

Hierarchy Level

```
[edit services inline-monitoring feature-profile]
```

Description

When you monitor packets, you need to specify what information you want to collect about them.

Options

- aggregate-intf-member-id** ID for a member of a link aggregation group (LAG) or an equal-cost multipath (ECMP) group.
- chip-delay** The amount of time the packet takes to transit the ASIC.
- egress-drop-reason** The reason the packet is dropped at egress.

flow-start-end-time	The flow start and end time.
ingress-drop-reason	The reason the packet is dropped at ingress.
inter-arrival-time	The time difference between two consecutive packets at ingress.
inter-departure-time	The time difference between two consecutive packets at egress.
queue-congestion-level	Queue congestion level
security-enable	Enable security analytics; specify that Denial-of-Service (DoS) attacks are reported to the collector.
shared-pool-congestion	Shared pool congestion level

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

[Flow-Based Telemetry \(EX4100, EX4100-F, and EX4400 Series\)](#) | 349

file (Sampling)

IN THIS SECTION

- [Syntax](#) | 1093
- [Hierarchy Level](#) | 1093
- [Description](#) | 1093

- Required Privilege Level | 1093
- Release Information | 1093

Syntax

```
file {  
    disable;  
    filename filename;  
    files number;  
    size bytes;  
    (stamp | no-stamp);  
    (world-readable | no-world-readable);  
}
```

Hierarchy Level

```
[edit forwarding-options sampling family inet output]
```

Description

Collect the traffic samples in a file.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

file (Trace Options)

IN THIS SECTION

- [Syntax](#) | 1094
- [Hierarchy Level](#) | 1094
- [Description](#) | 1094
- [Options](#) | 1094
- [Required Privilege Level](#) | 1095
- [Release Information](#) | 1095

Syntax

```
file filename <files number <size bytes> <world-readable | no-world-readable>;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions],  
[edit forwarding-options sampling traceoptions]
```

Description

Configure information about the files that contain trace logging information.

Options

filename—Name of the file containing the trace information.

- **Default:** /var/log/sampled

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

file-specification (File Format)

IN THIS SECTION

- [Syntax](#) | 1095
- [Hierarchy Level](#) | 1096
- [Description](#) | 1096
- [Required Privilege Level](#) | 1096
- [Release Information](#) | 1096

Syntax

```
file-specification {  
  variant variant-number {  
    data-format format;  
    name-format format;  
    transfer {  
      record-level number;
```

```

        timeout seconds;
    }
}

```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the file format for the flow collection files.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

file-specification (Interface Mapping)

IN THIS SECTION

- [Syntax](#) | 1097
- [Hierarchy Level](#) | 1097

- [Description | 1097](#)
- [Options | 1097](#)
- [Required Privilege Level | 1097](#)
- [Release Information | 1097](#)

Syntax

```
file-specification {  
    variant variant-number;  
}
```

Hierarchy Level

```
[edit services flow-collector interface-map]
```

Description

Configure the default file specification for interface mapping.

Options

variant-number—Default file format variant.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

filename

IN THIS SECTION

- [Syntax | 1098](#)
- [Hierarchy Level | 1098](#)
- [Description | 1098](#)
- [Options | 1098](#)
- [Required Privilege Level | 1098](#)
- [Release Information | 1099](#)

Syntax

```
filename filename;
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet |inet6 |mpls) output file]
```

Description

Configure the name of the output file.

Options

filename—Name of the file in which to place the traffic samples. All files are placed in the directory **/var/tmp**.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

filename-prefix

IN THIS SECTION

- [Syntax](#) | 1099
- [Hierarchy Level](#) | 1099
- [Description](#) | 1099
- [Options](#) | 1100
- [Required Privilege Level](#) | 1100
- [Release Information](#) | 1100

Syntax

```
filename-prefix prefix;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive]
```

Description

Configure the filename prefix for log files.

Options

prefix—Filename identifier.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

files

IN THIS SECTION

- [Syntax](#) | [1100](#)
- [Hierarchy Level](#) | [1101](#)
- [Description](#) | [1101](#)
- [Options](#) | [1101](#)
- [Required Privilege Level](#) | [1101](#)
- [Release Information](#) | [1101](#)

Syntax

```
files number;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],  
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],  
[edit forwarding-options sampling traceoptions file]
```

Description

Configure the total number of files to be saved with samples or trace data.

Options

number—Maximum number of traffic sampling or trace log files. When a file named *sampling-file* reaches its maximum size, it is renamed *sampling-file.0*, then *sampling-file.1*, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.

- **Range:** 1 through 100 files
- **Default:** 5 files for sampling output; 10 files for trace log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

filter

IN THIS SECTION

- [Syntax | 1102](#)
- [Hierarchy Level | 1102](#)
- [Description | 1102](#)
- [Options | 1102](#)
- [Required Privilege Level | 1103](#)
- [Release Information | 1103](#)

Syntax

```
filter {  
    input filter-name;  
    output filter-name;  
    group filter-group-number;  
}
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

Apply a firewall filter to an interface. You can also use filters for encrypted traffic.

Options

group filter-group-number—Use the specified interface to be part of a filter group. The default filter group number is 0.

input filter-name—Use the specified filter to evaluate when packets are received on the interface.

output *filter-name*—Use the specified filter to evaluate when packets are transmitted on the interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Routing Policies, Firewall Filters, and Traffic Policers User Guide](#)

[Configuring Flow Monitoring | 5](#)

flex-flow-sizing

IN THIS SECTION

- [Syntax | 1103](#)
- [Hierarchy Level | 1104](#)
- [Description | 1104](#)
- [Options | 1104](#)
- [Required Privilege Level | 1104](#)
- [Release Information | 1104](#)

Syntax

```
flex-flow-sizing;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services]
```

Description

Configure support for the service creation of flows for inline services sampling. This configuration results in a first-come-first-serve creation of flows. Whichever flow comes first, that is allowed to occupy the flow-table if there is space in the table. Otherwise, the flow is dropped and an error count is created.

NOTE: You cannot configure the `explicit flow-table-sizes` because `flex-flow-sizing` and `explicit flow-table-sizes` are mutually exclusive.

You need not perform `fpc reboot` to change from `flex` to `per family` configuration.

Options

- **Default:** 1K flows for IPv6 and VPLS flows each.
- **Range:** 15 through 256K flows for IPv4.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F5.

NOTE: Workaround for `flex-flow-sizing` on MX204 router for releases earlier than Junos OS Release 17.4R2 or Junos OS Release 18R3:

Replace `flex-flow-sizing` with below configuration and reload the box. Configure `flow-table-size` within a range of 1 through 15. If `flex-flow-sizing` is configured, deactivate or delete the same.

For example:

```
flow-table-size {  
    ipv4-flow-table-size 10;  
    ipv6-flow-table-size 4;  
}
```

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

[Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers | 633](#)

flow-active-timeout

IN THIS SECTION

- [Syntax | 1105](#)
- [Hierarchy Level | 1106](#)
- [Description | 1106](#)
- [Options | 1106](#)
- [Required Privilege Level | 1106](#)
- [Release Information | 1107](#)

Syntax

```
flow-active-timeout seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting name output],
[edit forwarding-options monitoring name output],
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls | vpls) output],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls) output],
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name],
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Set the interval after which an active flow is exported.

NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.

Options

seconds—Duration of the timeout period.

- **Range:** 60 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)
- **Default:** 1800 seconds (for forwarding-options configurations); 60 seconds (for services configurations)

NOTE: In active flow monitoring, the cflowd or flow monitoring version 9 records are exported after a time period that is a multiple of 60 seconds and greater than or equal to the configured active timeout value. For example, if the active timeout value is 90 seconds, the cflowd or flow monitoring version 9 records are exported at 120-second intervals. If the active timeout value is 150 seconds, the cflowd or flow monitoring version 9 records are exported at 180-second intervals, and so forth.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit services flow-monitoring version-ipfix template *template-name*] hierarchy level added in Junos OS Release 10.2.

Support at the [edit services flow-monitoring version9 template *template-name*] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

[Configuring Flow Monitoring | 5](#)

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

flow-collector

IN THIS SECTION

- [Syntax | 1107](#)
- [Hierarchy Level | 1108](#)
- [Description | 1109](#)
- [Required Privilege Level | 1109](#)
- [Release Information | 1109](#)

Syntax

```
flow-collector {
  analyzer-address address;
  analyzer-id name;
```

```

destinations {
  ftp:url {
    password "password";
  }
}
file-specification {
  variant variant-number {
    data-format format;
    name-format format;
    transfer {
      record-level number;
      timeout seconds;
    }
  }
}
interface-map {
  collector interface-name;
  file-specification variant-number;
  interface-name {
    collector interface-name;
    file-specification variant-number;
  }
}
retry number;
retry-delay seconds;
transfer-log-archive {
  archive-sites {
    ftp:url {
      password "password";
      username username;
    }
  }
  filename-prefix prefix;
  maximum-age minutes;
}
}

```

Hierarchy Level

[edit services]

Description

Define the flow collection.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Flow Collection Overview](#) | 225

flow-control-options

IN THIS SECTION

- [Syntax](#) | 1110
- [Hierarchy Level](#) | 1110
- [Description](#) | 1110
- [Usage Guidelines](#) | 1111
- [Required Privilege Level](#) | 1111
- [Release Information](#) | 1111

Syntax

```
flow-control-options {  
    down-on-flow-control;  
    dump-on-flow-control;  
    reset-on-flow-control;  
    up-on-flow-control;  
}
```

Hierarchy Level

[edit [interfaces](#) *mo-fpc/pic/port* [multiservice-options](#)]

Description

Configure the flow control options for application recovery in case of a prolonged flow control failure.

- `down-on-flow-control`—Bring interface down during prolonged flow control.
- `dump-on-flow-control`—Cause core dump during prolonged flow control.

NOTE: Starting with Junos OS Release 15.1, on MX Series routers with MS-MICs and MS-MPCs, instead of an eJunos kernel core file, the multiservices PIC management daemon (mspmnd) core file is generated when a prolonged flow control failure occurs and when you configure the setting to generate a core dump during prolonged flow control (by using the `dump-on-flow-control` option with the `flow-control-options` statement). The watchdog functionality continues to generate a kernel core file in such scenarios.

- `reset-on-flow-control`—Reset interface during prolonged flow control.

NOTE: Starting in Junos OS Release 16.1R7, the `reset-on-flow-control` option has no effect on the MS-MIC, MS-MPC, MS-DPC, MS-PIC 100, MS-PIC 400, and MS-PIC 500 line cards. This is because starting in Release 16.1R7, Junos OS restarts these line cards to recover them from stuck state due to prolonged flow control.

- `up-on-flow-control`—Cause interface to remain in stuck state until you manually restart the PICs.

NOTE: Starting in Junos OS Release 16.1R7, if interfaces on an MS-PIC or MS-DPC are in stuck state because of prolonged flow control, Junos OS restarts the service PICs to recover them from this state. However, if you want the PICs to remain in stuck state until you manually restart the PICs, configure the `up-on-flow-control` option. In releases before Release 16.1R7, there is no action taken to recover service PICs from this state unless one of the options for the `flow-control-options` statement is configured, or service PIC is manually restarted.

Usage Guidelines

See ["Configuring Flow Monitoring" on page 5](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 8.4.

flow-export-destination

IN THIS SECTION

- [Syntax | 1112](#)
- [Hierarchy Level | 1112](#)
- [Description | 1112](#)
- [Options | 1112](#)
- [Required Privilege Level | 1112](#)
- [Release Information | 1112](#)

Syntax

```
flow-export-destination {  
    (cflowd-collector | collector-pic);  
}
```

Hierarchy Level

```
[edit forwarding-options monitoring group-name family inet output]
```

Description

Configure flow collection.

Options

cflowd-collector—Use the cflowd collector.

collector-pic—Use the collector PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

flow-export-rate

IN THIS SECTION

- [Syntax | 1113](#)
- [Hierarchy Level | 1113](#)
- [Description | 1113](#)
- [Options | 1113](#)
- [Required Privilege Level | 1114](#)
- [Release Information | 1114](#)

Syntax

```
flow-export-rate rate;
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family family inet output inline-jflow]
```

Description

Specify the flow export rate of monitored packets in kpps. If you have multiple line cards of different types running on the same router, the `flow-export-rate` will be applied to each card. However, the rate applied to the PFEs on the card will vary in accordance with the number of PFEs that are on the card.

On the MX Series, the actual flow export rate might be less than the configured `flow-export-rate`. In addition, the maximum flow export rate for an FPC or MPC module cannot exceed the configured `flow-export-rate`, regardless of how many PICs or MICs are on the module.

Options

rate Flow export rate of monitored packets in kpps (from 1 through 400).

- **Default:** 1 kpps (applies to all PFEs on the FPC)

NOTE: The maximum rate per PFE is 100 kpps for LU, 800 kpps for XL/EA, so for an FPC with four LU PFEs (such as AS cards) you can set a maximum `flow-export-rate` of 400. For an FPC with two LU PFEs (such as the MPC2), the maximum `flow-export-rate` is 200. For an FPC with one LU PFE (such as the MPC5), the maximum is 100. The Junos CLI accepts as valid any value within the range of 1 to 3200, but when applied the value might trigger an error message such as “The configured flow export rate is higher than supported value/chip” in the Junos message log.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

[Configuring Discard Accounting | 435](#)

[Configuring Flow Monitoring | 5](#)

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

flow-export-timer

IN THIS SECTION

- [Syntax | 1115](#)
- [Hierarchy Level | 1115](#)
- [Description | 1115](#)
- [Default | 1115](#)

- Options | 1115
- Required Privilege Level | 1115

Syntax

```
flow-export-timer seconds
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

Configure the flow export period in seconds. The software exports a IPFIX packet periodically at the configured flow-export timer interval. After you configure the flow export period, you must reboot the device for it to take effect.

Default

10 seconds

Options

seconds Range: 10 to 600 seconds.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

flow-inactive-timeout

IN THIS SECTION

- [Syntax | 1116](#)
- [Hierarchy Level | 1116](#)
- [Description | 1116](#)
- [Options | 1117](#)
- [Required Privilege Level | 1117](#)
- [Release Information | 1117](#)

Syntax

```
flow-inactive-timeout seconds;
```

Hierarchy Level

```
[edit forwarding-options accounting name output],  
[edit forwarding-options monitoring name output],  
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls | vpls) output],  
[edit forwarding-options sampling family (inet |inet6 |mpls) output],  
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name],
```

Description

Set the interval of inactivity that marks a flow inactive.

NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.

Options

seconds—Duration of the timeout period.

- **Range:** 15 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)
- **Default:** 60 seconds (for forwarding-options configurations); 60 seconds (for services configurations)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit services flow-monitoring version-ipfix template *template-name*] hierarchy level added in Junos OS Release 10.2.

Support at the [edit services flow-monitoring version9 template *template-name*] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

[Configuring Flow Monitoring | 5](#)

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

flow-key (Flow Monitoring)

IN THIS SECTION

● [Syntax | 1118](#)

- [Hierarchy Level | 1118](#)
- [Description | 1118](#)
- [Options | 1118](#)
- [Required Privilege Level | 1119](#)
- [Release Information | 1119](#)

Syntax

```
flow-key {
    flow-direction;
    vlan-id;
    output-interface;
}
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Include VLAN IDs in both the ingress and egress directions in the flow key, enable flow direction information in a Version 9 or IPFIX flow template, or both, and configure the output-interface for bridge or VPLS family for inline flow monitoring on the MX Series.

Options

flow-direction	Enable reporting of the direction of the flow. The field contains 0x00 (ingress) or 0x01 (egress). The flow direction field in the output record contains the invalid value 0xFF if you do not configure flow-direction.
vlan-id	Include VLAN IDs in both the ingress and egress directions in the flow key.
output-interface	Configure the output-interface field as part of flow-key for bridge or VPLS family.

NOTE: If the output-interface (OIF) is configured under flow-key while the flow-monitoring is in progress, all the existing flows (where OIF was not part of flow-key) report OIF field as zero in the next export. Therefore, in progress configuration of output-interface as part of flow-key is not recommended. In order to configure output-interface as part of flow-key, it is recommended to disable the bridge or vpls sampling and wait for the active flows to become zero.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.2.

output-interface option added in Junos OS Release 18.2R1.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

flow-monitoring

IN THIS SECTION

- [Syntax | 1120](#)
- [Hierarchy Level | 1121](#)

- Description | 1121
- Required Privilege Level | 1121
- Release Information | 1122

Syntax

```

flow-monitoring {
  version9 {
    template template-name {
      options-template-id
      template-id
      source-id
      flow-active-timeout seconds;
      flow-inactive-timeout seconds;
      flow-key {
        flow-direction;
        vlan-id;
        output-interface;
      }
      (ipv4-template | ipv6-template | mpls-template label-position [ positions ] | mpls-ipv4-
template label-position [ positions ] | mpls-ipvx-template);
      peer-as-billing-template;
      option-refresh-rate packets packets seconds seconds;
      options-template-id
      source-id
      template-id
      template-refresh-rate packets packets seconds seconds;
      tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
    }
  }
  version-ipfix {
    template template-name {
      data-record-fields {
        source-prefix-as-path count;
        destination-prefix-as-path count;
        bgp-source-standard-community count;
        bgp-destination-standard-community count;
        bgp-source-extended-community count;

```

```

        bgp-destination-extended-community count;
        bgp-source-large-community count;
        bgp-destination-large-community count;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-key {
        flow-direction;
        vlan-id;
    }
    (ipv4-template | ipv6-template | mpls-ipv4-template | mpls-ipvx-template | vpls-template);
    nexthop-learning (enable | disable);
    observation-domain-id
    option-refresh-rate packets packets seconds seconds;
    options-template-id
    template-id
    template-refresh-rate packets packets seconds seconds;
    tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
}
}
}

```

Hierarchy Level

[edit services]

Description

Specify the active monitoring properties for flow aggregation version 9 or IPFIX.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

data-record-fields option introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

flow-monitoring (Inline Monitoring Services)

IN THIS SECTION

- [Syntax | 1122](#)
- [Hierarchy Level | 1123](#)
- [Description | 1123](#)
- [Default | 1123](#)
- [Options | 1123](#)
- [Required Privilege Level | 1124](#)
- [Release Information | 1124](#)

Syntax

```
flow-monitoring {  
    counter-profile profile-identifier;  
    flow-rate kbps burst-size bytes;  
    flow-limit number;  
    sampling-profile profile-name;  
    sampling-rate bytes;
```

```
security-enable;  
}
```

Hierarchy Level

```
[edit services inline-monitoring template template-name]
```

Description

Configures optional parameters for flow-based telemetry (FBT) for the EX4100, EX4100-F, and EX4400 switches. FBT enables per-flow-level analytics, using inline monitoring services to create flows, collect them, and export them to a collector. A flow is a sequence of packets that have the same source IP, destination IP, source port, destination port, or protocol on an interface. For each flow, various parameters are collected and sent to a collector using the open standard IPFIX template to organize the flow. Once there is no active traffic for a flow, the flow is aged out after the configured inactive-timeout period (flow-inactive-timeout at the [edit services inline-monitoring template *template-name*] hierarchy level).

Default

Disabled

Options

counter-profile <i>profile-identifier</i>	<p>Specify which counters should be exported to the collector, by selecting one of the pre-defined profile names.</p> <ul style="list-style-type: none">Per_flow_6_counters: packet range counter (3 counters), time-to-live (TTL) range counter, TCP window range counter, DoS attack (2 counters)Per_flow_4_counters: flow packet range (packet size 64 to 1500 bytes) counter (1 count), TTL range counter, TCP window range counter, DoS counter
flow-rate <i>kbps</i> burst-size <i>bytes</i>	<p>Specify the meter rate for each flow, in kbps, and configure the maximum number of bytes allowed for incoming packets to burst above the flow meter rate.</p> <p>Range: 8 to 10000000 kbps (flow-rate); 512 to 256000000 bytes (burst-size)</p>
flow-limit <i>number</i>	<p>Specify the maximum number of flows allowed.</p>

Range: 0 through 32000

sampling-profile
profile-identifier

Configure one of the following sampling profiles:

- **First_N_Pkt** : Sample contains the first N packets of a flow
- **Deterministic**: Sample contains every Nth packet of a flow
- **Random**: Sample contains randomly chosen packets from a flow
- **Combo1**: Sample contains the first N packets of a flow, followed by a random packet at the configured interval
- **Combo2**: Sample contains the first N packets of a flow, followed by a Deterministic 1 packet at the configured interval

sampling-rate
bytes

Specify the rate at which packets are sampled to create flows, in bytes

Range: 1 to 65535 bytes

security-enable

Enable security analytics; specify that Denial-of-Service (DoS) attacks are reported to the collector.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

[Flow-Based Telemetry \(EX4100, EX4100-F, and EX4400 Series\)](#) | 349

flow-server

IN THIS SECTION

- [Syntax | 1125](#)
- [Hierarchy Level | 1126](#)
- [Description | 1126](#)
- [Options | 1126](#)
- [Required Privilege Level | 1126](#)
- [Release Information | 1127](#)

Syntax

```
flow-server hostname {  
    aggregation {  
        autonomous-system;  
        destination-prefix;  
        protocol-port;  
        source-destination-prefix {  
            caida-compliant;  
        }  
        source-prefix;  
    }  
    autonomous-system-type (origin | peer);  
    dscp dscp-value;  
    forwarding-class class-name;  
    (local-dump | no-local-dump);  
    port port-number;  
    source-address address;  
    version format;  
    version9 {  
        template template-name;  
    }  
}
```


Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls | vpls |
bridge) output],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls | bridge) output]
```

Description

Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility `cfddcollect`. Specify a host system to collect sampled flows using the version 9 format.

You can configure up to one version 5 and one version 8 flow format at the `[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]` hierarchy level. For the same configuration, you can specify only either version 9 flow record formats or formats using versions 5 and 8, not both types of formats.

Options

hostname—IP address—IPv4 or IPv6—or identifier of the host system (the workstation either running the `cflowd` utility or collecting traffic flows using version 9).

NOTE: Only host systems running IPv4 are supported on QFX10000 switches.

You can configure only one host system for version 9.

NOTE: IPv6 configuration for `flow-server` is supported only in Junos OS Release 12.3 and later. Note that when you configure an IPv6 address for the `flow-server` statement, you must also configure an IPv6 address for the `inline-jflow source-address` statement at the `[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls | vpls | bridge) output]` hierarchy level.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

version9 statement introduced in Junos OS Release 8.3.

Support at the following hierarchy levels introduced in Junos OS Release 18.2R1: [edit forwarding-options sampling instance instance-name family bridge],, [edit forwarding-options sampling family bridge].

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

flow-table-size

IN THIS SECTION

- [Syntax](#) | 1127
- [Hierarchy Level](#) | 1128
- [Description](#) | 1128
- [Required Privilege Level](#) | 1128
- [Release Information](#) | 1129

Syntax

```
flow-table-size {  
    ipv4-flow-table-size units;  
    ipv6-extended-attrib;  
    ipv6-flow-table-size units;  
    mpls-flow-table-size units;  
    vpls-flow-table-size units;  
}
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services]
```

Description

Configure the size of hash tables for inline services sampling.

Starting with Junos OS Release 15.1F2, by default, the software allocates one 1K IPv4 flow table. To allocate 15 256K IPv4 flow tables, the former default, you can enter this configuration from the [edit] hierarchy level:

```
[edit]
user@router# set chassis fpc inline-services flow-table-size ipv4-flow-table-size 15
```

The maximum supported flow table size for a combination of both IPv4 and IPv6 is 15. For example, you can set the flow table size for IPv4 to 10 and set the size for IPv6 to 5. Verify that you have sized the flow tables adequately for IPv4 and IPv6 flow sampling.

NOTE: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more; however, the system will issue a warning message.

NOTE: We recommend that you configure this statement during a maintenance window:

- Prior to Junos OS 16.1R1 and 15.1F2, the FPC reboots automatically after you commit this configuration change.
- Starting from Junos OS Release 16.1R1 and 15.1F2, any changes in the configured size of the flow table do not require a reboot of the FPC.

The remaining statements are defined separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

ipv6-extended-attrib option added in Junos OS Release 14.2 for MX Series routers.

vpls-flow-table-size option added in Junos OS Release 13.2 for MX Series routers.

bridge-flow-table-size option added in Junos OS Release 18.2R1 for MX Series routers.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

[Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers | 633](#)

flow-table-size (Chassis)

IN THIS SECTION

- [Syntax | 1129](#)
- [Hierarchy Level | 1130](#)
- [Description | 1130](#)
- [Options | 1130](#)
- [Required Privilege Level | 1130](#)
- [Release Information | 1130](#)

Syntax

```
flow-table-size size;
```

Hierarchy Level

```
[edit chassis fpc slot inline-video-monitoring]
```

Description

Configure the number of video flows that can be measured per Packet Forwarding Engine by an MPC at a given time. This value takes effect the next time the MPC is rebooted.

Options

size Number of video flows per Packet Forwarding Engine.

- **Range:** 16 through 8192

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

| [Configuring Inline Video Monitoring on MX Series Routers](#) | 934

flow-tap

IN THIS SECTION

● [Syntax](#) | 1131

- [Hierarchy Level | 1131](#)
- [Description | 1131](#)
- [Options | 1131](#)
- [Required Privilege Level | 1132](#)
- [Release Information | 1132](#)

Syntax

```
flow-tap {
    (interface interface-name | tunnel-interface interface-name);
    family (inet | inet6);
}
```

Hierarchy Level

```
[edit services]
```

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 <i>interface-name</i>	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. <ul style="list-style-type: none">• 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 Information

Statement introduced in Junos OS Release 8.1.

ccc option introduced in Junos OS Release 17.2.

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](#)

forwarding-class (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1133](#)
- [Hierarchy Level | 1133](#)

- [Description | 1133](#)
- [Options | 1133](#)
- [Required Privilege Level | 1133](#)
- [Release Information | 1133](#)

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the forwarding class to be used for test frames. The forwarding class specifies the manner in which the test frames are processed by the Packet Forwarding Engine of the router. If you do not configure this parameter, test frames are treated as best-effort traffic.

Options

class-name Name of the forwarding class. You must have previously configured this forwarding class by including the forwarding-class statement at the [edit class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

forwarding-class (Sampling)

IN THIS SECTION

- [Syntax | 1134](#)
- [Hierarchy Level | 1134](#)
- [Description | 1134](#)
- [Default | 1135](#)
- [Options | 1135](#)
- [Required Privilege Level | 1135](#)
- [Release Information | 1135](#)

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet | inet6) output flow-server hostname]
```

Description

Specify the forwarding class to which exported packets for inline active flow monitoring are sent.

Default

If you do not include the *forwarding-class* statement, exported packets are sent to the best effort queue.

Options

forwarding-class *class-name*

Name of the forwarding class:

- assured-forwarding
- best-effort
- expedited-forwarding
- network-control

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches](#) | 433

ftp (Flow Collector Files)

IN THIS SECTION

- [Syntax](#) | 1136
- [Hierarchy Level](#) | 1136
- [Description](#) | 1136

- Options | 1136
- Required Privilege Level | 1137
- Release Information | 1137

Syntax

```
ftp: url;
```

Hierarchy Level

```
[edit services flow-collector destination]
```

Description

Specify the primary and secondary destination FTP server addresses.

Options

url—FTP server address. The URL can include the following macros, typed in braces:

- {%D}—Date
- {%T}—Time when the file is created
- {%I}—Description string for the logical interface configured using the collector *interface-name* statement at the [edit services flow-collector interface-map] hierarchy
- {%N}—Unique, sequential number for each new file created
- {am_pm}—AM or PM
- {date}—Current date using the {year} {month} {day} macros
- {day}—From 01 through 31
- {day_abbrev}—Sun through Sat
- {day_full}—Sunday through Saturday

- {generation number}—Unique, sequential number for each new file created
- {hour_12}—From 01 through 12
- {hour_24}—From 00 through 23
- {ifalias}—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy
- {minute}—From 00 through 59
- {month}—From 01 through 12
- {month_abbr}—Jan through Dec
- {month_full}—January through December
- {num_zone}—From -2359 to +2359; this macro is not supported
- {second}—From 00 through 60
- {time}—Time the file is created, using the {hour_24} {minute} {second} macros
- {time_zone}—Time zone code name of the locale; for example, gmt (this macro is not supported).
- {year}—In the format YYYY; for example, 1970
- {year_abbr}—From 00 through 99

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

ftp (Transfer Log Files)

IN THIS SECTION

- [Syntax | 1138](#)
- [Hierarchy Level | 1138](#)
- [Description | 1138](#)
- [Options | 1138](#)
- [Required Privilege Level | 1138](#)
- [Release Information | 1139](#)

Syntax

```
ftp: url;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive archive-sites]
```

Description

Specify the primary and secondary destination FTP server addresses.

Options

url—FTP server address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

g-duplicates-dropped-periodicity

IN THIS SECTION

- [Syntax](#) | [1139](#)
- [Hierarchy Level](#) | [1139](#)
- [Description](#) | [1140](#)
- [Default](#) | [1140](#)
- [Options](#) | [1140](#)
- [Required Privilege Level](#) | [1140](#)
- [Release Information](#) | [1140](#)

Syntax

```
g-duplicates-dropped-periodicity seconds;
```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Specify the frequency for sending notifications to affected control sources when transmission of duplicate sets of data is restricted because the `g-max-duplicates` threshold has been reached. This setting is applied globally; the `duplicates-dropped-periodicity` setting applied at the capture-group level overrides the global setting.

Default

The default period for sending notifications is 30 seconds.

Options

seconds—Period for sending DuplicatesDropped notifications.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[duplicates-dropped-periodicity](#) | 1071

[Configuring Junos Capture Vision](#) | 289

g-max-duplicates

IN THIS SECTION

● [Syntax](#) | 1141

● [Hierarchy Level](#) | 1141

- [Description | 1141](#)
- [Default | 1141](#)
- [Options | 1141](#)
- [Required Privilege Level | 1141](#)
- [Release Information | 1142](#)

Syntax

```
g-max-duplicates number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Specify the maximum number of content destinations to which DFC PICs can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting is applied globally; the `max-duplicates` setting applied at the `capture-group` level overrides the global setting.

Default

If no value is configured, a default setting of 3 is used.

Options

number—Maximum number of content destinations.

- **Range:** 1 through 64

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[max-duplicates](#) | [1216](#)

[Configuring Junos Capture Vision](#) | [289](#)

generate-snmp-traps

IN THIS SECTION

- [Syntax](#) | [1142](#)
- [Hierarchy Level](#) | [1142](#)
- [Description](#) | [1143](#)
- [Required Privilege Level](#) | [1143](#)
- [Release Information](#) | [1143](#)

Syntax

```
generate-snmp-traps;
```

Hierarchy Level

```
[edit services]  
[edit services video-monitoring]
```

Description

If this statement is configured, the service generates SNMP traps for severity levels such as Info, Warning, Critical, or Cleared. For example, if DF alarm changes from info to warning, or from warning to critical, mdiDFAAlarm trap be triggered.

NOTE: SNMP traps are not generated if SNMP trap generation is not enabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers | 927](#)
[alarms | 973](#)

halt-on-prefix-down (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1144](#)
- [Hierarchy Level | 1144](#)
- [Description | 1144](#)
- [Required Privilege Level | 1144](#)
- [Release Information | 1144](#)

Syntax

```
halt-on-prefix-down;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

By default, the RFC 2544-based benchmarking test ignores a prefix-down event (when the prefix associated with the test goes down) and continues to run. If this parameter is specified, a prefix that moves to the down state causes the corresponding tests to be stopped. The `show` command output for the test displays that the test was terminated due to the prefix going down.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[rfc2544-benchmarking](#) | 1347

hard-limit

IN THIS SECTION

- [Syntax | 1145](#)
- [Hierarchy Level | 1145](#)
- [Description | 1145](#)
- [Options | 1145](#)
- [Required Privilege Level | 1145](#)
- [Release Information | 1146](#)

Syntax

```
hard-limit bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which the dynamic flow capture application begins deleting criteria, until the bandwidth falls below the `hard-limit-target` value.

Options

bandwidth—Hard limit threshold, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[hard-limit-target | 1146](#)

[Configuring Junos Capture Vision | 289](#)

hard-limit-target

IN THIS SECTION

- [Syntax | 1146](#)
- [Hierarchy Level | 1146](#)
- [Description | 1146](#)
- [Options | 1147](#)
- [Required Privilege Level | 1147](#)
- [Release Information | 1147](#)

Syntax

```
hard-limit-target bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which the dynamic flow capture application stops deleting criteria.

Options

bandwidth—Target value, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[hard-limit | 1145](#)

[Configuring Junos Capture Vision | 289](#)

hardware-timestamp

IN THIS SECTION

- [Syntax | 1147](#)
- [Hierarchy Level | 1148](#)
- [Description | 1148](#)
- [Required Privilege Level | 1148](#)
- [Release Information | 1148](#)

Syntax

```
hardware-timestamp;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with icmp-ping, icmp-ping-timestamp, udp-ping, and udp-ping-timestamp probe types.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

Statement applied to MX Series routers in Junos OS Release 10.0.

Statement introduced in Junos OS Release 19.1 for PTX Series routers.

history-size

IN THIS SECTION

- [Syntax | 1149](#)
- [Junos OS Hierarchy Levels | 1149](#)
- [Junos OS Evolved Hierarchy Level | 1149](#)
- [Description | 1149](#)
- [Options | 1149](#)
- [Required Privilege Level | 1149](#)
- [Release Information | 1149](#)

Syntax

```
history-size size;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the number of stored history entries.

Options

size—Value from 0 to 255.

- **Default:** 50

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

host-outbound media-interface

IN THIS SECTION

- [Syntax | 1150](#)
- [Hierarchy Level | 1150](#)
- [Description | 1150](#)
- [Required Privilege Level | 1151](#)
- [Release Information | 1151](#)

Syntax

```
host-outbound media-interface;
```

Hierarchy Level

```
[edit chassis]
```

Description

Enable Layer 2 port mirroring of host-generated outbound packets only on MPCs on MX Series 5G Universal Routing Platforms.

This statement enables all Routing Engine-generated Layer 2 injections to execute egress logical interface filters.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

Examples: Layer 2 Port Mirroring at Multiple Levels of the Chassis

Configuring Port Mirroring

Understanding Layer 2 Port Mirroring

icmp

IN THIS SECTION

- [Junos OS | 1151](#)
- [Junos OS Evolved | 1152](#)
- [Junos OS Hierarchy Level | 1152](#)
- [Junos OS Evolved Hierarchy Level | 1152](#)
- [Description | 1152](#)
- [Required Privilege Level | 1152](#)
- [Release Information | 1152](#)

Junos OS

```
icmp {  
    destination-interface interface-name;  
}
```

Junos OS Evolved

```
icmp;
```

Junos OS Hierarchy Level

```
[edit services rpm probe-server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm probe-server]
```

Description

(Required for Junos OS Evolved and for J Series routers running Junos OS) Enable ICMP requests for the RPM probe server. ICMP requests are enabled by default on Junos OS (except for J Series routers); you do not need to explicitly configure them.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

NOTE: The destination-interface statement is not supported on PTX Series routers or for Junos OS Evolved.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#) | [646](#)

in-service

IN THIS SECTION

- [Syntax](#) | [1153](#)
- [Junos OS Hierarchy Level](#) | [1153](#)
- [Junos OS Evolved Hierarchy Level](#) | [1153](#)
- [Description](#) | [1154](#)
- [Default](#) | [1154](#)
- [Required Privilege Level](#) | [1154](#)
- [Release Information](#) | [1154](#)

Syntax

```
in-service;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Runs the test in the in-service mode. In this mode, while the test is running, the rest of the data traffic sent to and from the UNI port under test on the service are not interrupted. Control protocol packets and control protocol peering are not interrupted.

If this mode is not configured, the test runs in the default out-of-service mode. In the out-of-service mode, while the test is running, all the data traffic sent to and from the UNI port under test on the service is interrupted. Control protocol peering is not interrupted whereas control protocol packets such as CFM sessions are interrupted.

Default

The default service mode for the reflecting egress interface for an E-LAN service is out-of-service mode.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[rfc2544-benchmarking](#) | [1347](#)

[rfc2544](#) | [1345](#)

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | [855](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices](#) | [728](#)

[Configuring an RFC 2544-Based Benchmarking Test](#) | [739](#)

inactivity-timeout (Services RPM)

IN THIS SECTION

- [Syntax | 1155](#)
- [Hierarchy Level | 1155](#)
- [Description | 1155](#)
- [Options | 1155](#)
- [Required Privilege Level | 1155](#)
- [Release Information | 1156](#)

Syntax

```
inactivity-timeout seconds;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Inactivity timeout period, in seconds.

Options

seconds—Length of time the session is inactive before it times out.

- **Default:** 1800 seconds

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | [694](#)

inet6-options (Services)

IN THIS SECTION

- [Syntax](#) | [1156](#)
- [Hierarchy Level](#) | [1156](#)
- [Description](#) | [1157](#)
- [Options](#) | [1157](#)
- [Required Privilege Level](#) | [1157](#)
- [Release Information](#) | [1157](#)

Syntax

```
inet6-options {  
    source-address address;  
}
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Specify the source IPv6 address used for probes. If the source IPv6 address is not one of the devices' assigned addresses, the packet uses the outgoing interface's address as its source.

Options

`inet6-options`—Use the specified base IPv6 protocol-related settings to be used for RPM probes

`source-address ipv6-address`—Specify the base IPv6 address for sending the RPM probes from the client to the server (for example, 2001:db8::a:b:c:d).

Required Privilege Level

`system`—To view this statement in the configuration.

`system-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R4.

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | [651](#)

inband-flow-telemetry

IN THIS SECTION

- [Syntax](#) | [1158](#)
- [Hierarchy Level](#) | [1158](#)
- [Description](#) | [1158](#)
- [Options](#) | [1159](#)
- [Required Privilege Level](#) | [1161](#)
- [Release Information](#) | [1161](#)

Syntax

```

inband-flow-telemetry {
  clock-source (ntp|ptp);
  device-id {
    (id-number|auto);
  }
  flow-type (l3|vxlan);
  hop-limit value;
  meta-data-stack-length value;
  no-ipv6-address-match;
  profile {
    ifa-profile-name {
      sample-rate value;
      collector {
        source-address ipv4-address;
        destination-address ipv4-address;
        destination-port port-number;
        maximum-clip-length length;
        mtu size;
      }
    }
  }
}

```

Hierarchy Level

[edit services]

Description

Configure Inband Flow Analyzer 2.0 (IFA 2.0). IFA 2.0 provides insights about complex networks by collecting per-hop flow data on the data plane. IFA uses probe packets to collect network-wide flow data. IFA generates probe packets by sampling the traffic flow of interest. IFA probe packets are representative packets of the original flow and possess the exact same characteristics as the original flow. This means that IFA packets traverse the same path in the network and the same queues in the networking element as the original packet would. Because the IFA probe packets traverse the same network path as the original flow, the packets experience similar latency and congestion.

IFA uses the following processing nodes to monitor and analyze flows:

- IFA initiator node (also known as ingress node)
- IFA transit node
- IFA terminating node (also known as egress node)

Use the `inband-flow-telemetry` configuration options to configure the IFA nodes.

Options

clock-source (ntp|ptp) Configure the clock source protocol to enable more accurate timestamping. The QFX5120-48YM switch supports PTP as well as NTP; all of the other QFX5120 switch models support NTP only.

Default: ntp

device-id (id-number|auto) (Mandatory for all IFA nodes) Specify a unique device identifier for each hop within an IFA zone. You must configure this value for all three IFA node types: IFA initiator, IFA transit, and IFA terminating. If you configure `auto`, then the device ID is internally generated from the router ID or the management IP address.

Range: 1-1,048,575

flow-type (l3|vxlan) (Mandatory for IFA initiator node and terminating node and optional for transit node) Specify the IFA flow type—`l3` or `vxlan`. If you configure the flow type as `vxlan`, and the incoming traffic is L3, or vice versa, then the IFA nodes do not behave as expected. This is because you cannot initiate IFA probe packets with invalid fields.

You cannot configure both L3 or VXLAN flows on the same device. This restriction is applicable for the IFA initiator and terminating nodes (generally leaf nodes).

You don't need to configure `flow-type` for a transit node (generally spine nodes).

Default: l3

hop-limit value (Optional) Configure the maximum allowed hops in an IFA zone. The initiator node initializes this field. The hop limit is decremented at each hop. If the hop limit of the incoming packets is 0, the current node does not insert the metadata.

You can avoid the metadata insertion at the transit node by using the `hop-limit` configuration.

The IFA terminating node does not perform a hop-limit check. Even if the incoming IFA packet has hop-limit set to 0, the IFA terminating node inserts the metadata and reduces the hop limit by 1. In this case, the `hop-limit` value resets to 255. The `hop-limit` option cannot have a negative value.

- **Range:** 1-250

- **Default:** 250

meta-data-stack-length value (Optional) Configure the maximum allowed length of the metadata stack in multiples of four octets. The initiator node initializes this field. Each node in the path compares the current length with the maximum allowed length. If the current length equals or exceeds the maximum length, the transit node must stop inserting the metadata.

- **Range:** 8-255
- **Default:** 240 (for 30 hops)

no-ipv6-address-match (Optional) Optimize IFA filter group processing by removing the IPv6 source and destination address match qualifiers from the IFA filter group. IFA cannot be initiated or terminated with these two qualifiers, but you can initiate or terminate IFA with the remaining qualifiers.

Default: off (this statement is not included in the IFA configuration unless you specifically configure it)

profile ifa-profile-name IFA profile name.

sample-rate value Configure the average number of samples obtained in one second. For example, if you configure the sample rate as 1000, then out of 1000 packets one packet is sampled per second. You cannot have different sample rates for different flows on an IFA initiator node enabled on a port. All flows within a port must have the same sample rate.

Range: 1-16,777,215

collector Configure a collector for IFA 2.0 probe packets. The monitored packets are exported to the collector in IPFIX format. By default, Junos OS supports a maximum packet length of 256 bytes starting with the Ethernet header. An IFA IPFIX packet contains IFA headers (8 bytes), IFA metadata (variable length), and the originally monitored packet (256 bytes).

Configure the following collector-related options:

- **source-address** *ipv4-address*—IPv4 source address.
- **destination-address** *ipv4-address*—IPv4 destination address.
- **destination-port** *port-number*—Destination port value.

Range: 1-65,535.

- `maximum-clip-length length`—Number of bytes of the original flow packet that should be exported in the IPFIX packet. Because the maximum MTU is 9000 bytes at the IFA termination node, the maximum clip length for the IPFIX packet is equal to or less than: 9000 bytes - (IFA header length + IFA metadata header length + IFA metadata stack length).

Range: 64 to 9000 bytes

Default: 256 bytes

- `mtu size`—Size in bytes of the maximum transmission unit for IPFIX packets leaving the IFA termination node. Any packet exceeding 9000 bytes in length is dropped.

Range: 256 to 9000 bytes

Default: 9000 bytes

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

`clock-source`, `maximum-clip-length`, `mtu`, and `no-ipv6-address-match` options introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

[inband-flow-telemetry](#) | 1157

[clear inband-flow-telemetry stats](#) | 1560

[show services inband-flow-telemetry](#) | 1701

[Inband Flow Analyzer \(IFA\) 2.0 Probe for Real-Time Flow Monitoring](#) | 370

inline-jflow

IN THIS SECTION

- [Syntax | 1162](#)
- [Hierarchy Level | 1162](#)
- [Description | 1162](#)
- [Required Privilege Level | 1162](#)
- [Release Information | 1163](#)

Syntax

```
inline-jflow {  
    source-address address;  
    flow-export-rate rate;  
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family inet output]
```

Description

Specify inline flow monitoring for traffic from the designated address.

The remaining statements are explained separately. See [CLI Explorer](#).

NOTE: If you configure inline flow monitoring with `inline-jflow` then you have to disable it before performing ISSU. For more information, see [Before You Begin a Unified ISSU](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#) | 74

inline-monitoring

IN THIS SECTION

- [Syntax](#) | 1163
- [Hierarchy Level](#) | 1164
- [Description](#) | 1164
- [Options](#) | 1164
- [Required Privilege Level](#) | 1164
- [Release Information](#) | 1165

Syntax

```
inline-monitoring {  
  instance instance-name;  
  observation-cloud-id observation-cloud-identifier;  
  template template-name;  
  traceoptions;  
}
```

Hierarchy Level

[edit services]

Description

Configure inline monitoring services. When you enable inline monitoring, you can monitor actual IPv4 and IPv6 packets at different sampling rates, and export the actual packet up to the configured clip length. By default, Junos OS supports a maximum packet length of 126 bytes starting with the Ethernet header. The monitored packets are exported to an collector for further processing. The packets are exported in an IPFIX format, which includes information on the original packet size, and incoming or outgoing interface.

Options

instance <i>instance-name</i>	Configure parameters for an inline-monitoring instance. See "instance" on page 1165 for more information.
observation-cloud-id <i>observation-cloud-identifier</i>	<p>Observation cloud ID—an identifier for a particular observation cloud. An observation cloud is the largest set of observation domains. Per RFC 5101, an observation domain is the largest set of observation points for which flow information can be aggregated by a metering process. For example, a router line card may be an observation domain if it is composed of several interfaces, each of which is an observation point. The observation domain ID is unique per exporting process. Also per RFC 5101, an observation point is a location in the network where IP packets can be observed. Examples include: a line to which a probe is attached, a shared medium, such as an Ethernet-based LAN, a single port of a router, or a set of interfaces (physical or logical) of a router.</p> <ul style="list-style-type: none"> • Range: 1 through 255
template <i>template-name</i>	Configure templates for inline packet monitoring. See "template" on page 1442 for more information.
traceoptions	(Optional) Configure traceoptions for the inline monitoring process. See "traceoptions" on page 1493 for more information.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

observation-cloud-id option introduced in Junos OS Release 21.2R1.

Statement introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

[Inline Monitoring Services Configuration | 334](#)

[Flow-Based Telemetry \(EX4100, EX4100-F, and EX4400 Series\) | 349](#)

instance

IN THIS SECTION

- [Syntax \(Junos OS\) | 1165](#)
- [Syntax \(Junos OS Evolved\) | 1166](#)
- [Hierarchy Level | 1166](#)
- [Description | 1166](#)
- [Options | 1166](#)
- [Required Privilege Level | 1167](#)
- [Release Information | 1167](#)

Syntax (Junos OS)

```
instance name {
  collector name;
  maximum-clip-length maximum-clip-length;
```



```

    template-name template-name;
}

```

Syntax (Junos OS Evolved)

```

instance name {
    collector name;
    controller name;
    maximum-clip-length maximum-clip-length;
    sampling-rate sampling-rate;
    template-name template-name;
}

```

Hierarchy Level

```
[edit services inline-monitoring]
```

Description

Configure inline-monitoring instance parameters. You can use these instances along with firewall filters to monitor different streams of traffic at different sampling rates from the same interface.

You can configure a maximum of sixteen (Junos OS) or seven (Junos OS Evolved) inline-monitoring instances with four collectors each.

Options

name	Name of instance.
collector <i>name</i>	Configure an collector for the inline-monitoring instance. See "collector" on page 1006 for more information.
controller	(Junos OS Evolved only) Configure inline-monitoring services for packet redirects to the P4 controller or to the Routing Engine. See "controller" on page 1025 for more information.
maximum-clip-length	Maximum packet length.

- **Range:** 64 through 126 bytes (Junos OS)
Range: 64 through 256 bytes (Junos OS Evolved)
- **Default:** 126 bytes (Junos OS)
Default: 128 bytes (Junos OS Evolved)

sampling-rate (Junos OS Evolved) Rate at which the traffic is sampled. In Junos OS, you do not specify the sampling rate here; you specify the sampling rate at the [edit services inline-monitoring instance *instance-name* collector *collector-name*] hierarchy level.

For example, if you specify 1000, then 1 packet out of every 1000 packets is sampled.

- **Range:** 1 through 32000000
- **Default:** 1

template-name Name of the template.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 19.4R1.

Statement introduced in Junos Evolved OS Release 22.2R1.

RELATED DOCUMENTATION

| [Understanding Inline Monitoring Services](#) | 334

input (Sampling)

IN THIS SECTION

- [Syntax | 1168](#)
- [Hierarchy Level | 1168](#)
- [Description | 1168](#)
- [Required Privilege Level | 1168](#)
- [Release Information | 1169](#)

Syntax

```
input {  
    max-packets-per-second number;  
    rate number;  
    run-length number;  
    maximum-packet-length bytes;  
}
```

Hierarchy Level

```
[edit forwarding-options sampling],  
[edit forwarding-options sampling instance instance-name]
```

Description

Configure traffic sampling on a logical interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

input-interface-index

IN THIS SECTION

- [Syntax](#) | 1169
- [Hierarchy Level](#) | 1169
- [Description](#) | 1169
- [Options](#) | 1170
- [Required Privilege Level](#) | 1170
- [Release Information](#) | 1170

Syntax

```
input-interface-index number;
```

Hierarchy Level

```
[edit forwarding-options monitoring name output interface interface-name]
```

Description

Specify a value for the input interface index that overrides the default supplied by SNMP.

Options

number—Input interface index value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

input-packet-rate-threshold

IN THIS SECTION

- [Syntax](#) | 1170
- [Hierarchy Level](#) | 1171
- [Description](#) | 1171
- [Options](#) | 1171
- [Required Privilege Level](#) | 1171
- [Release Information](#) | 1171

Syntax

```
input-packet-rate-threshold rate;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify a packet rate threshold value that triggers a system log warning message.

Options

rate—Threshold value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

instance (Sampling)

IN THIS SECTION

- [Syntax](#) | 1172
- [Hierarchy Level](#) | 1173
- [Description](#) | 1173
- [Required Privilege Level](#) | 1173

Syntax

```

instance instance-name {
    disable;
    family (bridge | inet | inet6 | mpls | vpls) {
        disable;
        output {
            aggregate-export-interval seconds;
            flow-active-timeout seconds;
            flow-inactive-timeout seconds;
            extension-service service-name;
            flow-server hostname {
                aggregation {
                    autonomous-system;
                    destination-prefix;
                    protocol-port;
                    source-destination-prefix {
                        caida-compliant;
                    }
                    source-prefix;
                }
                autonomous-system-type (origin | peer);
                dscp dscp-value;
                forwarding-class class-name;
                (local-dump | no-local-dump);
                port port-number;
                source-address address;
                version format;
                version9 {
                    template template-name;
                }
                version-ipfix {
                    template template-name;
                }
            }
        }
        interface interface-name {
            engine-id number;

```

```

        engine-type number;
        source-address address;
    }
    inline-jflow {
        source-address address;
        flow-export-rate rate;
    }
}
}
input {
    rate number;
    run-length number;
    max-packets-per-second number;
    maximum-packet-length bytes;
}
}

```

Hierarchy Level

```
[edit forwarding-options sampling]
```

Description

Configure a sampling instance.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.

RELATED DOCUMENTATION

[Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches](#) | 433

interface (Accounting or Sampling)

IN THIS SECTION

- [Syntax](#) | 1174
- [Hierarchy Level](#) | 1174
- [Description](#) | 1174
- [Options](#) | 1175
- [Required Privilege Level](#) | 1175
- [Release Information](#) | 1175

Syntax

```
interface interface-name {  
    engine-id number;  
    engine-type number;  
    source-address address;  
}
```

Hierarchy Level

```
[edit forwarding-options accounting name output],  
[edit forwarding-options sampling family (inet | inet6 | mpls) output],  
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output]
```

Description

Specify the output interface for monitored traffic.

Options

interface-name—Name of the interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Discard Accounting | 435](#)

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic

interfaces

IN THIS SECTION

- [Syntax | 1176](#)
- [Hierarchy Level | 1176](#)
- [Description | 1176](#)
- [Default | 1176](#)
- [Required Privilege Level | 1176](#)
- [Release Information | 1176](#)

Syntax

```
interfaces { ... }
```

Hierarchy Level

```
[edit]
```

Description

Configure interfaces on the router.

Default

The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Junos OS Network Interfaces Library for Routing Devices](#)

interface (Services Flow Tap)

IN THIS SECTION

- [Syntax | 1177](#)
- [Hierarchy Level | 1177](#)
- [Description | 1177](#)
- [Options | 1177](#)
- [Required Privilege Level | 1177](#)
- [Release Information | 1178](#)

Syntax

```
interface sp-fpc/pic/port.logical-unit-number;
```

Hierarchy Level

```
[edit services flow-tap]
```

Description

Specify the AS PIC interface used with the flow-tap application. Any AS PIC available in the router can be assigned, and any logical interface on the AS PIC can be used.

Options

sp-fpc/pic/port.logical-unit-number Use the specified services interface for flow-tap service.

You cannot configure flow-tap services on channelized interfaces.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

[Configuring Junos Packet Vision on MX, M and T Series Routers](#) | 303

interface-map

IN THIS SECTION

- [Syntax](#) | 1178
- [Hierarchy Level](#) | 1179
- [Description](#) | 1179
- [Required Privilege Level](#) | 1179
- [Release Information](#) | 1179

Syntax

```
interface-map {  
    collector interface-name;  
    file-specification variant-number;  
    interface-name {  
        collector interface-name;  
        file-specification variant-number;  
    }  
}
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Match an input interface with a flow collector interface and apply the preset file specifications to the input interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

interfaces (Services Dynamic Flow Capture)

IN THIS SECTION

- [Syntax](#) | 1180
- [Hierarchy Level](#) | 1180
- [Description](#) | 1180
- [Options](#) | 1180
- [Required Privilege Level](#) | 1180

Syntax

```
interfaces interface-name;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify the DFC interface used with the control source configured in the same capture group.

Options

interface-name—Name of the DFC interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision | 289](#)

interfaces (Video Monitoring)

IN THIS SECTION

- [Syntax | 1181](#)
- [Hierarchy Level | 1183](#)
- [Description | 1183](#)
- [Options | 1183](#)
- [Required Privilege Level | 1184](#)
- [Release Information | 1184](#)

Syntax

```

interfaces {
  interface-name {
    family {
      inet {
        input-flows {
          input-flow-name {
            destination-address [ address ];
            destination-port [ port ];
            source-address [ address ];
            source-port [ port ];
            template template-name;
          }
        }
        output-flows {
          output-flow-name {
            destination-address [ address ];
            destination-port [ port ];
            source-address [ address ];
            source-port [ port ];
            template template-name;
          }
        }
      }
    }
  }
  inet6 {

```



```

        input-flows {
            input-flow-name {
                destination-address [ address ];
                destination-port [ port ];
                source-address [ address ];
                source-port [ port ];
                template template-name;
            }
        }
        output-flows {
            output-flow-name {
                destination-address [ address ];
                destination-port [ port ];
                source-address [ address ];
                source-port [ port ];
                template template-name;
            }
        }
    }
    mpls {
        input-flows {
            input-flow-name {
                (destination-address [ address ] | source-address [ address ]);
                destination-port [ port ];
                payload-type (ipv4 | ipv6);
                source-port [ port ];
                template template-name;
            }
        }
        output-flows {
            output-flow-name {
                (destination-address [ address ] | source-address [ address ]);
                destination-port [ port ];
                payload-type (ipv4 | ipv6);
                source-port [ port ];
                template template-name;
            }
        }
    }
}

```

Hierarchy Level

```
[edit services video-monitoring]
```

Description

Define video monitoring for specified input or output flows on selected interfaces. You can configure a maximum of 256 flows for an interface.

Options

destination-address *address* Destination IPv4 or IPv6 address or prefix value of a flow that you want to monitor. You can use up to 32 addresses.

For IPv4-over-MPLS flows, if you configure multiple addresses for both the destination and source, then either all the destination or all the source values must have the same prefix length. For example, the following is allowed, because all the destination addresses have the same prefix length.

```
[edit services video-monitoring interfaces ge-0/2/2.0 family mpls]
user@host# set input-flows input-flow-name destination-address
[203.0.13.0/24 198.51.100.0/24]
user@host# set input-flows input-flow-name source-address [172.16.0.0/12
192.0.2.11/32]
```

For IPv6-over-MPLS flows, if you configure both the destination and source address, you can use multiple addresses for either the destination or the source IP address, but not for both.

destination-port *port* Destination port number of a flow that you want to monitor. You can use multiple port numbers and port ranges.

- **Range:** 0 through 65,535

input-flows *input-flow-name* Name of an input flow you are defining.

interface-name Name of the interface to monitor.

output-flows *output-flow-name* Name of an output flow you are defining.

payload-type ipv4	Monitor video stream for IPv4-over-MPLS traffic.
payload-type ipv6	Monitor video stream for IPv6-over-MPLS traffic.
source-address address	Source IPv4 or IPv6 address or prefix value of a flow that you want to monitor. You can use up to 32 addresses.

For IPv4-over-MPLS flows, if you configure multiple addresses for both the destination and source, then either all the destination or all the source values must have the same prefix length. For example, the following is allowed, because all the destination addresses have the same prefix length.

```
[edit services video-monitoring interfaces ge-0/2/2.0 family mpls]
user@host# set input-flows input-flow-name destination-address
[203.0.13.0/24 198.51.100.0/24]
user@host# set input-flows input-flow-name source-address [172.16.0.0/12
192.0.2.11/32]
```

For IPv6-over-MPLS flows, if you configure both the destination and source address, you can use multiple addresses for either the destination or the source IP address, but not for both.

source-port port	Source port number of a flow that you want to monitor. You can use multiple port numbers and port ranges.
-------------------------	---

- **Range:** 0 through 65,535

template-name	Name of the template used to monitor the input flows or output flows on an interface. The template contains the measurement parameters for video monitoring, and is configured at the [edit services video-monitoring templates] hierarchy level.
----------------------	---

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

mpls option introduced in Junos OS Release 17.2.

payload-type ipv6 option introduced in Junos OS Release 17.4.

RELATED DOCUMENTATION

| [Configuring Inline Video Monitoring on MX Series Routers](#) | 934

inet6-options (Services)

IN THIS SECTION

- [Syntax](#) | 1185
- [Hierarchy Level](#) | 1185
- [Description](#) | 1185
- [Options](#) | 1186
- [Required Privilege Level](#) | 1186
- [Release Information](#) | 1186

Syntax

```
inet6-options {  
    source-address address;  
}
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Specify the source IPv6 address used for probes. If the source IPv6 address is not one of the devices' assigned addresses, the packet uses the outgoing interface's address as its source.

Options

`inet6-options`—Use the specified base IPv6 protocol-related settings to be used for RPM probes

`source-address ipv6-address`—Specify the base IPv6 address for sending the RPM probes from the client to the server (for example, 2001:db8::a:b:c:d).

Required Privilege Level

`system`—To view this statement in the configuration.

`system-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1R4.

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | 651

ipfix-sw-mode

IN THIS SECTION

- [Syntax](#) | 1186
- [Hierarchy Level](#) | 1187
- [Description](#) | 1187
- [Required Privilege Level](#) | 1187
- [Release Information](#) | 1187

Syntax

```
ipfix-sw-mode;
```

Hierarchy Level

```
[edit system packet-forwarding-options]
```

Description

(QFX5120 only) Enable software-based IPFIX, used to enable flow-based telemetry for VXLANs. You must reboot the device after you commit this statement to the configuration.

Required Privilege Level

admin—To view this statement in the configuration.

admin-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

[Flow-Based Telemetry for VXLANs \(QFX5120\) | 361](#)

ip-swap

IN THIS SECTION

- [Syntax | 1188](#)
- [Junos OS Hierarchy Level | 1188](#)
- [Junos OS Evolved Hierarchy Level | 1188](#)
- [Description | 1188](#)
- [Required Privilege Level | 1188](#)
- [Release Information | 1188](#)

Syntax

```
ip-swap;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Swaps source and destination IPv4 addresses. This statement is applicable only for family bridge or family ethernet-switching.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[rfc2544-benchmarking | 1347](#)

[rfc2544 | 1345](#)

ipv4-flow-table-size

IN THIS SECTION

- [Syntax | 1189](#)
- [Hierarchy Level | 1189](#)
- [Description | 1189](#)
- [Options | 1190](#)
- [Required Privilege Level | 1190](#)

Syntax

```
ipv4-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services flow-table-size]
```

Description

Configure the size of the IPv4 flow table in units of 256K entries.

NOTE: Prior to Junos OS Release 16.1R1 and 15.1F2, any changes in the configured size of the flow table initiates an automatic reboot of the FPC, and we recommend that you run this command in a maintenance window.

NOTE: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more, however, the system will issue a warning message.

Starting with Junos OS Release 16.1R1 and 15.1F2, by default, the software allocates 1K entries for IPv4 flow tables. To allocate fifteen 256K IPv4 flow tables, the former default, you can enter this configuration from the [edit] hierarchy level:

```
[edit]
user@router# set chassis fpc slot-number inline-services flow-table-size ipv4-flow-table-size 15
```

Starting with Junos OS Release 17.3R1, for LU-based platforms, the maximum number of units is 15. For XL-based platforms, the maximum is 220. For EA-based platforms, the maximum is 48 for MPC7E and MPC9E. For MPC8E, the maximum is 97.

Options

- units** Number of 256K flow entries available for the IPv4 flow table.
- **Range:** 1 through 245. On the MPC6E, the range is 1 through 220.
 - **Default:** 1024 (1K)—Starting with Junos OS Release 16.1R1 and 15.1F2
 - **Default:** 3,932,160 (3840K)—Prior to Junos OS Release 16.1R1 and 15.1F2

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 with Junos OS Release 17.3R1, for LU-based platforms, the maximum number of units is 15. For XL-based platforms, the maximum is 220. For EA-based platforms, the maximum is 48 for MPC7E and MPC9E. For MPC8E, the maximum is 97.
16.1R1	Starting with Junos OS Release 16.1R1 and 15.1F2, by default, the software allocates 1K entries for IPv4 flow tables.

RELATED DOCUMENTATION

ipv4-template

IN THIS SECTION

- [Syntax | 1191](#)
- [Hierarchy Level | 1191](#)
- [Description | 1191](#)
- [Required Privilege Level | 1192](#)
- [Release Information | 1192](#)

Syntax

```
ipv4-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify the version 9 or IPFIX template properties for one of the following:

- Template for monitoring IPv4 flows.
- Template for inline monitoring an MPLS-over-UDP flow that is carried between IPv4 endpoints on PTX Series routers. This monitoring looks past the tunnel header to report the inner payload of the packets. To use the template for MPLS-over-UDP flows, you must also configure `tunnel-observation mpls-over-udp` at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

NOTE: For an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP, configure `mpls-ipvx-template` in Junos OS Release 18.1 or `mpls-template` starting in Junos OS 18.2R1 at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level added in Junos OS Release 10.2.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers | 615](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

[Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers | 550](#)

ipv6-flow-table-size

IN THIS SECTION

- [Syntax | 1193](#)
- [Hierarchy Level | 1193](#)
- [Description | 1193](#)
- [Options | 1193](#)

● Required Privilege Level | 1194

Syntax

```
ipv6-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services ipv6 flow-table-size]
```

Description

Configure the size of the IPv6 flow table in units of 256K entries.

NOTE: Prior to Junos OS Release 15.1F2, any changes in the configured size of the flow table initiates an automatic reboot of the FPC.

NOTE: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more, however, the system will issue a warning message.

NOTE: Starting with Junos OS Release 17.3R1, the maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 256 on MPC5Es and MPC6Es with 4 GB DDR3 memory or higher. The maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 245 on MPC5Es and MPC6Es with DDR3 memory lower than 4 GB.

Options

units Number of 256K flow entries available for the IPv6 flow table.

- **Range:** 1 through 245
- **Default:** If number of units is not specified, 1024 flow entries are allocated for IPv6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#) | 74

ipv6-extended-attrib

IN THIS SECTION

- [Syntax](#) | 1194
- [Hierarchy Level](#) | 1194
- [Description](#) | 1195
- [Required Privilege Level](#) | 1195

Syntax

```
ipv6-extended-attrib;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services ipv6 flow-table-size]
```

Description

Enable the inclusion of element ID, 54, fragmentIdentification, and element ID, 64, ipv6ExtensionHeaders, in IPFIX flow templates that are exported to the flow collector

NOTE: Collection of IPv4 fragmentation IDs occurs automatically without having to configure this setting explicitly.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#) | 74

ipv6-template

IN THIS SECTION

- [Syntax](#) | 1195
- [Hierarchy Level](#) | 1196
- [Description](#) | 1196
- [Required Privilege Level](#) | 1196
- [Release Information](#) | 1196

Syntax

```
ipv6-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify that the flow aggregation version 9 or IPFIX template is used only for IPv6 records.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.4.

Support at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level added in Junos OS Release 10.2.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers | 615](#)

ivlan-cfi (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1197](#)
- [Hierarchy Level | 1197](#)
- [Description | 1197](#)

- Required Privilege Level | 1197
- Release Information | 1197

Syntax

```
ivlan-cfi;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

CFI bit to be used in the inner VLAN header of the frames generated.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

ivlan-id (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1198](#)
- [Hierarchy Level | 1198](#)
- [Description | 1198](#)
- [Options | 1198](#)
- [Required Privilege Level | 1198](#)
- [Release Information | 1199](#)

Syntax

```
ivlan-id number;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Configure the inner VLAN ID for the test frames. This parameter is applicable for dual-tagged packets.

Options

- number* VLAN ID number.
- **Range:** 0 through 4094

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

ivlan-priority (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1199](#)
- [Hierarchy Level | 1199](#)
- [Description | 1200](#)
- [Options | 1200](#)
- [Required Privilege Level | 1200](#)
- [Release Information | 1200](#)

Syntax

```
ivlan-priority value;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Description

Configure the priority value for the IEEE 802.1p bit in the inner VLAN tag. This parameter is valid only for a bridge family for an Ethernet LAN (ELAN) or an Ethernet Line (E-LINE) service.

Options

value IEEE 802.1p priority value in the inner VLAN tag

- **Range:** 0 through 7

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

jflow-log (Interfaces)

IN THIS SECTION

- [Syntax | 1201](#)
- [Hierarchy Level | 1201](#)
- [Description | 1201](#)
- [Required Privilege Level | 1201](#)

Syntax

```
jflow-log {  
    message-rate-limit messages-per-second;  
}
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Configure generation of log messages or template records in flow monitoring format for NAT error events. These records for NAT error events are generated when addresses for allocation from the NAT pool are not available, when ports for allocation to a subscriber are not available, or when the allocated quota is exceeded for NAT events (more than the configured number of ports is requested).

The remaining statement is described separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250](#) | 241

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

jflow-log (Services)

IN THIS SECTION

- [Syntax | 1202](#)
- [Hierarchy Level | 1203](#)
- [Description | 1203](#)
- [Required Privilege Level | 1203](#)
- [Release Information | 1203](#)

Syntax

```
jflow-log {
    collector collector-name {
        source-ip address;
        destination-address address;
        destination-port port-number;
    }
    collector-group collector-group-name {
        [collector-name1 collector-name2];
    }
    template-profile template-profile-name {
        collector collector-name ;
        collector-group collector-group-name ;
        template-type nat;
        version (ipfix | v9);
        refresh-rate packets packets seconds seconds;
        message-rate-limit messages-per-second
```

Hierarchy Level

[edit services]

Description

Enable the mechanism to record logging messages in flow monitoring format for NAT events. For this transmission of flow monitoring logs to work properly, the services PIC interface must have an IP address and appropriate logging options configured.

You can configure MX Series routers with MS-MPCs and MS-MICs to log network address translation (NAT) events using the Junos Traffic Vision (previously known as Jflow) version 9 or IPFIX (version 10) template format. This method of generating flow monitoring records for NAT events, such as NAT44 and NAT64 session creation and deletion, and NAT44 and NAT64 binding information base events, enables cohesive and streamlined analysis of NAT traffic and troubleshooting of NAT-related problems.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.2R2.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

label-position

IN THIS SECTION

- [Syntax | 1204](#)
- [Hierarchy Level | 1204](#)
- [Description | 1204](#)
- [Default | 1204](#)
- [Options | 1204](#)
- [Required Privilege Level | 1205](#)
- [Release Information | 1205](#)

Syntax

```
label-position [ positions ];
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name mpls-ipv4-template],  
[edit services flow-monitoring version9 template template-name mpls-template]
```

Description

Specify positions for up to three labels in the active flow monitoring version 9 template.

Default

[1 2 3]

Options

positions—Numbered positions for the labels.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates](#) | 583

license-server

IN THIS SECTION

- [Syntax](#) | 1205
- [Hierarchy Level](#) | 1206
- [Description](#) | 1206
- [Options](#) | 1206
- [Required Privilege Level](#) | 1206
- [Release Information](#) | 1206

Syntax

```
license-server {  
    ip-address address;  
    log-interval seconds;  
    services (jflow | cgnat | firewall);  
}
```


Hierarchy Level

[edit]

Description

On MX Series routers with MS-MICs and MS-MPCs, configure the capability to transmit the throughput details per service for the Junos Address Aware, Junos Traffic Vision, and Junos Network Secure services in the last time interval to an external log collector.

Options

ip-address <i>address</i>	Use the specified IP address of the license log server.
log-interval <i>seconds</i>	Use the specified frequency at which throughput data must be sent from the router to the log collector. <ul style="list-style-type: none"> • Range: 60 through 86,400 seconds
services	Specify the services for which throughput data must be exported. <ul style="list-style-type: none"> • jflow—Use inline flow monitoring service or Junos Traffic Vision. • cgnat—Use carrier-grade NAT service or Junos Address Aware. • firewall—Use stateful firewall or Junos Network Secure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[License Server Management for Throughput Data Export on MX Series Routers for NAT, Firewall, and Inline Flow Monitoring Services | 724](#)

light

IN THIS SECTION

- [Syntax | 1207](#)
- [Junos OS Hierarchy Level | 1207](#)
- [Junos OS Evolved Hierarchy Level | 1207](#)
- [Description | 1208](#)
- [Options | 1208](#)
- [Required Privilege Level | 1208](#)
- [Release Information | 1208](#)

Syntax

```
light {  
    port number;  
}
```

Junos OS Hierarchy Level

```
[edit services rpm twamp server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

Enable the Two-Way Active Measurement Protocol (TWAMP) server for light control on the UDP port that reflects the test packets. If you want the default IANA port for TWAMP (862), you do not need to configure the port option. To complete the configuration for light control, you also need to configure the `control-type light` statement at either the Junos OS `[edit services rpm twamp client control-connection control-connection-name]` hierarchy level or the Junos OS Evolved `[edit services monitoring twamp client control-connection control-connection-name]` hierarchy level.

You configure the `light` statement because you want a stateless version of TWAMP where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away. If you want to have a stateful TWAMP server for Junos OS Evolved, you should configure the `managed` statement at the `[edit services monitoring twamp server]` hierarchy level instead. For Junos OS, if you want to have a stateful TWAMP server, do not configure the `light` statement for the TWAMP server, because managed mode is the default. Because TWAMP light servers are stateless, information about them is not included in the output of the `show services rpm twamp server connection` (Junos OS) or the `show services monitoring twamp server control-info` (Junos OS Evolved) operational mode command; only information about managed servers is included.

Options

- port *number*** Specify the UDP port that reflects the TWAMP test packets.
- **Range:** You can specify any port from 862 through 65535.
 - **Default:** 862 (IANA port for TWAMP)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.3R1.

Support added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

| [Understand Two-Way Active Measurement Protocol](#) | 686

local-dump

IN THIS SECTION

- [Syntax | 1209](#)
- [Hierarchy Level | 1209](#)
- [Description | 1209](#)
- [Options | 1209](#)
- [Required Privilege Level | 1209](#)
- [Release Information | 1210](#)

Syntax

```
(local-dump | no-local-dump);
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server  
hostname],  
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```

Description

Enable collection of cflowd records in a log file.

Options

`no-local-dump`—Do not dump cflowd records to a log file before exporting.

`local-dump`—Dump cflowd records to a log file before exporting.

Required Privilege Level

`interface`—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Flow Aggregation](#) | 577

logical-system

IN THIS SECTION

- [Syntax](#) | 1210
- [Hierarchy Level](#) | 1210
- [Description](#) | 1211
- [Options](#) | 1211
- [Required Privilege Level](#) | 1211
- [Release Information](#) | 1211

Syntax

```
logical-system logical-system-name {  
    [ routing-instances instance-name ];  
}
```

Hierarchy Level

```
[edit services rpm bgp]
```

Description

Specify the logical system used by the probes.

Options

logical-system-name—Logical system name.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

| [Configuring BGP Neighbor Discovery Through RPM](#) | 671

managed

IN THIS SECTION

- [Syntax](#) | 1212
- [Hierarchy Level](#) | 1212
- [Description](#) | 1212
- [Options](#) | 1212
- [Required Privilege Level](#) | 1213
- [Release Information](#) | 1213

Syntax

```
managed {
    client-limit limit;
    client-list list-name {
        address address <routing-instance [instance-name...]>;
    }
    control-inactivity-timeout seconds;
    control-maximum-duration seconds;
    control-per-client-limit number;
    port number;
    test-per-client-limit limit;
}
```

Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

Enable the Two-Way Active Measurement Protocol (TWAMP) server for managed control on the TCP port that reflects the test packets.

Options

client-limit <i>limit</i>	Specify the maximum number of TWAMP clients. <ul style="list-style-type: none"> • Range: 0 to 1000 • Default: 0 (disabled, which means there is no limit to the number of clients)
client-list <i>list-name</i> { address <i>address</i> < routing-instance [<i>instance-name...</i>]> }	Specify a list of clients and their corresponding addresses and routing instances. If you do not specify a routing instance, the default routing table is used (inet.0). <ul style="list-style-type: none"> • Syntax: address <i>address</i> <routing-instance [<i>instance-name...</i>]>—Specify one address statement for each client, configuring an IPv4 address and, optionally, a list of routing instances for each one, to filter packets accordingly.
control-inactivity-timeout <i>seconds</i>	Specify the number of seconds to wait after the control connection becomes inactive before timing out.

- **Range:** 0 through 86400 seconds; specify 0 to disable.
- **Default:** 900 seconds

control-maximum-duration *seconds* Specify how long the control connection remains up.

- **Range:** 0 through 86400 seconds
- **Default:** 0 (disabled, which means there is no limit to the length of time the control connection remains up)

control-per-client-limit *number* Specify the maximum number of control connections allowed per client.

- **Range:** 0 to 1000
- **Default:** 0 (disabled, which means there is no limit to the number of control connections allowed)

port *number* Specify which port can accept TWAMP control connections.

- **Range:** You can specify port 862, or any port from 49152 through 65535.
- **Default:** 862 (IANA port for TWAMP)

test-per-client-limit *limit* Specify the maximum number of test sessions allowed per client.

- **Range:** 0 to 1000
- **Default:** 0 (disabled, which means there is no limit to the number of test sessions allowed per client)

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

Understand Two-Way Active Measurement Protocol | 686

match

IN THIS SECTION

- [Syntax | 1214](#)
- [Hierarchy Level | 1214](#)
- [Description | 1214](#)
- [Required Privilege Level | 1214](#)
- [Release Information | 1214](#)

Syntax

```
match expression;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],  
[edit forwarding-options sampling traceoptions file]
```

Description

Specify the regular expression for lines to be logged for tracing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

max-connection-duration

IN THIS SECTION

- [Syntax](#) | 1215
- [Hierarchy Level](#) | 1215
- [Description](#) | 1215
- [Options](#) | 1215
- [Required Privilege Level](#) | 1216
- [Release Information](#) | 1216

Syntax

```
max-connection-duration hours;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Specify the maximum time a connection can exist between a client and the server.

Options

hours Number of hours a connection can exist between a client and the server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

max-duplicates

IN THIS SECTION

- [Syntax](#) | 1216
- [Hierarchy Level](#) | 1217
- [Description](#) | 1217
- [Default](#) | 1217
- [Options](#) | 1217
- [Required Privilege Level](#) | 1217
- [Release Information](#) | 1217

Syntax

```
max-duplicates number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify the maximum number of content destinations to which the DFC PIC can send data from a single input set of packets. Limiting the number of duplicates reduces the load on the PIC. This setting overrides the globally applied `g-max-duplicates` setting.

Default

If no value is configured, a default setting of 3 is used.

Options

number—Maximum number of content destinations.

- **Range:** 1 through 64

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[g-max-duplicates](#) | [1140](#)

[Configuring Junos Capture Vision](#) | [289](#)

max-packets-per-second

IN THIS SECTION

- [Syntax | 1218](#)
- [Hierarchy Level | 1218](#)
- [Description | 1218](#)
- [Options | 1219](#)
- [Required Privilege Level | 1219](#)
- [Release Information | 1219](#)

Syntax

```
max-packets-per-second number;
```

Hierarchy Level

```
[edit forwarding-options sampling input],  
[edit forwarding-options sampling instance instance-name input]
```

Description

Specify the traffic threshold that must be exceeded before packets are dropped. A value of 0 instructs the Packet Forwarding Engine not to sample any traffic.

NOTE: The max-packets-per-second statement is not supported when you configure inline flow monitoring (by including the inline-jflow statement at the [edit forwarding-options sampling instance *instance-name* family (inet | inet6) output] hierarchy level).

NOTE: When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the output statement, the `max-packets-per-second` value is ignored.

Options

number—Maximum number of packets per second.

- **Range:** 0 through 65,535
- **Default:** 1000

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

maximum-age

IN THIS SECTION

- [Syntax](#) | 1220
- [Hierarchy Level](#) | 1220
- [Description](#) | 1220
- [Options](#) | 1220
- [Required Privilege Level](#) | 1220

Syntax

```
maximum-age minutes;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive]
```

Description

Maximum age of transfer log file.

Options

minutes—Transfer log file age.

- **Range:** 1 through 360

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

maximum-connections

IN THIS SECTION

- [Syntax | 1221](#)
- [Hierarchy Level | 1221](#)
- [Description | 1221](#)
- [Options | 1221](#)
- [Required Privilege Level | 1222](#)
- [Release Information | 1222](#)

Syntax

```
maximum-connections count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed connections between the server and all control client hosts.

NOTE: The maximum number of connections between the server and all control client hosts must be greater than or equal to the number of connections between the server and a single controlled client host.

Options

count—Maximum number of connections.

- **Range:** 1 through 1000

- **Default:** 64

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

maximum-connections-per-client

IN THIS SECTION

- [Syntax](#) | 1222
- [Hierarchy Level](#) | 1223
- [Description](#) | 1223
- [Options](#) | 1223
- [Required Privilege Level](#) | 1223
- [Release Information](#) | 1223

Syntax

```
maximum-connections-per-client count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed connections between the server and a single control client host.

NOTE: The maximum number of connections between the server and all control client hosts must be greater than or equal to the number of connections between the server and a single controlled client host.

Options

count—Maximum number of connections.

- **Range:** 1 through 500
- **Default:** 64

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

maximum-packet-length

IN THIS SECTION

- [Syntax | 1224](#)
- [Hierarchy Level | 1224](#)
- [Description | 1224](#)
- [Options | 1225](#)
- [Required Privilege Level | 1225](#)
- [Release Information | 1225](#)

Syntax

```
maximum-packet-length bytes;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input],  
[edit forwarding-options port-mirroring input],  
[edit forwarding-options port-mirroring instance instance-name input],  
[edit forwarding-options sampling input],  
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the maximum packet length to be used for port mirroring or traffic sampling. Packets longer than the maximum are truncated. This statement cannot be used with inline flow monitoring ([edit forwarding-options sampling instance *instance-name* family (inet | inet6) output inline-jflow]).

NOTE: For MX Series routers with Modular Port Interface Concentrators (MPCs), when `maximum-packet-length (clip length)` is configured for port-mirrored packets and the mirror-destination

interface is a next-hop-group, the clip length is effective only for the first member interface of the next-hop-group. The mirrored packet copy sent to the rest of the interfaces is not clipped. In addition, native analyzer sessions (that is, the [edit forwarding-options analyzer analyzer-name input] hierarchy level for MX Series routers) can be configured without specifying input parameters. As such, these instances use the following input values by default: rate = 1, and maximum-packet-length = 0.

Options

bytes—Maximum length (in bytes) of the mirrored packet or the sampled packet.

Set the maximum-packet-length value to zero to disable truncation; that is, to mirror or sample the entire packet. Otherwise, Juniper recommends that you configure the packet length to be equal to, or greater than, the IP header length. For IPv4, set the maximum length to at least 20, and for IPv6, set the maximum length to at least 40.

- **Range:** 0 through 9216. For MX Series routers with MPCs, and for EX9200 switches, the range is 1 through 255 bytes.
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

For MX Series routers except the MX 80, support at the [edit forwarding-options analyzer analyzer-name input] hierarchy level was introduced in Junos OS Release 14.1

RELATED DOCUMENTATION

Configuring Port Mirroring

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

maximum-sessions

IN THIS SECTION

- [Syntax | 1226](#)
- [Hierarchy Level | 1226](#)
- [Description | 1226](#)
- [Options | 1226](#)
- [Required Privilege Level | 1227](#)
- [Release Information | 1227](#)

Syntax

```
maximum-sessions count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed test sessions the server can have running at one time.

NOTE: The maximum number of test sessions running on the server at one time must be greater than or equal to the maximum number of sessions the server can open on a single client connection.

Options

count—Maximum number of sessions.

- **Range:** 1 through 2048

- **Default:** 64

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

maximum-sessions-per-connection

IN THIS SECTION

- [Syntax](#) | 1227
- [Hierarchy Level](#) | 1228
- [Description](#) | 1228
- [Options](#) | 1228
- [Required Privilege Level](#) | 1228
- [Release Information](#) | 1228

Syntax

```
maximum-sessions-per-connection count;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the maximum number of allowed sessions the server can open on a single client connection.

NOTE: The maximum number of test sessions running on the server at one time must be greater than or equal to the maximum number of sessions the server can open on a single client connection.

Options

count—Maximum number of sessions.

- **Default:** 64

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

media-loss-rate

IN THIS SECTION

- [Syntax | 1229](#)
- [Hierarchy Level | 1229](#)
- [Description | 1229](#)
- [Required Privilege Level | 1230](#)
- [Release Information | 1230](#)

Syntax

```
media-loss-rate {  
    no-syslog-generation;  
    generate-snmp-traps;  
    storm-control {  
        count number;  
        interval number;  
    }  
    alarm-mode {  
        immediate;  
    }  
}
```

Hierarchy Level

[edit services]

Description

Configure the media loss rate. The media loss rate is the number of media packets lost over a configurable time interval (interval-duration) where the flow packets are packets carrying streaming application information. A single IP packet can contain zero or more streaming packets.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers | 927](#)
[alarms | 973](#)

media-rate-variation

IN THIS SECTION

- [Syntax | 1230](#)
- [Hierarchy Level | 1231](#)
- [Description | 1231](#)
- [Required Privilege Level | 1231](#)
- [Release Information | 1231](#)

Syntax

```
media-rate-variation {  
  no-syslog-generation;  
  generate-snmp-traps;  
  storm-control {  
    count number;  
    interval number;  
  }  
}
```

```
alarm-mode {  
    mdi-records-count number;  
    average;  
}  
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the media rate variation. The media rate variation is the difference between the expected packet rate and actual packet rate expressed as a percentage of the expected packet rate.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers](#) | 927

[alarms](#) | 973

message-rate-limit (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax | 1232](#)
- [Hierarchy Level | 1232](#)
- [Description | 1232](#)
- [Options | 1233](#)
- [Required Privilege Level | 1233](#)
- [Release Information | 1233](#)

Syntax

```
message-rate-limit messages-per-second
```

Hierarchy Level

```
[edit interfaces interface-name services-options jflow-log]
```

Description

Define the maximum number of logs or template records in flow monitoring format to be generated for NAT error events per second from the specified interface. These records for NAT error events are generated when addresses for allocation from the NAT pool are not available, when ports for allocation to a subscriber are not available, or when the allocated quota is exceeded for NAT events (more than the configured number of ports is requested).

NOTE: The `message-rate-limit` option can be configured only for multiservices interfaces (ms-*x/x/x*) and not with other interface types.

Options

messages-per-second —Maximum number of flow monitoring log messages per second for NAT error events that can be formatted and sent from the PIC to an external collector. The default rate is 10,000 for an external collector.

- **Range:** 1 through 2,147,483,647

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

minimum-priority

IN THIS SECTION

- [Syntax | 1234](#)
- [Hierarchy Level | 1234](#)
- [Description | 1234](#)

- Options | 1234
- Required Privilege Level | 1234
- Release Information | 1234

Syntax

```
minimum-priority value;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Specify the minimum priority for the control source.

Options

value—Minimum priority value; if not specified, defaults to 0.

- **Range:** 0 through 254

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[Configuring Junos Capture Vision](#) | 289

mode

IN THIS SECTION

- [Syntax](#) | 1235
- [Junos OS Hierarchy Level](#) | 1235
- [Junos OS Evolved Hierarchy Level](#) | 1235
- [Description](#) | 1235
- [Options](#) | 1236
- [Required Privilege Level](#) | 1236
- [Release Information](#) | 1236

Syntax

```
mode (initiate-and-terminate | reflect);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the test mode for the packets that are sent during the benchmarking test.

Options

initiate-and-terminate	(Most ACX routers, except for the ACX7100, ACX7509, ACX7024, ACX5448, ACX5048, and ACX5096) Initiate and terminate the test on the chosen service (IPv4 or Ethernet).
reflect	(MX routers and the ACX7100, ACX7509, ACX7024, ACX5448, ACX5048, and ACX5096 routers) Reflect the test frames on the chosen service (IPv4 or Ethernet).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1 for reflect mode only.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[rfc2544-benchmarking | 1347](#)

[rfc2544 | 1345](#)

monitoring (Forwarding Options)

IN THIS SECTION

- [Syntax | 1237](#)
- [Hierarchy Level | 1237](#)
- [Description | 1237](#)

- Required Privilege Level | 1238
- Release Information | 1238

Syntax

```
monitoring name {  
    family inet {  
        output {  
            cflowd hostname port-number;  
            export-format cflowd-version-5;  
            flow-active-timeout seconds;  
            flow-export-destination {  
                (cflowd-collector | collector-pic);  
            }  
            flow-inactive-timeout seconds;  
            interface interface-name {  
                number;  
                engine-type number;  
                input-interface-index number;  
                output-interface-index number;  
                source-address address;  
            }  
        }  
    }  
}
```

Hierarchy Level

[edit forwarding-options]

Description

Specify the flow monitoring instance name and properties.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

monitoring (Services)

IN THIS SECTION

- [Syntax](#) | 1238
- [Hierarchy Level](#) | 1241
- [Description](#) | 1241
- [Required Privilege Level](#) | 1241
- [Release Information](#) | 1241

Syntax

```
monitoring {  
  rfc2544 {  
    tests{  
      test-name test-name {  
        test-interface interface-name;  
        mode reflect;  
        family inet;  
        destination-ipv4-address address;
```

```

        destination-udp-port port-number;
        source-ipv4-address address;
        source-udp-port port-number;
    }
}
}
rpm {
    owner name {
        test test-name {
            data-fill data;
            data-size size;
            destination-port port;
            dscp-code-points (RPM) dscp-bits;
            history-size size;
            moving-average-size number;
            offload-type {
                none;
                pfe-timestamp;
            }
            probe-count count;
            probe-interval seconds;
            probe-type type;
            routing-instance instance-name;
            source-address address;
            target (url url | address address);
            test-interval interval;
            thresholds thresholds;
            ttl hop-count
        }
    }
}
twamp {
    client {
        control-connection name {
            control-type (managed | light);
            destination-port destination-port;
            routing-instance routing-instance;
            source-address source-address;
            target target-address;
            test-start (auto | manual);
            test-interval seconds;
            test-session name {

```

```

        data-size data-size;
        destination-port destination-port;
        dscp-code-points dscp-code-points;
        history-size history-size;
        moving-average-size moving-average-size;
        offload-type (none | pfe-timestamp);
        probe-count probe-count;
        probe-interval seconds;
        source-address source-address;
        target target-address;
        thresholds {
            control-failure (on | off);
            successive-loss number;
            total-loss number;
            threshold-type (microseconds | average);
        }
        ttl hop-count;
    }
}

server {
    managed {
        client-limit limit;
        client-list {
            address address <routing-instance [instance-name...]>;
        }
        control-inactivity-timeout seconds;
        control-per-client-limit limit;
        control-maximum-duration seconds;
        port port;
        test-per-client-limit limit;
    }
    light {
        port port;
    }
}
}
}

```

Hierarchy Level

[edit services]

Description

Configure monitoring services.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

twamp statement introduced in Junos OS Evolved Release 20.3R1.

rfc2544 statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

moving-average-size

IN THIS SECTION

● [Syntax | 1242](#)

- [Junos OS Hierarchy Levels | 1242](#)
- [Junos OS Evolved Hierarchy Level | 1242](#)
- [Description | 1242](#)
- [Options | 1242](#)
- [Required Privilege Level | 1243](#)
- [Release Information | 1243](#)

Syntax

```
moving-average-size number;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],
[edit services rpm probe owner test test-name],
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Enable statistical calculation operations to be performed across a configurable number of the most recent samples.

Options

number—Number of samples to be used in calculations.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

mpls-flow-table-size

IN THIS SECTION

- [Syntax | 1244](#)
- [Hierarchy Level | 1244](#)
- [Description | 1244](#)
- [Options | 1244](#)
- [Required Privilege Level | 1244](#)
- [Release Information | 1245](#)

Syntax

```
mpls-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services flow-table-size]
```

Description

Configure the size of the MPLS flow table in units of 256,000 entries.

NOTE: Starting with Junos OS Release 17.3R1, the maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 256 on MPC5Es and MPC6Es with 4 GB DDR3 memory or higher. The maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 245 on MPC5Es and MPC6Es with DDR3 memory lower than 4 GB.

NOTE: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more, however, the system will issue a warning message.

Options

units Number of 256,000 flow entries available for the MPLS flow table.

- **Range:** 1 through 245
- **Default:** 15 (3,840,000)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

| [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#) | 74

mpls-ipv4-template

IN THIS SECTION

- [Syntax](#) | 1245
- [Hierarchy Level](#) | 1245
- [Description](#) | 1246
- [Required Privilege Level](#) | 1246
- [Release Information](#) | 1246

Syntax

```
mpls-ipv4-template {  
    label-position [ positions ];  
}
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name]
```


Description

Specify the flow aggregation version 9 or IPFIX properties for templates that combine IPv4 and MPLS records. The remaining statement is explained separately.

Starting in Junos OS Release 18.4R1, use `mpls-template` instead of `mpls-ipv4-template` for inline flow monitoring of MPLS-IPv4 flows on the MX Series. You must also configure `tunnel-observation ipv4` at the `[edit services flow-monitoring (version-ipfix | version9) template template-name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the `[edit services flow-monitoring version-ipfix template template-name]` hierarchy level introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers | 550](#)

[Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers | 550](#)

mpls-ipvx-template

IN THIS SECTION

- [Syntax | 1247](#)
- [Hierarchy Level | 1247](#)
- [Description | 1247](#)
- [Required Privilege Level | 1247](#)

Syntax

```
mpls-ipvx-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name]
```

Description

In Junos OS Release 18.1, specify the version 9 or IPFIX template for inline monitoring an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP on PTX Series routers. This monitoring looks past the tunnel header to report the inner payload of the packets. To use the template for MPLS-over-UDP flows, you must also configure `tunnel-observation mpls-over-udp` at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

Starting in Junos OS Release 18.2R1, use `mpls-template` instead of `mpls-ipvx-template`.

NOTE: For an MPLS-over-UDP flow that is carried between IPv4 endpoints, configure `ipv4-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template template-name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

RELATED DOCUMENTATION

| [Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers](#) | 550

mpls-template

IN THIS SECTION

- [Syntax](#) | 1248
- [Hierarchy Level](#) | 1248
- [Description](#) | 1248
- [Required Privilege Level](#) | 1249
- [Release Information](#) | 1249

Syntax

```
mpls-template {  
    label-position [positions];  
}
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify the flow aggregation IPFIX or version 9 properties for templates used only for MPLS records.

Starting in Junos OS Release 18.2R1, you can also use `mpls-template` to specify the version 9 or IPFIX template for inline monitoring an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP on PTX Series routers. (In Junos OS Release 18.1, use `mpls-ipvx-template` instead of `mpls-template`.) This monitoring looks past the tunnel header to report the inner payload of the packets. To use the template for MPLS-

over-UDP flows, you must also configure tunnel-observation mpls-over-udp at the [edit services flow-monitoring (version 9 | version-ipfix) template *template-name*] hierarchy level.

NOTE: For an MPLS-over-UDP flow that is carried between IPv4 endpoints, configure ipv4-template at the [edit services flow-monitoring (version9 | version-ipfix) template *template-name*] hierarchy level.

Starting in Junos OS Release 18.4R1, use mpls-template instead of mpls-ipv4-template for inline flow monitoring of MPLS-IPv4 flows on the MX Series. You must also configure tunnel-observation ipv4 at the [edit services flow-monitoring (version-ipfix | version9) template *template-name*] hierarchy level.

The remaining statement is explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the [edit services [flow-monitoring version-ipfix template *template-name*](#)] hierarchy level introduced in Junos OS Release 16.1.

Statement introduced in Junos OS Release 18.2R1 on PTX Series routers.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates](#) | 583

multiservice-options

IN THIS SECTION

- [Syntax | 1250](#)
- [Hierarchy Level | 1250](#)
- [Description | 1250](#)
- [Required Privilege Level | 1251](#)
- [Release Information | 1251](#)

Syntax

```
multiservice-options {  
    (core-dump | no-core-dump);  
    (syslog | no-syslog);  
    flow-control-options {  
        down-on-flow-control;  
        dump-on-flow-control;  
        reset-on-flow-control;  
    }  
}
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port]
```

Description

For flow-monitoring interfaces only, configure multiservice-specific interface properties.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

name-format

IN THIS SECTION

- [Syntax](#) | 1251
- [Hierarchy Level](#) | 1251
- [Description](#) | 1252
- [Options](#) | 1252
- [Required Privilege Level](#) | 1253
- [Release Information](#) | 1253

Syntax

```
name-format "format";
```

Hierarchy Level

```
[edit services flow-collector file-specification variant variant-number]
```

Description

Specify the name format for a specific file format. The files can include supported macros. Use macros to organize files on the external machine to which they are exported from the collector PIC.

Options

format—Specify the filename format, within quotation marks. The name format can include the following macros, typed in braces:

- `{%D}`—Date
- `{%T}`—Time when the file is created
- `{%I}`—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy level
- `{%N}`—Unique, sequential number for each new file created
- `{am_pm}`—AM or PM
- `{date}`—Current date using the `{year}` `{month}` `{day}` macros
- `{day}`—From 01 through 31
- `{day_abbrev}`—Sun through Sat
- `{day_full}`—Sunday through Saturday
- `{generation number}`—Unique, sequential number for each new file created
- `{hour_12}`—From 01 through 12
- `{hour_24}`—From 00 through 23
- `{ifalias}`—Description string for the logical interface configured using the collector statement at the [edit services flow-collector interface-map] hierarchy level
- `{minute}`—From 00 through 59
- `{month}`—From 01 through 12
- `{month_abbrev}`—Jan through Dec
- `{month_full}`—January through December
- `{num_zone}`—From -2359 through +2359; this macro is not supported

- {second}—From 00 through 60
- {time}—Time the file is created, using the {hour_24} {minute} {second} macros
- {time_zone}—Time zone code name of the locale; for example,gmt (this macro is not supported).
- {year}—In the format YYYY; for example, 1970
- {year_abbr}—From 00 through 99

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

next-hop (Forwarding Options)

IN THIS SECTION

- [Syntax](#) | 1254
- [Hierarchy Level](#) | 1254
- [Description](#) | 1254
- [Options](#) | 1254
- [Required Privilege Level](#) | 1254
- [Release Information](#) | 1254

Syntax

```
next-hop address;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring family (inet | inet6) output interface interface-name]
```

Description

Specify the next-hop address for sending copies of packets to an analyzer.

Options

address—IP address of the next-hop router.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

next-hop (RPM)

IN THIS SECTION

- [Syntax | 1255](#)
- [Hierarchy Level | 1255](#)
- [Description | 1255](#)
- [Required Privilege Level | 1255](#)
- [Release Information | 1255](#)

Syntax

```
next-hop next-hop;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Specify the next-hop address to which the probe should be sent.

Required Privilege Level

services—To view this statement in the configuration.

services-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.4.

RELATED DOCUMENTATION

[probe](#) | [1313](#)

next-hop-group (Forwarding Options)

IN THIS SECTION

- [Syntax](#) | [1256](#)
- [Hierarchy Level](#) | [1256](#)
- [Description](#) | [1256](#)
- [Options](#) | [1257](#)
- [Required Privilege Level](#) | [1257](#)
- [Release Information](#) | [1257](#)

Syntax

```
next-hop-group group-name {  
    interface interface-name {  
        next-hop address;  
    }  
}
```

Hierarchy Level

```
[edit forwarding-options]
```

Description

Specify the next-hop address for sending copies of packets to an analyzer.

It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.

NOTE: In Junos OS releases earlier through Release 14.2, the `next-hop-group` statement is present in the forwarding-options stanza for a routing instance, but the `next-hop-group` statement is not allowed in a routing instance. In other words, in a routing instance, `[edit routing-instances routing-instance-name forwarding-options next-hop-group]` is not supported. You will get an error message if you try to commit this type of configuration. Starting in Junos OS Release 14.2, the `next-hop-group` statement is not present in `[edit routing-instances routing-instance-name forwarding-options]`.

Options

address—IP address of the next-hop router. Each next-hop group supports up to 16 next-hop addresses. Up to 30 next-hop groups are supported. Each next-hop group must have at least two next-hop addresses.

group-name—Name of next-hop group. Up to 30 next-hop groups are supported for the router. Each next-hop group is expected to have at least two next-hop addresses.

interface-name—Name of interface used to reach the next-hop destination.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

nexthop-learning

IN THIS SECTION

- [Syntax | 1258](#)
- [Hierarchy Level | 1258](#)
- [Description | 1258](#)
- [Options | 1259](#)
- [Required Privilege Level | 1259](#)
- [Release Information | 1260](#)

Syntax

```
nexthop-learning (disable | enable);
```

Hierarchy Level

```
[edit services flow-monitoring (version-ipfix | version9) template template-name]
```

Description

Enable learning of next-hop addresses to correctly report the next hop address, output SNMP, destination IP address, destination IP mask values, the packet loss priority, and the first two characters of the configured forwarding class name in the flow records when a destination is reachable through multiple paths. By default, this behavior of learning the next-hop addresses is disabled for inline active flow monitoring.

For PTX Series, starting in Junos OS Evolved 21.2R1 and in Junos OS Release 21.3R1, every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. The effect of the `nexthop-learning` statement on this behavior varies depending upon the operating system. For Junos OS Evolved, we do not recommend that you configure the `nexthop-learning` statement, as it reduces the number of packets that can be processed. For Junos OS, you can configure the `nexthop-learning` statement

to change this default no-flow behavior and once again create and maintain flows, then attach the template to all sampling instances associated with FPCs that require the previous behavior.

When learning of next-hop addresses is disabled, data is reported as follows:

- If the destination address of the sampled IPv4 flow is reachable through multiple paths, the `ipNextHopIPv4Address` (Element ID 15) and `egressInterface` (Element ID 14) in the IPv4 flow record are set to the gateway IP address and SNMP index of the first path seen in the forwarding table.
- If the destination address of the sampled IPv6 flow is reachable through multiple paths, the `ipNextHopIPv6Address` (Element ID 62) and `egressInterface` (Element ID 14) in the IPv6 flow records are set to 0.
- The Incoming Interface (IIF) and Outgoing Interface (OIF) should be part of the same VRF. If the OIF is in a different VRF, `destinationIPv4PrefixLength` (Element ID 13), `bgpDestinationAsNumber` (Element ID 17), `ipNextHopIPv4Address` (Element ID 15), and `egressInterface` (Element ID 14) are set to 0 in IPv4 flow records and `destinationIPv6PrefixLength` (Element ID 30), `bgpDestinationAsNumber` (Element ID 17), `ipNextHopIPv6Address` (Element ID 62), and `egressInterface` (Element ID 14) are set to 0 in IPv6 flow records.
- The packet loss priority and forwarding class information is not reported for the PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016 routers.

When learning of next-hop addresses is enabled, output SNMP, destination IP address, destination IP mask values, the packet loss priority, and the first two characters of the configured forwarding class in the flow records are reported correctly. In addition, when enabled, `mplsTopLabelIPv4Address` (Element ID 47) in IPv4 flow records reports correctly when MPLS ingress sampling is enabled.

NOTE: Nexthop learning is supported only when sampling is implemented on the PFE. This is known as inline active flow monitoring. Nexthop learning does not work when sampling is configured on the MS-DPC/MS-MPC/MS-MIC service cards.

Options

`disable`—Disable the learning of next-hop information required for inline active flow monitoring.

`enable`—Enable the learning of next-hop information required for inline active flow monitoring.

- **Default:** `disable`

Required Privilege Level

`system`—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1F2.

nexthop-learning is supported for the bridge-template statement in Junos OS Release 18.2R1 for MX Series routers.

Support introduced in Junos OS Release 20.3R1 to correctly report the packet loss priority and the first two characters of the configured forwarding class in the IPv4 and IPv6 IPFIX templates for the PTX1000, PTX10008 (without the JNP10008-SF3), and PTX10016 routers.

RELATED DOCUMENTATION

[Configuring Next-Hop Address Learning on MX Series and PTX Series Routers for Destinations Accessible Over Multiple Paths](#) | 641

no-remote-trace (Trace Options)

IN THIS SECTION

- [Syntax](#) | 1260
- [Hierarchy Level](#) | 1261
- [Description](#) | 1261
- [Required Privilege Level](#) | 1261
- [Release Information](#) | 1261

Syntax

```
no-remote-trace;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions],  
[edit forwarding-options sampling traceoptions]
```

Description

Disable remote tracing.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

no-syslog

IN THIS SECTION

- [Syntax](#) | 1262
- [Hierarchy Level](#) | 1262
- [Description](#) | 1262
- [Required Privilege Level](#) | 1262
- [Release Information](#) | 1262

Syntax

```
no-syslog;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Disable system logging of control protocol requests and responses. By default, these messages are logged.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

no-syslog-generation

IN THIS SECTION

- [Syntax](#) | 1263
- [Hierarchy Level](#) | 1263

- [Description | 1263](#)
- [Required Privilege Level | 1263](#)
- [Release Information | 1263](#)

Syntax

```
no-syslog-generation;
```

Hierarchy Level

```
[edit services]
```

Description

Disable system log generation.

NOTE: If this statement is not configured, `edit services` generates a system log with respective severity level for values not within the configured range.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers | 927](#)

notification-targets

IN THIS SECTION

- [Syntax | 1264](#)
- [Hierarchy Level | 1264](#)
- [Description | 1264](#)
- [Options | 1264](#)
- [Required Privilege Level | 1265](#)
- [Release Information | 1265](#)

Syntax

```
notification-targets address port port-number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

List the destination IP addresses and User Datagram Protocol (UDP) ports to which DFC PICs log exception information and control protocol state transitions, such as timeout values.

Options

address—Allowed destination IP address.

port *port-number*—Allowed destination UDP port number.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

observation-domain-id

IN THIS SECTION

- [Syntax](#) | 1265
- [Hierarchy Level](#) | 1265
- [Description](#) | 1266
- [Options](#) | 1266
- [Required Privilege Level](#) | 1266
- [Release Information](#) | 1266

Syntax

```
observation-domain-id domain-id;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template template-name]
```

Description

For IPFIX flows, an identifier of an observation domain is locally unique to an exporting process of the templates. The export process uses the observation domain ID to uniquely identify to the collection process in which the flows were metered. We recommend that you configure this ID to be unique for each IPFIX flow. A value of 0 indicates that no specific Observation Domain is identified by this information element. The ID configured here is contained within Information Element 149.

MX and QFX Series: Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an IPFIX device. If you configure the same observation domain ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base observation domain ID is transmitted in the flow. The actual observation domain ID is derived from the value you configure and plus other parameters such as the slot number, lookup chip (LU) instance, and Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two IPFIX devices.

PTX Series: When you configure the observation domain ID, the software attaches the ID to a particular template type. If you configure the same observation domain ID for two different template types, such as for IPv4 and IPv6, this does not impact flow monitoring, because the configured ID is not what is being sent. The value sent in the packets is derived from that configured value and the FPC slot value. This method ensures two IPFIX devices can never have the same value of observation domain ID.

Options

domain-id—Identifier for the observation domain for IPFIX flows.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

[Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows | 620](#)

[Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows | 625](#)

offload-type

IN THIS SECTION

- [Syntax | 1267](#)
- [Hierarchy Level | 1267](#)
- [Description | 1267](#)
- [Options | 1268](#)
- [Required Privilege Level | 1268](#)
- [Release Information | 1268](#)

Syntax

```
offload-type {  
    none;  
    pfe-timestamp;  
}
```

Hierarchy Level

```
[edit services monitoring rpm owner name test test-name]
```

Description

Enable timestamping of RPM probe messages in the Packet Forwarding Engine (PFE) host processor, by offloading the processing of RPM probes to the PFE. This feature is supported only with `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp` probe types.

Options

none	Timestamping performed on the Routing Engine
pfe-timestamp	Timestamping performed on the PFE

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#) | 646

one-way-hardware-timestamp

IN THIS SECTION

- [Syntax](#) | 1269
- [Hierarchy Level](#) | 1269
- [Description](#) | 1269
- [Required Privilege Level](#) | 1269
- [Release Information](#) | 1269

Syntax

```
one-way-hardware-timestamp;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Enable timestamping of RPM probe messages for one-way delay and jitter measurements. You must configure this statement along with the `destination-interface` statement to invoke timestamping. This feature is supported only with `icmp-ping`, `icmp-ping-timestamp`, `udp-ping`, and `udp-ping-timestamp` probe types.

NOTE: Starting in Junos OS Evolved Release 20.1R1, the function provided by this statement has been replaced by the `offload-type (none|pfe-timestamp)` statement.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.5.

Statement introduced in Junos OS Release 19.1 for PTX Series routers.

RELATED DOCUMENTATION

[Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches](#) | 663

[destination-interface](#) | 1046

[hardware-timestamp](#) | 1147

option-refresh-rate

IN THIS SECTION

- [Syntax | 1270](#)
- [Hierarchy Level | 1270](#)
- [Description | 1270](#)
- [Options | 1270](#)
- [Required Privilege Level | 1271](#)
- [Release Information | 1271](#)

Syntax

```
option-refresh-rate packets packets seconds seconds;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name],
```

Description

Specify the frequency at which the flow generator sends updates about template options, like sampling rate, to the flow collector. Specify the refresh rate, in either the number of packets or seconds.

Options

packets—Refresh rate, in number of packets.

- **Range:** 1 through 480,000
- **Default:** 4800

seconds—Refresh rate, in number of seconds.

- **Range:** 10 through 600
- **Default:** 600

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the [edit services flow-monitoring version-ipfix template *template-name*] hierarchy level added in Junos OS Release 10.2.

Support at the [edit services flow-monitoring version9 template *template-name*] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

options-template-id

IN THIS SECTION

- [Syntax | 1272](#)
- [Hierarchy Level | 1272](#)
- [Description | 1272](#)
- [Options | 1273](#)
- [Required Privilege Level | 1273](#)
- [Release Information | 1273](#)

Syntax

```
options-template-id id;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Define a unique options template ID to be used for flow aggregation of version 9 and IPFIX flows. After a re-start of the exporting process, the software may re-assign option template IDs.

MX and QFX Series

The options template ID can be a value in the range of 1024 through 65535. If you do not configure values for the options template ID, a default value is assumed for this ID, which is different for the various address families. If you configure the same options template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same options template ID value for both IPv4 and IPv6, the collector validates the export data based on the options template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4.

PTX Series

If you choose to configure the options template ID, the range is 1024 to 65520. If you do not configure this ID, the default value that is set is in the range 256-511 and is different for each template.

You can configure the `option-template-id` statement for family `inet`, `inet6`, and `mpls` only.

You must not configure the same ID for different templates (option or data template). The operating system does not check to ensure that you do not configure the same ID value for different templates. If you configure the same ID value, such a setting is not processed properly and might cause unexpected behavior.

The options template ID range [configured `options-template-id` value + 20) is reserved and you must not configure any another ID in this range. The difference between configured options template IDs across families should be at least 20; for example, if `options-template-id 1056` is configured for family `inet`, you should not configure template IDs in the range of 1056 to 1075 for any other families.

For Junos OS, if you change the options template ID, all flows are inactively timed out. New flows are learned afresh.

For Junos OS Evolved, if you change the options template ID, this change does not impact the flows.

Options

id—Unique identifier for the options template to be used for version 9 or IPFIX flows. The range 0-255 is reserved and is used for template sets, options template sets, and for future expansion. If you do not configure an options template ID, the software generates a template ID for you.

- **Range:** 1024 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows | 620](#)

[Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows | 625](#)

outer-tag-protocol-id (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1274](#)
- [Hierarchy Level | 1274](#)
- [Description | 1274](#)

- Required Privilege Level | 1274
- Release Information | 1274

Syntax

```
outer-tag-protocol-id;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

TPID to be used in the outer VLAN tag. Supported values are 0x8100, 0x88a8, 0x9100, 0x9200

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

output (Accounting)

IN THIS SECTION

- [Syntax | 1275](#)
- [Hierarchy Level | 1276](#)
- [Description | 1276](#)
- [Required Privilege Level | 1276](#)
- [Release Information | 1276](#)

Syntax

```
output {  
    aggregate-export-interval seconds;  
    cflowd hostname {  
        aggregation {  
            autonomous-system;  
            destination-prefix;  
            protocol-port;  
            source-destination-prefix {  
                caida-compliant;  
            }  
            source-prefix;  
        }  
        autonomous-system-type (origin | peer);  
        (local-dump | no-local-dump);  
        port port-number;  
        source-address address;  
        version format;  
    }  
    flow-active-timeout seconds;  
    flow-inactive-timeout seconds;  
    interface interface-name {  
        engine-id number;  
        engine-type number;  
        source-address address;  
    }  
}
```

```
}
}
```

Hierarchy Level

```
[edit forwarding-options accounting name]
```

Description

Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Discard Accounting](#) | 435

output (Monitoring)

IN THIS SECTION

- [Syntax](#) | 1277
- [Hierarchy Level](#) | 1277
- [Description](#) | 1277

- Required Privilege Level | 1277
- Release Information | 1278

Syntax

```
output {
  cflowd hostname port port-number;
  export-format format;
  flow-active-timeout seconds;
  flow-export-destination {
    (cflowd-collector | collector-pic);
  }
  flow-inactive-timeout seconds;
  interface interface-name {
    engine-id number;
    engine-type number;
    input-interface-index number;
    output-interface-index number;
    source-address address;
  }
}
```

Hierarchy Level

```
[edit forwarding-options monitoring name family inet]
```

Description

Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

output (Sampling)

IN THIS SECTION

- [Syntax](#) | 1278
- [Hierarchy Level](#) | 1279
- [Description](#) | 1280
- [Required Privilege Level](#) | 1280
- [Release Information](#) | 1280

Syntax

```
output {  
    aggregate-export-interval seconds;  
    flow-active-timeout seconds;  
    flow-inactive-timeout seconds;  
    extension-service service-name;  
    flow-server hostname {  
        aggregation {  
            autonomous-system;  
            destination-prefix;  
            protocol-port;  
            source-destination-prefix {  
                caida-compliant;  
            }  
        }  
    }  
}
```

```

    }
    source-prefix;
  }
  autonomous-system-type (origin | peer);
  dscp dscp-value;
  forwarding-class class-name;
  (local-dump | no-local-dump);
  port port-number;
  source-address address;
  version format;
  version9 {
    template template-name;
  }
}
interface interface-name {
  engine-id number;
  engine-type number;
  source-address address;
}
file {
  disable;
  filename filename;
  files number;
  size bytes;
  (stamp | no-stamp);
  (world-readable | no-world-readable);
}
inline-jflow {
  source-address address;
  flow-export-rate rate;
}
}

```

Hierarchy Level

```

[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls)],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls)]

```

Description

Configure cflowd or flow monitoring, output files and interfaces, and flow properties.

The remaining statements are explained separately. See [CLI Explorer](#).

NOTE: The `inline-jflow` statement is valid only under the `[edit forwarding-options sampling instance instance-name family inet output]` hierarchy level. The `file` statement is valid only under the `[edit forwarding-options sampling family inet output]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

output-interface-index

IN THIS SECTION

- [Syntax](#) | 1281
- [Hierarchy Level](#) | 1281
- [Description](#) | 1281
- [Options](#) | 1281
- [Required Privilege Level](#) | 1281

Syntax

```
output-interface-index number;
```

Hierarchy Level

```
[edit forwarding-options monitoring name output interface interface-name]
```

Description

Specify a value for the output interface index that overrides the default supplied by SNMP.

Options

number—Output interface index value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

ovlan-cfi (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1282](#)
- [Hierarchy Level | 1282](#)
- [Description | 1282](#)
- [Required Privilege Level | 1282](#)
- [Release Information | 1282](#)

Syntax

```
ovlan-cfi;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

CFI bit to be used in the outer VLAN tag.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[rfc2544-benchmarking](#) | 1347

ovlan-id (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax](#) | 1283
- [Hierarchy Level](#) | 1283
- [Description](#) | 1283
- [Options](#) | 1284
- [Required Privilege Level](#) | 1284
- [Release Information](#) | 1284

Syntax

```
ovlan-id number;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Configure the outer VLAN ID for the test frames. This parameter is applicable for single tagged and dual-tagged packets.

Options

number

VLAN ID number.

- **Range:** 0 through 4094

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

ovlan-priority (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1285](#)
- [Hierarchy Level | 1285](#)
- [Description | 1285](#)
- [Options | 1285](#)
- [Required Privilege Level | 1285](#)
- [Release Information | 1285](#)

Syntax

```
ovlan-priority value;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Description

Configure the priority value for the IEEE 802.1p bit in the outer VLAN tag. This parameter is valid only for a bridge family for an Ethernet LAN (ELAN) or an Ethernet Line (E-LINE) service.

Options

value IEEE 802.1p priority value in the outer VLAN tag

- **Range:** 0 through 7

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | [855](#)

[Configuring RFC 2544-Based Benchmarking Tests](#) | [864](#)

[rfc2544-benchmarking](#) | [1347](#)

owner

IN THIS SECTION

- [Syntax | 1286](#)
- [Hierarchy Level | 1287](#)
- [Description | 1287](#)
- [Options | 1287](#)
- [Required Privilege Level | 1287](#)
- [Release Information | 1287](#)

Syntax

```
owner name {
  test test-name {
    data-fill data;
    data-size size;
    destination-port port;
    dscp-code-points (RPM) dscp-bits;
    history-size size;
    moving-average-size number;
    offload-type {
      none;
      pfe-timestamp;
    }
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instance instance-name;
    source-address address;
    target (url url | address address);
    test-interval interval;
    thresholds thresholds;
    traps traps;
    ttl hop-count
```

```
}
}
```

Hierarchy Level

```
[edit services monitoring rpm]
```

Description

Specify an owner name and test options. The owner name combined with the test options for a particular test name represents a single RPM configuration instance.

Options

owner
name You can configure any name, up to 32 characters in length, or specify one of the pre-defined owner names:

- icmp-evo
- icmp-junos
- udp-evo
- udp-junos

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.1R1.

traps option introduced in Junos OS Evolved Release 21.2R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

packet-loss-priority (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1288](#)
- [Hierarchy Level | 1288](#)
- [Description | 1288](#)
- [Required Privilege Level | 1288](#)
- [Release Information | 1289](#)

Syntax

```
packet-loss-priority;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the packet loss priority (PLP) value to be used for the test frames generated in the packet forwarding pipeline. If a value is not configured, then the default value of `low` is used.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[rfc2544-benchmarking](#) | 1347

packet-size (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax](#) | 1289
- [Hierarchy Level](#) | 1289
- [Description](#) | 1290
- [Options](#) | 1290
- [Required Privilege Level](#) | 1290
- [Release Information](#) | 1290

Syntax

```
packet-size bytes;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile-name]
```

Description

Define up to 10 packet sizes that are used sequentially for the test.

Options

bytes Size of the test packet. If you enter multiple packet sizes, you must separate each number with a space.

- **Range:** 64 through 9136 bytes

NOTE: The valid packet sizes are 64, 68, 72, 128, 256, 512, 768, 1024, 1280, 1518, 1522, 1600, 1728, 2496, 3584, 4016, 9104, and 9136 bytes. If you specify a packet size other than the ones listed here as valid sizes, the configuration is saved when you commit the setting and no error message is displayed. However, when you start the test by entering the test services rpm rfc2544-benchmarking test *test-name* start command, an error message is displayed if you configured an invalid packet size in the test profile associated with the test name.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[rfc2544-benchmarking](#) | 1347

passive-monitor-mode

IN THIS SECTION

- [Syntax | 1291](#)
- [Hierarchy Level | 1291](#)
- [Description | 1291](#)
- [Required Privilege Level | 1291](#)
- [Release Information | 1292](#)

Syntax

```
passive-monitor-mode;
```

Hierarchy Level

```
[edit interfaces interface-name (ATM, Fast Ethernet, and Gigabit Ethernet)  
[edit interfaces interface-name unit logical-unit-number] (SONET/SDH)
```

Description

For Asynchronous Transfer Mode (ATM), SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, monitor packet flows from another router. If you include this statement in the configuration, the SONET/SDH interface does not send keepalives or alarms, and does not participate actively on the network.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers | 157](#)
[multiservice-options | 1250](#)

password (Flow Collector File Servers)

IN THIS SECTION

- [Syntax | 1292](#)
- [Hierarchy Level | 1292](#)
- [Description | 1292](#)
- [Options | 1293](#)
- [Required Privilege Level | 1293](#)
- [Release Information | 1293](#)

Syntax

```
password "password";
```

Hierarchy Level

```
[edit services flow-collector destination ftp:uri]
```

Description

Specify the primary and secondary destination FTP server password.

Options

password—FTP server password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

password (Transfer Log File Servers)

IN THIS SECTION

- [Syntax](#) | [1293](#)
- [Hierarchy Level](#) | [1294](#)
- [Description](#) | [1294](#)
- [Options](#) | [1294](#)
- [Required Privilege Level](#) | [1294](#)
- [Release Information](#) | [1294](#)

Syntax

```
password "password";
```


Hierarchy Level

```
[edit services flow-collector transfer-log-archive archive-sites]
```

Description

Specify the primary and secondary destination FTP server password.

Options

password—FTP server password.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

peer-as-billing-template

IN THIS SECTION

- [Syntax](#) | 1295
- [Hierarchy Level](#) | 1295
- [Description](#) | 1295
- [Required Privilege Level](#) | 1295

Syntax

```
peer-as-billing-template;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

Enables the extraction of bandwidth usage information for billing purposes in PIC-based sampling configurations. This capability is supported on routers and applies only to IPv4 and IPv6 traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.4.

RELATED DOCUMENTATION

Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583

persistent-results

IN THIS SECTION

- [Syntax | 1296](#)
- [Hierarchy Level | 1296](#)
- [Description | 1296](#)
- [Required Privilege Level | 1296](#)
- [Release Information | 1296](#)

Syntax

```
persistent-results;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection control-client-name]
```

Description

When enabled, allows to display the old and current tests' results after a network recovery or after TWAMP server reachability when you execute the `show services rpm twamp client probe-results` command. By default, the option is disabled.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches | 694](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

pfe

IN THIS SECTION

- [Syntax | 1297](#)
- [Syntax \(PTX10004, PTX10008, PTX10016 with PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards, and PTX10001-36MR\) | 1298](#)
- [Hierarchy Level | 1298](#)
- [Description | 1298](#)
- [Options | 1298](#)
- [Required Privilege Level | 1299](#)
- [Release Information | 1299](#)

Syntax

```
pfe pfe-instance {
  exception-reporting {
    category category-name; {
      inline-monitoring-instance inline-monitoring-instance;
    }
  }
  forwarding-packages {
    mobility;
  }
  power (off | on);
  tunnel-services;
}
```

Syntax (PTX10004, PTX10008, PTX10016 with PTX10K-LC1201-36CD and PTX10K-LC1202-36MR line cards, and PTX10001-36MR)

```
pfe pfe-instance {
  exception-reporting {
    category category-name; {
      inline-monitoring-instance inline-monitoring-instance;
    }
  }
  forwarding-packages {
    mobility;
  }
  power (off | on);
  temp-perf-throttle-disable;
  temp-volt-reduction-disable;
  tunnel-services;
}
```

Hierarchy Level

```
[edit chassis fpc slot-number],
[edit chassis lcc name fpc slot-number],
[edit chassis member name fpc slot-number]
```

Description

Configure options for the Packet Forwarding Engine (PFE).

Options

- | | |
|----------------------------|--|
| <i>pfe-instance</i> | PFE identifier or PFE number that represents the PFE ASIC ID ranging from 0-4 for PTX10004 and 0-7 for PTX10008. |
| power | Power PFEs on or off |

- off—Power off PFE to save power in scenarios where full system capacity is not required. You can power off the ASICs on PTX10001-36MR device using the set chassis fpc <slot number> pfe <pfe-instance number> power off command.

- on—Power on PFE

By default, PFE is on considered power on. FPC power off supersedes PFE power on.

temp-perf-throttle-disable

Disable temperature based PFE performance throttling. By default, the system detects overheat condition at individual PFE level and gradually reduces the performance of the affected PFE. This results in reduction of power consumption, heat dissipation, PFE operating temperature, and prevents line card shutdown.

- **Default**

By default, automatic performance throttling is enabled.

temp-volt-reduction-disable

Disable temperature based PFE voltage reduction. By default, the system detects the ASIC temperature at individual PFE level that reaches the safe operating limit and dynamically manages the voltage of the DC-DC converters. This results in reduction of power consumption.

- **Default**

By default, temperature voltage reduction is enabled.

tunnel-services

Tunnel services configuration

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

exception-reporting option added in Junos OS Release 21.2R1.

temp-perf-throttle-disable option added in Junos OS Evolved Release 21.4R1.

temp-volt-reduction-disable option added in Junos OS Evolved Release 22.2R1.

exception-reporting option added in Junos OS Evolved Release 22.2R1.

pic-memory-threshold

IN THIS SECTION

- [Syntax | 1300](#)
- [Hierarchy Level | 1300](#)
- [Description | 1300](#)
- [Options | 1300](#)
- [Required Privilege Level | 1300](#)
- [Release Information | 1301](#)

Syntax

```
pic-memory-threshold percentage percentage;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name]
```

Description

Specify a PIC memory usage percentage that triggers a system log warning message.

Options

percentage—PIC memory threshold value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

pop-all-labels

IN THIS SECTION

- [Syntax](#) | 1301
- [Hierarchy Level](#) | 1301
- [Description](#) | 1302
- [Default](#) | 1302
- [Required Privilege Level](#) | 1302
- [Release Information](#) | 1302

Syntax

```
pop-all-labels {  
    required-depth number;  
}
```

Hierarchy Level

```
[edit interfaces interface-name atm-options mpls],  
[edit interfaces interface-name fastether-options mpls],  
[edit interfaces interface-name gigether-options mpls],  
[edit interfaces interface-name sonet-options mpls]
```


Description

For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, removes up to two MPLS labels from incoming IP packets. For passive monitoring on T Series devices, removes up to five MPLS labels from incoming IP packets. For passive monitoring on MX Series routers with MPCs, all labels are popped by default and the `required-depth` statement is ignored.

Except for MX Series routers with MPCs, this statement has no effect on IP packets with more than two MPLS labels, or IP packets with more than five MPLS labels on T Series devices. Packets with MPLS labels cannot be processed by the monitoring PIC; if packets with MPLS labels are forwarded to the monitoring PIC, they are discarded.

The remaining statement is explained separately. See [CLI Explorer](#).

Default

If you omit this statement, the MPLS labels are not removed, and the packet is not processed by the monitoring PIC.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers | 157](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

port (Flow Monitoring)

IN THIS SECTION

- [Syntax | 1303](#)
- [Hierarchy Level | 1303](#)
- [Description | 1303](#)
- [Options | 1303](#)
- [Required Privilege Level | 1304](#)
- [Release Information | 1304](#)

Syntax

```
port port-number;
```

Hierarchy Level

```
[edit forwarding-options accounting name output cflowd hostname],  
[edit forwarding-options monitoring name family inet output cflowd hostname],  
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls) output flow-server  
hostname],  
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
```

Description

Specify the User Datagram Protocol (UDP) port number on the cflowd host system or flow server.

Options

port-number—Any valid UDP port number on the host system.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Flow Aggregation](#) | 577

port (RPM)

IN THIS SECTION

- [Syntax](#) | 1304
- [Junos OS Hierarchy Level](#) | 1305
- [Junos OS Evolved Hierarchy Level](#) | 1305
- [Description](#) | 1305
- [Options](#) | 1305
- [Required Privilege Level](#) | 1305
- [Release Information](#) | 1305

Syntax

```
port number;
```

Junos OS Hierarchy Level

```
[edit services rpm probe-server (tcp | udp)]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name udp]
```

Description

Specify the port number for the RPM probe server.

Options

number—Port number for the probe server. The value can be 7 or 49,160 through 65,535.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

| [Configuring RPM Receiver Servers](#) | 662

port (TWAMP)

IN THIS SECTION

- [Syntax | 1306](#)
- [Junos OS Hierarchy Level | 1306](#)
- [Junos OS Evolved Hierarchy Level | 1306](#)
- [Description | 1306](#)
- [Options | 1306](#)
- [Required Privilege Level | 1307](#)
- [Release Information | 1307](#)

Syntax

```
port number;
```

Junos OS Hierarchy Level

```
[edit services rpm twamp server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring twamp server]
```

Description

TWAMP server listening port.

Options

number—Port number.

- **Range:** 1 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Statement introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

post-cli-implicit-firewall

IN THIS SECTION

- [Syntax](#) | 1307
- [Hierarchy Level](#) | 1308
- [Description](#) | 1308
- [Default](#) | 1308
- [Required Privilege Level](#) | 1308
- [Release Information](#) | 1308

Syntax

```
post-cli-implicit-firewall;
```

Hierarchy Level

```
[edit services rpm twamp]
```

Description

Ensure that the CLI configured (`explicit firewall`) takes precedence over the implicit firewall. The inline TWAMP client or server uses implicit firewall to achieve its functionality.

NOTE: Wrong configuration of CLI firewall can lead to improper functioning of inline TWAMP client or server. After you enable or disable this configuration statement, you must restart the router, or restart remote operation using the command `restart remote-operations`, for the operation to be effective.

When you issue the command `restart remote-operations`, all TWAMP sessions (both client and server) are forced to stop. You must restart all the RPM sessions and all TWAMP sessions (both client and server).

Default

The default for this configuration statement is in disabled status.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

| [Understand Two-Way Active Measurement Protocol](#) | 686

pre-rewrite-tos

IN THIS SECTION

- [Syntax | 1309](#)
- [Hierarchy Level | 1309](#)
- [Description | 1309](#)
- [Required Privilege Level | 1309](#)
- [Release Information | 1309](#)

Syntax

```
pre-rewrite-tos;
```

Hierarchy Level

```
[edit forwarding-options sampling]
```

Description

Preserve prenormalized type-of-service (ToS) value for egress sampled or mirrored packets. This configuration preserves the prerewrite ToS value for all forms of sampling, such as Routing Engine-based sampling, port mirroring, flow monitoring, and so on. This statement is effective for egress sampling only.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

primary-data-record-fields

IN THIS SECTION

- [Syntax](#) | 1310
- [Hierarchy Level](#) | 1310
- [Description](#) | 1310
- [Default](#) | 1311
- [Options](#) | 1311
- [Additional Information](#) | 1312
- [Required Privilege Level](#) | 1312
- [Release Information](#) | 1312

Syntax

```
primary-data-record-fields {  
    name;  
}
```

Hierarchy Level

```
[edit services inline-monitoring template template-name]
```

Description

Configure which IPFIX information elements (IEs) to include in the primary data record, which is exported by the data path and contains the specific properties seen by the data path, either from the sampled packet or the packet notification header or context. With this statement, each inline-monitoring

instance can export IPFIX packets using a different, customized data template, resulting in a custom set of IEs in the IPFIX packets for each instance. The fields are encoded using one or more Common Properties IDs. As per RFC 5473, the primary data record is exported using the IPFIX data template.

Default

datalink-frame-section—Datalink Frame Section (IE 315) is always included in the primary data record.

Options

name Configure which IPFIX information elements (IEs) to include in the primary data record. Per RFC 5102, the Common Properties ID (CPID) is an identifier of a set of common properties that is locally unique per observation domain and transport session.

- cpid-egress-interface-index—CPID Egress Interface Index (IE 137); reports zero value when ingress monitoring
- cpid-forwarding-class-drop-priority—CPID Forwarding Class Drop Priority (IE 137)
- cpid-forwarding-exception-code—CPID Forwarding Exception Code (IE 137)
- cpid-forwarding-nexthop-id—(Junos OS only) CPID Forwarding Nexthop ID (IE 137)
- cpid-ingress-interface-index—CPID Ingress Interface Index (IE 137)
- cpid-underlying-ingress-interface-index—(Junos OS only) CPID Underlying Ingress Interface Index (IE 137); reports a zero value when egress monitoring or if this value is not available in the data path.
- datalink-frame-size—(Junos OS only) Datalink Frame Size (IE 312)
- direction—Flow direction (IE 61)
- egress-interface-snmp-id—Egress Interface SNMP ID (IE 14); reports zero value when ingress monitoring
- ingress-interface-snmp-id—Ingress Interface SNMP ID (IE 10)

Additional Information

Table 137: Values Reported for Aggregated Ethernet (AE) Interfaces

CPID	Ingress on AE Interface	Egress on AE Interface
ingressInterfaceIndex	Logical interface index of the AE interface	Logical interface index of the AE interface
underlyingIngressInterfaceIndex (Junos OS only)	Logical interface index of member link	0
egressInterfaceIndex	0	Logical interface index of member link

Table 138: Values Reported for IRB Interfaces

CPID	Ingress on IRB Interface	Egress on IRB interface
ingressInterfaceIndex	Logical interface index of the IRB interface	Logical interface index of ingress interface
underlyingIngressInterfaceIndex (Junos OS only)	Logical interface index of vlan-bridge encapsulated interface	0
egressInterfaceIndex	0	Logical interface index of vlan-bridge encapsulated interface

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

Statement introduced in Junos Evolved OS Release 22.2R1.

RELATED DOCUMENTATION

[Juniper Resiliency Interface](#) | 409

probe

IN THIS SECTION

- [Syntax](#) | 1313
- [Hierarchy Level](#) | 1315
- [Description](#) | 1315
- [Options](#) | 1315
- [Required Privilege Level](#) | 1315
- [Release Information](#) | 1315

Syntax

```
probe owner {  
  test test-name {  
    data-fill data;  
    data-size size;  
    delegate-probes probes;  
    destination-interface interface-name;  
    destination-port port;  
    dscp-code-point dscp-bits;  
    history-size size;  
    inet6-options source-address ipv6-address;  
    moving-average-size number;  
    next-hop next-hop;  
    offload-type {  
      none;  
      pfe-timestamp;  
    }  
    probe-count count;  
    probe-interval seconds;
```

```

probe-type type;
routing-instance instance-name;
rpm-scale {
    destination {
        interface interface-name.logical-unit-number;
        subunit-cnt subunit-cnt;
    }
    source {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    source-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    target {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    target-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    tests-count tests-count;
}
source-address address;
target (url url | address ipv4-address | inet6-url url | inet6-address ipv6-address);
test-interval interval;
thresholds
{
    egress-time microseconds;
    ingress-time microseconds;
    jitter-egress microseconds;
    jitter-ingress microseconds;
    jitter-rtt microseconds;
    rtt microseconds;
    std-dev-egress microseconds;
    std-dev-ingress microseconds;
    std-dev-rtt microseconds;
}

```

```

        successive-loss count;
        total-loss count;
    }
    traps [trap-names];
    ttl [hop-count];
}
}

```

Hierarchy Level

```
[edit services rpm]
```

Description

Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.

Options

owner—Owner name up to 32 characters in length.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

probe-count

IN THIS SECTION

- [Syntax | 1316](#)
- [Junos OS Hierarchy Level | 1316](#)
- [Junos OS Evolved Hierarchy Level | 1316](#)
- [Description | 1316](#)
- [Options | 1317](#)
- [Required Privilege Level | 1317](#)
- [Release Information | 1317](#)

Syntax

```
probe-count count;
```

Junos OS Hierarchy Level

```
[edit services rpm bgp],  
[edit services rpm probe owner test test-name],  
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the number of probes within a test.

Options

count 1 through 15 for RPM, for TWAMP 1 through 4,294,967,290.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

probe-interval

IN THIS SECTION

- [Syntax | 1318](#)
- [Junos OS Hierarchy Level | 1318](#)
- [Junos OS Evolved Hierarchy Level | 1318](#)
- [Description | 1318](#)
- [Options | 1318](#)

- Required Privilege Level | 1318
- Release Information | 1318

Syntax

```
probe-interval interval;
```

Junos OS Hierarchy Level

```
[edit services rpm bgp],  
[edit services rpm probe owner test test-name],  
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the time to wait between sending packets, in seconds.

Options

interval—Number of seconds, from 1 through 255.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the `[edit services rpm twamp client control-connection control-client-name]` hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

probe-limit

IN THIS SECTION

- [Syntax | 1319](#)
- [Hierarchy Level | 1319](#)
- [Description | 1320](#)
- [Options | 1320](#)
- [Required Privilege Level | 1320](#)
- [Release Information | 1320](#)

Syntax

```
probe-limit limit;
```

Hierarchy Level

```
[edit services rpm]
```

Description

Configure the maximum number of concurrent probes allowed.

Options

limit—Maximum number of concurrent probes allowed.

- **Range:** (MX Series routers only) Starting in Junos OS Release 17.2R2 and 17.3R1, 1 through 2000. In Junos releases earlier than 17.2R1, the range is 1 through 500.
- **Range:** (PTX Series Packet Transport routers only) 1 through 200
- **Range:** (Other platforms) 1 through 500
- **Default:** 100

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Limiting the Number of Concurrent RPM Probes on M, MX, T and PTX Routers and EX Series Switches | 663](#)

probe-server

IN THIS SECTION

- [JUNOS OS | 1321](#)
- [Junos OS Evolved | 1321](#)

- [Junos OS Hierarchy Level | 1321](#)
- [Junos OS Evolved Hierarchy Level | 1322](#)
- [Description | 1322](#)
- [Required Privilege Level | 1322](#)
- [Release Information | 1322](#)

JUNOS OS

```
probe-server {  
    icmp {  
        destination-interface destination-interface;  
    }  
    tcp {  
        destination-interface interface-name;  
        port number;  
    }  
    udp {  
        destination-interface interface-name;  
        port number;  
    }  
}
```

Junos OS Evolved

```
probe-server {  
    icmp;  
    udp {  
        port number;  
    }  
}
```

Junos OS Hierarchy Level

```
[edit services rpm]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm]
```

Description

Enable the RPM server to act as a receiver for the probes, specified per protocol.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

NOTE: The destination-interface statement is not supported on PTX Series routers or for Junos OS Evolved.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Junos OS Evolved Syntax and hierarchy level introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

| [Configuring RPM Receiver Servers](#) | 662

probe-type

IN THIS SECTION

- [Syntax | 1323](#)
- [Junos OS Hierarchy Levels | 1323](#)
- [Junos OS Evolved Hierarchy Level | 1323](#)
- [Description | 1323](#)
- [Options | 1323](#)
- [Required Privilege Level | 1324](#)
- [Release Information | 1324](#)

Syntax

```
probe-type type;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],  
[edit services rpm probe owner test test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the packet and protocol contents of a probe.

Options

type—One of the following probe type values:

- `http-get`—(Junos OS only, not available at the `[edit services rpm bgp]` hierarchy level.) Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.
- `http-metadata-get`—(Junos OS only, not available at the `[edit services rpm bgp]` hierarchy level.) Sends an HTTP get request for metadata to a target URL.
- `icmp-ping`—Sends ICMP echo requests to a target address.
- `icmp-ping-timestamp`—Sends ICMP timestamp requests to a target address.
- `tcp-ping`—(Junos OS only) Sends TCP packets to a target.
- `udp-ping`—(Junos OS only) Sends UDP packets to a target.
- `udp-ping-timestamp`—Sends UDP timestamp requests to a target address.

Required Privilege Level

`system`—To view this statement in the configuration.

`system-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

profiles (RFC 2544 Benchmarking)

IN THIS SECTION

● [Syntax | 1325](#)

- Hierarchy Level | 1325
- Description | 1325
- Options | 1325
- Required Privilege Level | 1326
- Release Information | 1326

Syntax

```
profiles {
  test-profile profile-name {
    test-type (throughput | latency | frame-loss | back-back-frames);
    packet-size bytes;
    step-percent percent;
    bandwidth-kbps kpbs;
  }
}
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking]
```

Description

Configure the test profile to specify attributes, such as the period for the test and the type of test to be performed, for the RFC 2544-based benchmarking test. The test profile is referenced in the test interface to perform a specific type of benchmarking test and compute statistics to describe the performance characteristics of a network interconnecting device.

Options

profiles Define the test profile for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

rate (Interface Services)

IN THIS SECTION

- [Syntax | 1326](#)
- [Hierarchy Level | 1327](#)
- [Description | 1327](#)
- [Options | 1327](#)
- [Required Privilege Level | 1327](#)
- [Release Information | 1327](#)

Syntax

```
rate new-sessions-per-second;
```

Hierarchy Level

```
[edit interfaces interface-name services-options session-limit]
```

Description

Specify the maximum number of new sessions allowed per second on services cards.

Options

rate *new-sessions-per-second* Specify the maximum number of new sessions allowed per second.

- **Range:** 0, which indicates no limit, or greater.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support 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.

rate (Forwarding Options)

IN THIS SECTION

- [Syntax | 1328](#)
- [Hierarchy Level | 1328](#)
- [Description | 1328](#)
- [Options | 1328](#)

- Required Privilege Level | 1329
- Release Information | 1329

Syntax

```
rate number;
```

Hierarchy Level

```
[edit forwarding-options analyzer analyzer-name input],
[edit forwarding-options port-mirroring family (inet | inet6) input],
[edit forwarding-options port-mirroring input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the ratio of the number of packets to be sampled. For example, if you specify a rate of 10, every tenth packet (1 packet out of 10) is sampled.

Native analyzer sessions (that is, the [edit forwarding-options analyzer *analyzer-name* *input*] hierarchy level for MX Series routers) can be configured without specifying input parameters, which means that the instance uses default input values: rate = 1 and maximum-packet-length = 0.

NOTE: The recommended sampling rate for the MX150 is 1000 or greater. If you configure less than 1000, a warning is issued.

Options

number—Denominator of the ratio.

- **Range:** 1 through 16000000 (16M)

For QFX Series switches, the maximum sampling rate for inline Junos Traffic Vision (J-Flow) is 65535.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit forwarding-options analyzer analyzer-name input] hierarchy level for MX Series routers introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Configuring Port Mirroring

Configuring Traffic Sampling

receive-failure-threshold (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1329](#)
- [Hierarchy Level | 1330](#)
- [Description | 1330](#)
- [Required Privilege Level | 1330](#)
- [Release Information | 1330](#)

Syntax

```
receive-failure-threshold;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specifies the failure threshold value of the received test frames.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

receive-options-packets

IN THIS SECTION

- [Syntax | 1331](#)
- [Hierarchy Level | 1331](#)
- [Description | 1331](#)
- [Required Privilege Level | 1331](#)
- [Release Information | 1331](#)

Syntax

```
receive-options-packets;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

When you enable passive monitoring, this statement is required for conformity with cflowd records structure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers](#) | 157

receive-ttl-exceeded

IN THIS SECTION

- [Syntax](#) | 1332
- [Hierarchy Level](#) | 1332

- [Description | 1332](#)
- [Required Privilege Level | 1332](#)
- [Release Information | 1332](#)

Syntax

```
receive-ttl-exceeded;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet]
```

Description

When you enable passive monitoring, this statement is required for conformity with cflowd records structure.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers | 157](#)

reflect-etype

IN THIS SECTION

- [Syntax | 1333](#)
- [Junos OS Hierarchy Level | 1333](#)
- [Junos OS Evolved Hierarchy Level | 1333](#)
- [Description | 1333](#)
- [Options | 1334](#)
- [Required Privilege Level | 1334](#)
- [Release Information | 1334](#)

Syntax

```
reflect-etype ethertype-value;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the EtherType to be used for reflection of the test frames. EtherType is a two-octet field in an Ethernet frame that defines the protocol in the frame payload. This statement is valid only if you configure the test mode to be a reflector. If you do not configure this statement, all EtherTypes are reflected.

Options

ethertype-value Identifier for the EtherType. The EtherType value appears in the Ethernet type field of the packet. It specifies the protocol being transported in the Ethernet frame. For instance, the EtherType for IPv4 is 0x0800. So, if you specify the value as 2048, IPv4 packets are reflected.

- **Range:** 1 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Supported RFC 2544-Based Benchmarking Statements on MX Series Routers | 737](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

reflect-mode

IN THIS SECTION

- [Junos OS Syntax | 1335](#)
- [Junos OS Evolved Syntax | 1335](#)
- [Junos OS Hierarchy Level | 1335](#)

- [Junos OS Evolved Hierarchy Level | 1335](#)
- [Description | 1335](#)
- [Options | 1336](#)
- [Required Privilege Level | 1336](#)
- [Release Information | 1336](#)

Junos OS Syntax

```
reflect-mode (mac-rewrite | mac-swap | no-mac-swap );
```

Junos OS Evolved Syntax

```
reflect-mode (mac-swap | no-mac-swap);
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the reflection mode for the benchmarking test.

Options

mac-rewrite (Junos OS ACX Series routers only) Enable rewriting of the MAC address on the reflected frames. You configure which MAC addresses the software should use with the `source-mac-address` and `destination-mac-address` configuration statements.

mac-swap Swap the source and destination MAC addresses in the test frame. This is the default behavior.

NOTE: In bridge families, when the service type is ELAN, MAC addresses are swapped by default on the reflected frames. When the service type is ELINE, MAC addresses are not swapped by default.

no-mac-swap Do not swap the source and destination MAC addresses in the test frame. The frame is returned to the originator without any modification to the MAC addresses.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[Understanding RFC2544-Based Benchmarking Tests for E-LAN and E-Line Services on MX Series Routers | 733](#)

[rfc2544-benchmarking | 1347](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

refresh-rate (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax | 1337](#)
- [Hierarchy Level | 1337](#)
- [Description | 1337](#)
- [Options | 1337](#)
- [Required Privilege Level | 1338](#)
- [Release Information | 1338](#)

Syntax

```
refresh-rate packets packets seconds seconds;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the refresh rate for transmitting flow template records with version 9 and IPFIX templates for NAT events to the collector, in either packets or seconds.

Options

packets— Number of packets after which templates are sent to the collector.

- **Range:** 1 through 480,000
- **Default:** 4800

seconds—Number of seconds after which templates are sent to the collector

- **Range:** 10 through 600

- **Default:** 600

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

required-depth

IN THIS SECTION

- [Syntax | 1339](#)
- [Hierarchy Level | 1339](#)
- [Description | 1339](#)
- [Options | 1339](#)
- [Required Privilege Level | 1339](#)
- [Release Information | 1340](#)

Syntax

```
required-depth number;
```

Hierarchy Level

```
[edit interfaces interface-name atm-options mpls pop-all-labels],  
[edit interfaces interface-name fastether-options mpls pop-all-labels],  
[edit interfaces interface-name gigether-options mpls pop-all-labels],  
[edit interfaces interface-name sonet-options mpls pop-all-labels]
```

Description

For passive monitoring on ATM, SONET/SDH, Fast Ethernet, and Gigabit Ethernet interfaces only, specify the number of MPLS labels an incoming packet must have for the `pop-all-labels` statement to take effect. For passive monitoring on the MX304 router, MX Series routers with MPCs, and MX series routers with the LC9600 line card, all labels are popped by default and the `required-depth` statement is ignored.

If you include the `required-depth 1` statement, the `pop-all-labels` statement takes effect for incoming packets with one label only. If you include the `required-depth 2` statement, the `pop-all-labels` statement takes effect for incoming packets with two labels only.

Options

number—Number of MPLS labels on incoming IP packets.

- **Range:** 1 through 2 labels.
- **Default:** If you omit this statement, the `pop-all-labels` statement takes effect for incoming packets with one or two labels. The default is equivalent to including the `required-depth [1 2]` statement.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers | 157](#)

[Interfaces Fundamentals for Junos OS](#)

resiliency

IN THIS SECTION

- [Syntax | 1340](#)
- [Hierarchy Level | 1341](#)
- [Description | 1341](#)
- [Default | 1341](#)
- [Options | 1342](#)
- [Additional information | 1342](#)
- [Required Privilege Level | 1342](#)
- [Release Information | 1342](#)

Syntax

```
resiliency {  
  exceptions {  
    forwarding;  
    os;  
    routing;  
  }  
  store {  
    database;
```

```

    file name {
        files number;
        (no-world-readable | world-readable);
        size size;
    }
    fwding-file name{
        files number;
        (no-world-readable | world-readable);
        size size;
    }
}
traceoptions {
    file name {
        files number;
        match;
        (no-world-readable | world-readable);
        size size;
    }
    flag flag;
    no-remote-trace;
}
}

```

Hierarchy Level

[edit system]

Description

Configure the Juniper Resiliency Interface (JRI) on-box collector to detect exceptions from multiple modules and collect those exceptions in a file or a database. If you want to store exceptions data in a file, you need to configure a filename for the file and the size for the file.

Default

The on-box collector is disabled.

Options

The statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Additional information

To fully configure JRI for forwarding exceptions, you must also:

- Configure inline-monitoring services as you normally would, and configure the Juniper-specific IEs with the `primary-data-record-fields` statement at the `[edit services inline-monitoring templates template-name]` hierarchy level.
- Configure the Observation Cloud identifier with the `observation-cloud-id` statement at the `[edit services inline-monitoring]` hierarchy level.
- Configure exception reporting for a particular inline-monitoring instance with the `exception-reporting` statement at the `[edit chassis fpc name pfe name]` hierarchy level.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

[Juniper Resiliency Interface](#) | 409

[inline-monitoring](#) | 1163

[primary-data-record-fields](#) | 1310

[exception-reporting](#) | 1079

retry (Services Flow Collector)

IN THIS SECTION

- [Syntax | 1343](#)
- [Hierarchy Level | 1343](#)
- [Description | 1343](#)
- [Options | 1343](#)
- [Required Privilege Level | 1343](#)
- [Release Information | 1344](#)

Syntax

```
retry number;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the maximum number of attempts the flow collector interface make to transfer log files to the FTP server.

Options

number—Maximum number of transfer retry attempts.

- **Range:** 0 through 10

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

retry-delay

IN THIS SECTION

- [Syntax](#) | [1344](#)
- [Hierarchy Level](#) | [1344](#)
- [Description](#) | [1345](#)
- [Options](#) | [1345](#)
- [Required Privilege Level](#) | [1345](#)
- [Release Information](#) | [1345](#)

Syntax

```
retry-delay seconds;
```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the amount of time the flow collector interface waits between retry attempts.

Options

seconds—Amount of time between transfer retry attempts.

- **Range:** 0 through 60

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

rfc2544

IN THIS SECTION

- [Syntax](#) | [1346](#)
- [Hierarchy Level](#) | [1346](#)
- [Description](#) | [1346](#)
- [Default](#) | [1347](#)
- [Required Privilege Level](#) | [1347](#)
- [Release Information](#) | [1347](#)

Syntax

```

rfc2544 {

    tests{
        test-name test-name {
            test-interface interface-name;
            mode reflect;
            family (ccc | ethernet-switching | inet);
            direction (egress | ingress);
            disable-signature-check;
            in-service;
            ip-swap;
            reflect-etype;
            reflect-mode (mac-swap | no-mac-swap);
            destination-ipv4-address address;
            destination-mac-address mac-address;
            destination-udp-port port-number;
            source-ipv4-address address;
            source-mac-address mac-address;
            source-udp-port port-number;
            service-type service-type;
            udp-tcp-port-swap;
        }
    }
}

```

Hierarchy Level

```
[edit services monitoring]
```

Description

Configure the parameters for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device. The test name that you configure contains details, such as the address family and the test mode, for the test. You can associate the same test profile with multiple test names.

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

Disabled

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 21.1R1.

Support for the destination-mac-address, direction, disable-signature-check, in-service, ip-swap, reflect-etype, reflect-mode, service-type, source-mac-address, and udp-tcp-port-swap configuration statements added in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

| [RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

rfc2544-benchmarking

IN THIS SECTION

- [Junos OS Syntax \(except for SRX devices\)](#) | 1348
- [Junos OS Syntax \(SRX300 and SRX550HM\)](#) | 1349
- [Hierarchy Level](#) | 1349
- [Description](#) | 1350
- [Required Privilege Level](#) | 1350
- [Release Information](#) | 1350

Junos OS Syntax (except for SRX devices)

```

rfc2544-benchmarking {
  profiles {
    test-profile profile-name {
      test-type (throughput | latency | frame-loss | back-back-frames);
      packet-size bytes;
      step-percent percent;
      bandwidth-kbps kpbs;
    }
  }
  tests {
    test-name test-name {
      destination-ipv4-address address;
      destination-mac-address destination-mac-address;
      destination-udp-port port-number;
      direction (egress | ingress);
      disable-signature-check;
      dscp-code-points dscp-code-points;
      family (bridge| inet | ccc | vpls);
      forwarding-class forwarding-class;
      halt-on-prefix-down;
      in-service;
      ip-swap;
      ivlan-cfi ivlan-cfi;
      ivlan-id ivlan-id;
      ivlan-priority ivlan-priority;
      mode (initiate-and-terminate | reflect);
      outer-tag-protocol-id outer-tag-protocol-id;
      ovlan-cfi ovlan-cfi;
      ovlan-id ovlan-id;
      ovlan-priority ovlan-priority;
      packet-loss-priority (high | low | medium-high);
      receive-failure-threshold receive-failure-threshold;
      reflect-etype reflect-etype;
      reflect-mode (mac-rewrite | mac-swap | no-mac-swap);
      service-type (elan | eline);
      skip-arp-iteration;
      source-ipv4-address address;
      source-mac-address source-mac-address;
      source-udp-port port-number;
      test-finish-wait-duration test-finish-wait-duration;
    }
  }
}

```

```

        test-interface interface-name;
        test-iterator-duration test-iterator-duration;
        test-iterator-pass-threshold test-iterator-pass-threshold;
        test-profile test-profile;
        timestamp-format (microseconds | nanoseconds);
        transmit-failure-threshold transmit-failure-threshold;
        udp-tcp-port-swap;
        vlan-cfi vlan-cfi;
        vlan-id vlan-id;
        vlan-priority vlan-priority;
    }
}

```

Junos OS Syntax (SRX300 and SRX550HM)

```

rpm {
    rfc2544-benchmarking {
        tests {
            test-name test-name {
                destination-ipv4-address address;
                destination-udp-port port-number;
                disable-signature-check;
                family inet
                mode reflect;
                source-ipv4-address address;
                source-udp-port port-number;
                test-interface interface-name;
            }
        }
    }
}

```

Hierarchy Level

```
[edit services rpm]
```


Description

Configure the parameters for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device. You must configure a test profile for the initiator (ACX routers only), which specifies the type of test and the manner in which it must be performed, and associate the test profile with a test name. The test name that you configure contains details, such as the address family and the test mode (initiate-and-terminate), for the test. You can associate the same test profile with multiple test names. For a reflector, you need only configure a test name and its associated details, and set the test mode to reflect.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[show services rpm rfc2544-benchmarking | 1791](#)

[show services rpm rfc2544-benchmarking test-id | 1800](#)

routing-instance (RPM)

IN THIS SECTION

● [Syntax | 1351](#)

● [Junos OS Hierarchy Level | 1351](#)

- Junos OS Evolved Hierarchy Level | 1351
- Description | 1351
- Options | 1351
- Required Privilege Level | 1352
- Release Information | 1352

Syntax

```
routing-instance instance-name;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name]
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the routing instance used by the probes. The routing instance is also applicable for control connection.

NOTE: The media interface from where the TWAMP control and test or data packets arrive and exit the si- logical interface must be a part of the same routing instance.

Options

instance-name—Routing instance configured at the [edit routing-instance] hierarchy level.

- **Default:** Internet (IPv4) routing table inet.0.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

routing-instance (cflowd)

IN THIS SECTION

- [Syntax | 1353](#)
- [Hierarchy Level | 1353](#)
- [Description | 1353](#)
- [Options | 1353](#)
- [Required Privilege Level | 1353](#)
- [Release Information | 1353](#)

Syntax

```
routing-instance instance-name;
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output flow-server hostname]
```

Description

Configure a non-default VPN routing and forwarding (VRF) instance through which flow collectors can be reachable for inline flow monitoring. You cannot configure a flow collector to be reachable through non-default VRF instances for version 5 and version 8 flows. You must configure the routing instance to be a VRF instance by including the `instance-type vrf` statement at the `[edit routing-instances instance-name]` hierarchy level.

Options

instance-name—Name of a routing instance that has been configured at the `[edit routing-instance]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

RELATED DOCUMENTATION

| *Directing Traffic Sampling Output to a Server Running the cflowd Application*

routing-instance-list (TWAMP)

IN THIS SECTION

- [Syntax | 1354](#)
- [Hierarchy Level | 1354](#)
- [Description | 1354](#)
- [Options | 1355](#)
- [Required Privilege Level | 1355](#)
- [Release Information | 1355](#)

Syntax

```
routing-instance-list {  
    instance-name {  
        port number;  
    }  
}
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

Configure the Two-Way Active Measurement Protocol (TWAMP) servers on specific routing instances, instead of associating the TWAMP server at the system-level to apply to all routing instances configured on a router. The default routing instance is Internet routing table `inet.0`. If you do not specify a routing instance, the TWAMP probe applies to all routing instances. To apply the TWAMP probe to only the default routing instance, you must explicitly set the value of `instance-name` to `default`. If an interface is not part of any routing instance, the default port is used for TWAMP probes. You can configure up to 100 routing instances for a TWAMP server.

Options

instance-name—Name of the routing instance, a maximum of 31 characters.

number—Port number.

- **Range:** 1 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

routing-instances

IN THIS SECTION

- [Syntax](#) | 1356
- [Hierarchy Level](#) | 1356
- [Description](#) | 1356
- [Options](#) | 1356
- [Required Privilege Level](#) | 1356
- [Release Information](#) | 1356

Syntax

```
routing-instances instance-name;
```

Hierarchy Level

```
[edit services rpm bgp],  
[edit services rpm bgp logical-system logical-system-name]
```

Description

Specify the routing instance used by the probes.

Options

instance-name—A routing instance configured at the [edit routing-instances] hierarchy level.

- **Default:** Internet routing table inet.0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.6.

RELATED DOCUMENTATION

| [Configuring BGP Neighbor Discovery Through RPM](#) | 671

rpm (Interfaces)

IN THIS SECTION

- [Syntax | 1357](#)
- [Hierarchy Level | 1357](#)
- [Description | 1357](#)
- [Options | 1357](#)
- [Required Privilege Level | 1358](#)
- [Release Information | 1358](#)

Syntax

```
rpm (client client | server server | twamp-client twamp-client | twamp-server twamp-server);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number]
```

Description

Associate an RPM or TWAMP client (router or switch that originates RPM or TWAMP probes) or RPM or TWAMP server with a specified interface.

NOTE: The TWAMP client is applicable only for si- interfaces.

Options

client—Identifier for RPM client router or switch.

server—Identifier for RPM server.

twamp-client—Identifier for RPM TWAMP client router.

twamp-server—Identifier for RPM TWAMP server.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.1.

RELATED DOCUMENTATION

[Configuring RPM Timestamping on MX, M, T, and PTX Series Routers and EX Series Switches](#) | 663

rpm (Services)

IN THIS SECTION

- [Syntax \(Junos OS, except for QFX10000 switches, EX9200 switches, and SRX devices\)](#) | 1359
- [Syntax \(Junos OS, for EX9200 and QFX10000 switches\)](#) | 1364
- [Syntax \(Junos OS, for SRX300 and SRX550HM devices\)](#) | 1368
- [Syntax \(Junos OS Evolved\)](#) | 1369
- [Junos OS Hierarchy Level](#) | 1370
- [Junos OS Evolved Hierarchy Level](#) | 1370
- [Description](#) | 1370
- [Required Privilege Level](#) | 1371
- [Release Information](#) | 1371

Syntax (Junos OS, except for QFX10000 switches, EX9200 switches, and SRX devices)

```

rpm {
  bgp {
    data-fill data;
    data-size size;
    destination-port port;
    history-size size;
    logical-system logical-system-name [routing-instances routing-instance-name];
    moving-average-size number;
    probe-count count;
    probe-interval seconds;
    probe-type type;
    routing-instances instance-name;
    test-interval interval;
    ttl ttl;
  }
  probe owner {
    delegate-probes;
    test test-name {
      data-fill data;
      data-size size;
      destination-interface interface-name;
      destination-port port;
      dscp-code-points dscp-bits;
      hardware-timestamp;
      history-size size;
      inet6-options {
        source-address source-address;
      }
      moving-average-size number;
      next-hop next-hop;
      one-way-hardware-timestamp;
      probe-count count;
      probe-interval seconds;
      probe-type type;
      routing-instance instance-name;
      rpm-scale {
        target {
          address-base address-base;
          count count;

```

```

        step step;
    }
    target-inet6 {
        address-base address-base;
        count count;
        step step;
    }
    source {
        address-base address-base;
        count count;
        step step;
    }
    source-inet6 {
        address-base address-base;
        count count;
        step step;
    }
    destination {
        interface interface;
        subunit-cnt subunit-cnt;
    }
    tests-count tests-count;
}
source-address address;
target (url url | address address);
test-interval interval;
thresholds thresholds;
traps traps;
ttl tll;
}
}
probe-server {
    icmp {
        destination-interface destination-interface;
    }
    tcp {
        destination-interface interface-name;
        port number;
    }
    udp {
        destination-interface interface-name;
        port number;
    }
}

```

```

}
probe-limit limit;
rfc2544-benchmarking {
    profiles {
        test-profile profile-name {
            test-type (throughput | latency | frame-loss | back-back-frames);
            packet-size bytes;
            step-percent percent;
            bandwidth-kbps kpbs;
        }
    }
}
tests {
    test-name test-name {
        destination-ipv4-address address;
        destination-mac-address destination-mac-address;
        destination-udp-port port-number;
        direction (egress | ingress);
        disable-signature-check;
        dscp-code-points dscp-code-points;
        family (bridge| inet | ccc);
        forwarding-class forwarding-class;
        halt-on-prefix-down;
        in-service;
        ip-swap;
        ivlan-cfi ivlan-cfi;
        ivlan-id ivlan-id;
        ivlan-priority ivlan-priority;
        mode reflect;
        outer-tag-protocol-id outer-tag-protocol-id;
        ovlan-cfi ovlan-cfi;
        ovlan-id ovlan-id;
        ovlan-priority ovlan-priority;
        packet-loss-priority (high | low | medium-high);
        receive-failure-threshold receive-failure-threshold;
        reflect-etype reflect-etype;
        reflect-mode (mac-rewrite | mac-swap | no-mac-swap);
        reflector-port reflector-port;
        service-type (elan | eline);
        skip-arp-iteration;
        source-ipv4-address address;
        source-mac-address source-mac-address;
        source-udp-port port-number;
        test-finish-wait-duration test-finish-wait-duration;
    }
}

```

```

test-interface interface-name;
test-iterator-duration test-iterator-duration;
test-iterator-pass-threshold test-iterator-pass-threshold;
test-profile test-profile;
timestamp-format (microseconds | nanoseconds);
transmit-failure-threshold transmit-failure-threshold;
udp-tcp-port-swap;
vlan-cfi vlan-cfi;
vlan-id vlan-id;
vlan-priority vlan-priority;
}
}
}
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;
}
twamp {
    client {
        control-connection name {
            authentication-mode none;
            control-type (managed | light);
            destination-interface destination-interface;
            destination-port destination-port;
            history-size history-size;
            moving-average-size moving-average-size;
            persistent-results;
            routing-instance routing-instance;
            target-address target-address;
            tcp-keepcnt count;
            tcp-keepidle seconds;
            tcp-keepintvl seconds;
            test-count test-count;
            test-interval seconds;
            traps {
                control-connection-closed;
                test-iteration-done;
            }
            test-session name {
                data-fill-with zeros;
            }
        }
    }
}

```

```

data-size data-size;
dscp-code-points dscp-code-points;
probe-count probe-count;
probe-interval seconds;
source-address source-address;
target-address target-address local-link IPv6-link-local-interface-
name;

thresholds {
    egress-time microseconds;
    ingress-time microseconds;
    jitter-egress microseconds;
    jitter-ingress microseconds;
    jitter-rtt microseconds;
    max-rtt microseconds;
    rtt microseconds;
    std-dev-egress microseconds;
    std-dev-ingress microseconds;
    std-dev-rtt microseconds;
    successive-loss successive-loss;
    total-loss total-loss;
}
traps {
    egress-jitter-exceeded;
    egress-std-dev-exceeded;
    egress-time-exceeded;
    ingress-jitter-exceeded;
    ingress-std-dev-exceeded;
    ingress-time-exceeded;
    jitter-exceeded;
    max-rtt-exceeded;
    probe-failure;
    rtt-exceeded;
    std-dev-exceeded;
    test-completion;
    test-failure;
}
ttl hop-count;
}
}
}
post-cli-implicit-firewall;
server {
    authentication-key-chain name {

```

```

        key-id name {
            secret secret;
        }
    }
    authentication-mode <authenticated> <control-only-encrypted> <encrypted>
<none>;
    client-list {
        address address <routing-instance [instance-name...]>;
    }
    max-connection-duration hours;
    maximum-connections maximum-connections;
    maximum-connections-per-client maximum-connections-per-client;
    maximum-sessions maximum-sessions;
    maximum-sessions-per-connection maximum-sessions-per-connection;
    port port;
    routing-instance-list name {
        port port;
    }
    server-inactivity-timeout minutes;
    tcp-keepcnt count;
    tcp-keepidle seconds;
    tcp-keepintvl seconds;
    }
}
}

```

Syntax (Junos OS, for EX9200 and QFX10000 switches)

```

rpm {
    bgp {
        data-fill data;
        data-size size;
        destination-port port;
        history-size size;
        logical-system logical-system-name [routing-instances routing-instance-name];
        moving-average-size number;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instances instance-name;
        test-interval interval;
    }
}

```

```

    ttl t11;
}
probe owner {
    delegate-probes;
    test test-name {
        data-fill data;
        data-size size;
        destination-interface interface-name;
        destination-port port;
        dscp-code-points dscp-bits;
        hardware-timestamp;
        history-size size;
        inet6-options {
            source-address source-address;
        }
        moving-average-size number;
        next-hop next-hop;
        one-way-hardware-timestamp;
        probe-count count;
        probe-interval seconds;
        probe-type type;
        routing-instance instance-name;
        rpm-scale {
            target {
                address-base address-base;
                count count;
                step step;
            }
            target-inet6 {
                address-base address-base;
                count count;
                step step;
            }
            source {
                address-base address-base;
                count count;
                step step;
            }
            source-inet6 {
                address-base address-base;
                count count;
                step step;
            }
        }
    }
}

```



```

        destination {
            interface interface;
            subunit-cnt subunit-cnt;
        }
        tests-count tests-count;
    }
    source-address address;
    target (url url | address address);
    test-interval interval;
    thresholds thresholds;
    traps traps;
    ttl tll;
}
}
probe-server {
    icmp {
        destination-interface destination-interface;
    }
    tcp {
        destination-interface interface-name;
        port number;
    }
    udp {
        destination-interface interface-name;
        port number;
    }
}
probe-limit limit;
twamp {
    client {
        control-connection name {
            authentication-mode none;
            destination-interface destination-interface;
            destination-port destination-port;
            history-size history-size;
            moving-average-size moving-average-size;
            persistent-results;
            routing-instance routing-instance;
            target-address target-address;
            tcp-keepcnt count;
            tcp-keepidle seconds;
            tcp-keepintvl seconds;
            test-count test-count;

```

```

test-interval    seconds;
traps {
    control-connection-closed;
    test-iteration-done;
}
test-session name {
    data-fill-with zeros;
    data-size    data-size;
    dscp-code-points dscp-code-points;
    probe-count  probe-count;
    probe-interval seconds;
    target-address target-address;
    thresholds {
        egress-time microseconds;
        ingress-time microseconds;
        jitter-egress microseconds;
        jitter-ingress microseconds;
        jitter-rtt microseconds;
        max-rtt microseconds;
        rtt microseconds;
        std-dev-egress microseconds;
        std-dev-ingress microseconds;
        std-dev-rtt microseconds;
        successive-loss successive-loss;
        total-loss total-loss;
    }
    traps {
        egress-jitter-exceeded;
        egress-std-dev-exceeded;
        egress-time-exceeded;
        ingress-jitter-exceeded;
        ingress-std-dev-exceeded;
        ingress-time-exceeded;
        jitter-exceeded;
        max-rtt-exceeded;
        probe-failure;
        rtt-exceeded;
        std-dev-exceeded;
        test-completion;
        test-failure;
    }
    ttl hop-count;
}

```

```

    }
  }
  post-cli-implicit-firewall;
  server {
    authentication-key-chain name {
      key-id name {
        secret secret;
      }
    }
    authentication-mode <authenticated > <control-only-encrypted> <encrypted >
    <none>;
    client-list {
      address address <routing-instance [instance-name...]>;
    }
    max-connection-duration hours;
    maximum-connections maximum-connections;
    maximum-connections-per-client maximum-connections-per-client;
    maximum-sessions maximum-sessions;
    maximum-sessions-per-connection maximum-sessions-per-connection;
    port port;
    routing-instance-list name {
      port port;
    }
    server-inactivity-timeout minutes;
    tcp-keepcnt count;
    tcp-keepidle seconds;
    tcp-keepintvl seconds;
  }
}

```

Syntax (Junos OS, for SRX300 and SRX550HM devices)

```

rpm {
  rfc2544-benchmarking {
    tests {
      test-name test-name {
        destination-ipv4-address address;
        destination-udp-port port-number;
        disable-signature-check;
        family inet

```

```

        mode reflect;
        source-ipv4-address address;
        source-udp-port port-number;
        test-interface interface-name;
        test-iterator-duration test-iterator-duration;
    }
}
}
}

```

Syntax (Junos OS Evolved)

```

rpm {
  owner name {
    test test-name {
      data-fill data;
      data-size size;
      destination-port port;
      dscp-code-points dscp-bits;
      history-size size;
      moving-average-size number;
      offload-type {
        none;
        pfe-timestamp;
      }
      probe-count count;
      probe-interval seconds;
      probe-type type;
      routing-instance instance-name;
      source-address address;
      target (url url | address address);
      test-interval interval;
      thresholds thresholds;
      traps traps;
      ttl hop-count
    }
  }
}

```

Junos OS Hierarchy Level

```
[edit services]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring]
```

Description

Configure real-time performance monitoring (RPM). RPM enables you to configure active probes to track and monitor traffic. Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets. RPM provides Management Information Base (MIB) support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

Starting in Junos OS Evolved Release 20.1R1, you can configure RPM probes. For Junos OS Evolved, RPM is configured at the `[edit services monitoring rpm]` hierarchy level. The scope of support is limited to:

- Probe generation and reception (client) as well as reflection (server) for the following RPM probe types:
 - icmp-ping
 - icmp-timestamp
 - udp-ping
 - udp-timestamp
- Probe history management
- Reporting through syslog only

Starting in Junos OS Evolved 20.3R1, you can configure TWAMP probes. Starting in Junos OS Evolved 21.1R1, you can configure RFC 2544 benchmarking tests. For Junos OS Evolved, TWAMP is configured at the `[edit services monitoring twamp]` hierarchy level and RFC 2544 benchmarking tests are configured at the `[edit services monitoring rfc2544]` hierarchy level.

Starting in Junos OS Evolved Release 21.2R1, reporting through SNMP MIB objects is supported for RPM.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

traps option introduced in Junos OS Evolved Release 21.2R1.

source-address option and the local-link sub-option of the target-address option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

Release History Table

Release	Description
21.4R1	Starting in Junos OS Release 21.4R1, you can configure RPM and TWAMP probes for EX9200 switches.
21.4R1	Starting in Junos OS Release 21.4R1, for TWAMP Light test-sessions on MX and PTX routers, you can configure IPv6 link-local target addresses and their corresponding source addresses.
21.3R1	Starting in Junos OS 21.3R1, you can configure RPM and TWAMP probes for QFX10000 Series switches.
21.2R1 Evo	Starting in Junos OS Evolved Release 21.2R1, reporting through SNMP MIB objects is supported for RPM.
21.1R1 Evo	Starting in Junos OS Evolved 21.1R1, you can configure RFC 2544 benchmarking tests.
20.3R1 Evo	Starting in Junos OS Evolved 20.3R1, you can configure TWAMP probes.
20.1R1 Evo	Starting in Junos OS Evolved Release 20.1R1, you can configure RPM probes. For Junos OS Evolved, RPM is configured at the [edit services monitoring rpm] hierarchy level.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#) | 646

[Configuring BGP Neighbor Discovery Through RPM](#) | 671

rpm-scale

IN THIS SECTION

- [Syntax](#) | 1372
- [Hierarchy Level](#) | 1373
- [Description](#) | 1373
- [Options](#) | 1373
- [Required Privilege Level](#) | 1374
- [Release Information](#) | 1374

Syntax

```
rpm-scale {  
  destination {  
    interface interface-name.logical-unit-number;  
    subunit-cnt subunit-cnt;  
  }  
  source {  
    address-base ipv4-address-base;  
    count ipv4-count;  
    step ipv4-step;  
  }  
  source-inet6 {  
    address-base ipv6-address-base;  
    count ipv6-count;  
    step ipv6-step;  
  }  
  target {
```

```

        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    target-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    tests-count tests-count;
}

```

Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Description

Configure the generation of multiple IPv4 RPM tests for a probe owner. Starting in Junos OS Release 18.2R1, you can also configure the generation of multiple IPv6 RPM tests for a probe owner. Tests are generated for multiple combinations of source and target addresses, which are incremented based on your configuration. Tests are first generated for all the source addresses with the initial target address, then tests are generated for all the source addresses with the next available target address, and so on.

Options

<i>interface-name.logical-unit-number</i>	The services interface that is generating RPM probes and the logical unit number that is used for the first test that is generated.
<i>ipv4-address-base</i>	The IPv4 source or target address that is incremented to generate the addresses used in the RPM tests.
<i>ipv4-count</i>	The maximum number of IPv4 source or target addresses to use for the generated RPM tests.
<i>ipv4-step</i>	The amount to increment the IPv4 source or target address for each generated RPM test.
<i>ipv6-address-base</i>	The IPv6 source or target address that is incremented to generate the addresses used in the RPM tests.

<i>ipv6-count</i>	The maximum number of IPv6 source or target addresses to use for the generated RPM tests.
<i>ipv6-step</i>	The amount to increment the IPv6 source or target address for each generated RPM test.
<i>subunit-cnt</i>	The maximum number of logical units used by the services interface in the generated tests. The first generated test uses the logical unit specified in the <i>interface-name.logical-unit-number</i> option, and each successive test increments the logical unit number by one. Once the maximum number of logical units has been used, the next generated test cycles back to the logical unit that was used in the first test.
<i>tests-count</i>	<p>The maximum number of RPM tests to generate. This number must be less than or equal to the number of generated source addresses multiplied by the number of generated target addresses.</p> <ul style="list-style-type: none"> • Range: 1 through 500,000 for probes generated on an MS-MPC or MS-MIC. 1 through 2,000 for probes generated on the Routing Engine.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 17.4R1.

RELATED DOCUMENTATION

| [Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | 651

rpm-tracking

IN THIS SECTION

- [Syntax | 1375](#)
- [Hierarchy Level | 1375](#)
- [Description | 1375](#)
- [Options | 1376](#)
- [Required Privilege Level | 1377](#)
- [Release Information | 1377](#)

Syntax

```
rpm-tracking {  
    route destination-prefix {  
        metric metric;  
        next-hop next-hop {;  
            rpm-probe name rpm-test test-name;  
        }  
        preference preference;  
        tag tag;  
    }  
}
```

Hierarchy Level

```
[edit routing-instances name routing-options],  
[edit routing-options]
```

Description

RPM-tracked static routes are coupled with a given RPM test instance. Routes can be installed or removed according to the results of the given RPM test. When installed, these routes are automatically

given a preference of 1, and so are preferred over static routes that may already exist with the same prefix.

If the RPM test result is “success,” then all the RPM-tracked static routes that match the probe owner and test name of the successful test are added to the routing table. If the test result is “failure,” then all the RPM-tracked static routes that match the probe owner and test name of the failed test are removed, if present, from the routing table.

RPM route tracking supports both IPv4 and IPv6 routes. RPM-tracked static routes are configured individually; wildcards, ranges, and regular expressions are not supported.

Options

route <i>destination-prefix</i>	(Required) Configure an RPM-tracked static route. Must be a IPv4 or IPv6 destination prefix.
next-hop <i>next-hop</i>	(Required) Configure a next-hop address. Must be a IPv4 or IPv6 address. You can configure up to 16 multiple paths (next-hops) for any given RPM-tracked static route (RPM-tracked static routes with multiple next-hops can also be configured inside a routing instance).
metric <i>metric</i>	(Optional) Set the route metric for the destination prefix. The route with the lowest metric is active in the routing table. <ul style="list-style-type: none"> • Default: 1 • Range: 1 through 16.
preference <i>preference</i>	(Optional) Set the route preference value for the destination prefix. The route with the lowest preference value is active in the routing table. Qualified next hops are not supported with RPM-tracked static routes. Hence, the setting of the preference, metric, and tag options apply only to the RPM-tracked static route and not to the related next hops. <ul style="list-style-type: none"> • Default: 1
tag <i>tag</i>	(Optional) Set the route tag value for the destination prefix. The route with the lowest tag value is active in the routing table. <ul style="list-style-type: none"> • Default: 0
rpm-probe <i>name</i>	(Required) Must be a valid RPM probe owner from the [edit services rpm] configuration hierarchy.

rpm-test *test-name* (Required) Must be a valid RPM test name from the [edit services rpm] configuration hierarchy.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Support introduced in Junos OS Release 18.4 R1.

next-hop option added in Junos OS Release 19.1R1.

preference and tag options added in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[show route rpm-tracking | 1625](#)

run-length

IN THIS SECTION

- [Syntax | 1378](#)
- [Hierarchy Level | 1378](#)
- [Description | 1378](#)
- [Options | 1378](#)
- [Required Privilege Level | 1378](#)
- [Release Information | 1378](#)

Syntax

```
run-length number;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring input],
[edit forwarding-options port-mirroring instance port-mirroring-instance-name input],
[edit forwarding-options port-mirroring family (inet|inet6) input],
[edit forwarding-options sampling input],
[edit forwarding-options sampling instance instance-name input]
```

Description

Set the number of samples following the initial trigger event. The configuration enables you to sample packets following those already being sampled.

NOTE: The run-length statement is not supported when you configure inline flow monitoring (by including the inline-jflow statement at the [edit forwarding-options sampling instance *instance-name* family (inet | inet6) output] hierarchy level).

Options

number—Number of samples.

- **Range:** 0 through 20
- **Default:** 0

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Applying Forwarding Table Filters

[Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

sample-once

IN THIS SECTION

- [Syntax](#) | 1379
- [Hierarchy Level](#) | 1379
- [Description](#) | 1379
- [Required Privilege Level](#) | 1380
- [Release Information](#) | 1380

Syntax

```
sample-once;
```

Hierarchy Level

```
[edit forwarding-options sampling]
```

Description

Explicitly sample a packet for active monitoring only once. Setting this option avoids duplication of packets in cases where sampling is enabled at both the ingress and egress interfaces and simplifies analysis of the sampled traffic.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

RELATED DOCUMENTATION

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

sampling (Forwarding Options)

IN THIS SECTION

- [Syntax](#) | 1380
- [Hierarchy Level](#) | 1383
- [Description](#) | 1383
- [Required Privilege Level](#) | 1383
- [Release Information](#) | 1383

Syntax

```
sampling {  
  disable;  
  family (inet | inet6 | mpls | vpls) {  
    disable;  
    output {  
      aggregate-export-interval seconds;  
      extension-service service-name;  
      file {
```

```

        disable;
        filename filename;
        files number;
        size bytes;
        (stamp | no-stamp);
        (world-readable | no-world-readable);
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-server hostname {
        aggregation {
            autonomous-system;
            destination-prefix;
            protocol-port;
            source-destination-prefix {
                caida-compliant;
            }
            source-prefix;
        }
        autonomous-system-type (origin | peer);
        (local-dump | no-local-dump);
        port port-number;
        source-address address;
        version format;
        version9 {
            template template-name;
        }
    }
    interface interface-name {
        engine-id number;
        engine-type number;
        source-address address;
    }
}

input {
    max-packets-per-second number;
    maximum-packet-length bytes;
    rate number;
    run-length number;
}

instance instance-name {
    disable;

```



```

family (bridge | inet | inet6 | mpls | vpls) {
    disable;
    output {
        aggregate-export-interval seconds;
        extension-service service-name;
        flow-server hostname {
            aggregation {
                autonomous-system;
                destination-prefix;
                protocol-port;
                source-destination-prefix {
                    caida-compliant;
                }
                source-prefix;
            }
            autonomous-system-type (origin | peer);
            dscp dscp-value;
            forwarding-class class-name;
            (local-dump | no-local-dump);
            port port-number;
            source-address address;
            version format;
            version9 {
                template template-name;
            }
            version-ipfix {
                template template-name;
            }
        }
        inline-jflow {
            source-address address;
            flow-export-rate rate;
        }
        interface interface-name {
            engine-id number;
            engine-type number;
            source-address address;
        }
    }
}

input {
    max-packets-per-second number;
    maximum-packet-length bytes;
}

```

```

        rate number;
        run-length number;
    }
}
pre-rewrite-tos;
sample-once;
traceoptions {
    no-remote-trace;
    file filename <files number> <size bytes> <match expression> <world-readable | no-world-
readable>;
}
}

```

Hierarchy Level

[edit forwarding-options]

Description

Configure traffic sampling.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

Applying Forwarding Table Filters

Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format

Directing Traffic Sampling Output to a Server Running the cflowd Application

[Configuring Port Mirroring](#)[Tracing Traffic-Sampling Operations](#)[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

sampling (Interfaces)

IN THIS SECTION

- [Syntax | 1384](#)
- [Hierarchy Level | 1384](#)
- [Description | 1384](#)
- [Options | 1385](#)
- [Required Privilege Level | 1385](#)
- [Release Information | 1385](#)

Syntax

```
sampling direction;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number  
family inet]
```

Description

Configure the direction of traffic to be sampled.

Options

direction can be one of the following:

input—Configure at least one expected ingress point.

output—Configure at least one expected egress point.

input output—On a single interface, configure at least one expected ingress point and one expected egress point.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Junos OS Services Interfaces Library for Routing Devices](#)

[Configuring Flow Monitoring | 5](#)

sampling-instance

IN THIS SECTION

- [Syntax | 1386](#)
- [Hierarchy Level | 1386](#)
- [Description | 1386](#)
- [Required Privilege Level | 1386](#)
- [Release Information | 1386](#)

Syntax

```
sampling-instance instance-name;
```

Hierarchy Level

```
[edit chassis fpc slot-number],  
[edit chassis lcc number fpc slot-number] (Routing Matrix),  
[edit chassis member member-number fpc slot slot-number]
```

Description

Associate a defined sampling instance with a specific FPC, MPC, or DPC for active sampling instances configured at the [edit forwarding-options sampling] hierarchy level.

For M120 routers with FEB, this statement must also be configured under [edit chassis feb *slot-number*], in addition to the [edit forwarding-options sampling] hierarchy level.

In a two-member MX Series Virtual Chassis, the primary router (member 0) uses FPC slot numbers 0 through 11 with no offset; the backup router (member 1) uses FPC slot numbers 12 through 23, with an offset of 12.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.6.

Support at the [edit chassis member *member-number* fpc slot *slot-number*] hierarchy level introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

[Associating Sampling Instances for Active Flow Monitoring with a Specific FPC, MPC, or DPC | 132](#)
[Inline Flow Monitoring for Virtual Chassis Overview](#)

server (Junos OS)

IN THIS SECTION

- [Syntax | 1387](#)
- [Hierarchy Level | 1388](#)
- [Description | 1388](#)
- [Options | 1388](#)
- [Required Privilege Level | 1388](#)
- [Release Information | 1388](#)

Syntax

```
server {  
    authentication-mode none;  
    client-list list-name {  
        address address <routing-instance [instance-name...]>;  
    }  
    light {  
        port number;  
    }  
    max-connection-duration hours;  
    maximum-connections count;  
    maximum-connections-per-client count;  
    maximum-sessions count;  
    maximum-sessions-per-connection count;  
    port number;  
    server-inactivity-timeout seconds;  
    tcp-keepcnt count;  
    tcp-keepidle seconds;  
    tcp-intvl seconds;  
}
```

Hierarchy Level

```
[edit services rpm twamp]
```

Description

TWAMP server configuration settings.

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

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

light option added in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

server (Junos OS Evolved)

IN THIS SECTION

● [Syntax](#) | 1389

- [Hierarchy Level | 1389](#)
- [Description | 1389](#)
- [Options | 1390](#)
- [Required Privilege Level | 1390](#)
- [Release Information | 1390](#)

Syntax

```
server {
    light {
        port number;
    }
    managed {
        client-limit limit;
        client-list list-name {
            address address <routing-instance [instance-name...]>;
        }
        control-inactivity-timeout seconds;
        control-maximum-duration seconds;
        control-per-client-limit number;
        port number;
        test-per-client-limit limit;
    }
}
```

Hierarchy Level

```
[edit services monitoring twamp]
```

Description

Configure the Two-Way Active Measurement Protocol (TWAMP) server. The server listens in on all routing instances.

Options

client-limit <i>limit</i>	Maximum number of TWAMP clients Range: 0 to 1000 clients
control-inactivity-timeout <i>seconds</i>	Inactivity timeout on control connection Range: 0 to 86400 seconds
control-maximum-duration <i>seconds</i>	Hard limit on control connection duration Range: 0 to 86400 seconds
control-per-client-limit <i>number</i>	Maximum number of TWAMP control connections per client Range: 0 to 1000 control connections per client
test-per-client-limit <i>limit</i>	Maximum number of TWAMP test sessions per client Range: 0 to 1000 test sessions per 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

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

| [Understand Two-Way Active Measurement Protocol](#) | 686

server-inactivity-timeout

IN THIS SECTION

- [Syntax | 1391](#)
- [Hierarchy Level | 1391](#)
- [Description | 1391](#)
- [Options | 1391](#)
- [Required Privilege Level | 1392](#)
- [Release Information | 1392](#)

Syntax

```
server-inactivity-timeout minutes;
```

Hierarchy Level

```
[edit services rpm twamp server]
```

Description

The maximum time the Two-Way Active Measurement Protocol (TWAMP) server has to finish the TWAMP control protocol negotiation.

Options

minutes Number of minutes the TWAMP server has to finish the TWAMP control protocol negotiation.

- **Default:** 15 minutes
- **Range:** 1 through 30 minutes

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 11.1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

service-port

IN THIS SECTION

- [Syntax](#) | 1392
- [Hierarchy Level](#) | 1393
- [Description](#) | 1393
- [Options](#) | 1393
- [Required Privilege Level](#) | 1393
- [Release Information](#) | 1393

Syntax

```
service-port port-number;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

Identify the User Datagram Protocol (UDP) port number for control protocol requests.

Options

port-number—Port number for control protocol request messages.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

service-type

IN THIS SECTION

- [Syntax](#) | 1394
- [Hierarchy Level](#) | 1394
- [Description](#) | 1394
- [Options](#) | 1394

- Required Privilege Level | 1394
- Release Information | 1395

Syntax

```
service-type (elan | eline) ;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Configure the service under test. Possible values are `elan` and `eline`. This statement is applicable only for the bridge family, the ethernet-switching family, or when the `mode` is configured as `reflect`. When the service type is `elan`, MAC addresses are swapped by default on the reflected frames. The `no-mac-swap` option is not supported for this service type. When the service type is `eline`, MAC addresses are not swapped by default in the reflected frames. Use the `mac-swap` option to swap the addresses.

NOTE: When you configure Layer 2 reflection, you can specify the service type under test as `eline` if you want to simulate an E-Line service using bridge encapsulation.

Options

- | | |
|--------------|----------------------------------|
| elan | Specify the ELAN service type. |
| eline | Specify the E-Line service type. |

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[rfc2544-benchmarking | 1347](#)

[rfc2544 | 1345](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

services-options

IN THIS SECTION

- [Syntax | 1395](#)
- [Hierarchy Level | 1396](#)
- [Description | 1396](#)
- [Required Privilege Level | 1396](#)
- [Release Information | 1397](#)

Syntax

```
services-options {
  cgn-pic;
  close-timeout seconds;
  fragment-limit number-of-fragments;
  disable-global-timeout-override;
  ignore-errors <alg> <tcp>;
  inactivity-non-tcp-timeout seconds;
```

```

inactivity-tcp-timeout seconds;
inactivity-timeout seconds;
open-timeout seconds;
pba-interim-logging-interval seconds;
reassembly-timeout seconds;
session-limit {
    maximum number;
    rate (Interface Services) new-sessions-per-second;
    cpu-load-threshold percentage;
}
session-timeout seconds;
jflow-log {
    message-rate-limit messages-per-second;
}
syslog {
    host hostname {
        facility-override facility-name;
        log-prefix prefix-value;
        port port-number;
        services severity-level;
    }
    message-rate-limit messages-per-second;
}
tcp-tickles tcp-tickles;
trio-flow-offload minimum-bytes minimum-bytes;
}

```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Define the service options to be applied on an interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

Configuring Default Timeout Settings for Services Interfaces

Configuring System Logging for Services Interfaces

shared-key

IN THIS SECTION

- [Syntax | 1397](#)
- [Hierarchy Level | 1397](#)
- [Description | 1398](#)
- [Options | 1398](#)
- [Required Privilege Level | 1398](#)
- [Release Information | 1398](#)

Syntax

```
shared-key value;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```


Description

Configure the authentication key value.

Options

value—Secret authentication value shared between a control source and destination.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

size

IN THIS SECTION

- [Syntax](#) | 1399
- [Hierarchy Level](#) | 1399
- [Description](#) | 1399
- [Options](#) | 1399
- [Required Privilege Level](#) | 1399
- [Release Information](#) | 1399

Syntax

```
size bytes;
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],
[edit forwarding-options sampling family (inet |inet6 |mpls) output file],
[edit forwarding-options sampling traceoptions file]
```

Description

Specify the maximum size of each file containing sample or log data. The file size is limited by the number of files to be created and the available hard disk space.

When a traffic sampling file named `sampling-file` reaches the maximum size, it is renamed `sampling-file.0`. When the `sampling-file` again reaches its maximum size, `sampling-file.0` is renamed `sampling-file.1` and `sampling-file` is renamed `sampling-file.0`. This renaming scheme continues until the maximum number of traffic sampling files is reached. Then the oldest traffic sampling file is overwritten.

Options

bytes—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

- **Syntax:** *kb* to specify KB, *mb* to specify MB, or *gb* to specify GB
- **Range:** 10 KB through the maximum file size supported on your router
- **Default:** 1 MB for sampling data; 128 KB for log information

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

skip-arp-iteration (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax](#) | 1400
- [Hierarchy Level](#) | 1400
- [Description](#) | 1400
- [Required Privilege Level](#) | 1401
- [Release Information](#) | 1401

Syntax

```
skip-arp-iteration;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Disable the Address Resolution Protocol (ARP) test iteration for IPv4 or inet services during a benchmarking test. This parameter is valid only for an inet family. An ARP test iteration is a 3-second iteration that is run for all inet tests. The results of this iteration are disregarded in the test result calculations. The ARP test iteration is performed by sending test frames for 3 seconds to ensure that all devices on the path to destination add ARP entries in the cache of the corresponding devices. This parameter is not applicable for CCC and bridge families.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

slamon-services

IN THIS SECTION

- [Syntax | 1401](#)
- [Hierarchy Level | 1402](#)
- [Description | 1402](#)
- [Options | 1402](#)
- [Required Privilege Level | 1402](#)
- [Release Information | 1402](#)

Syntax

```
slamon-services rfc2544;
```

Hierarchy Level

```
[edit chassis fpc slot-number]
```

Description

(MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, and MX10003 (with the LC2103 card) routers only) Enable service-level agreement (SLA) monitoring services support for RFC 2544-based benchmarking tests on MX Series routers with the MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP), MPC3E (MX-MPC3E-3D), MPC3E-NG (MX-MPC3E-3D-NG), MPC4E (MPC4E-3D-32XGE-SFPP and MPC4E-3D-2CGE-8XGE), MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-100G10G, and MPC5EQ-100G10G), MPC6E (MX2K-MPC6E), MPC7E (MPC7E-MRATE or MPC7E-10G), MX2K-MPC8E or MX2K-MPC9E, MPC10E (MPC10E-15C-MRATE or MPC10E-10C-MRATE), and MX2K-MPC11E line cards that are hosting test interfaces. For aggregated interfaces, enable support for RFC 2544-based benchmarking tests on all MPCs hosting child links. A system log is generated when you enable support for RFC 2544-based benchmarking tests on unsupported MPCs.

NOTE: On MX104 and MX80 Series routers that have a single fixed FPC, this configuration is not required.

Options

`rfc2544`—Enable support for RFC 2544-based benchmarking tests.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1R1.

RELATED DOCUMENTATION

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices](#) | 728

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)[Enabling Support for RFC 2544-Based Benchmarking Tests on MX Series Routers | 747](#)

soft-limit

IN THIS SECTION

- [Syntax | 1403](#)
- [Hierarchy Level | 1403](#)
- [Description | 1403](#)
- [Options | 1403](#)
- [Required Privilege Level | 1404](#)
- [Release Information | 1404](#)

Syntax

```
soft-limit bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which congestion notifications are sent to each control source of the criteria that point to this content destination. If the control source is configured with the `syslog` statement, a log message also be generated.

Options

bandwidth—Soft limit threshold, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

soft-limit-clear

IN THIS SECTION

- [Syntax](#) | 1404
- [Hierarchy Level](#) | 1404
- [Description](#) | 1405
- [Options](#) | 1405
- [Required Privilege Level](#) | 1405
- [Release Information](#) | 1405

Syntax

```
soft-limit-clear bandwidth;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Specify a bandwidth threshold at which the latch set by the soft-limit threshold is cleared.

Options

bandwidth—Soft-limit clear threshold, in bits per second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

[Configuring Junos Capture Vision | 289](#)

[soft-limit | 1403](#)

source-address (Forwarding Options)

IN THIS SECTION

- [Syntax | 1406](#)
- [Hierarchy Level | 1406](#)
- [Description | 1406](#)
- [Options | 1406](#)
- [Required Privilege Level | 1406](#)
- [Release Information | 1406](#)

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit forwarding-options accounting name outputinterface interface-name],
[edit forwarding-options monitoring namefamilyfamily inet output interface interface-name],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls | vpls) output interface interface-name],
[edit forwarding-options sampling family (inet |inet6 |mpls) output interface interface-name],
[edit forwarding-options sampling instance instance-name family inet output inline-jflow]
```

Description

Specify the source address for monitored packets.

Options

address—Interface source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Discard Accounting | 435](#)

[Configuring Flow Monitoring | 5](#)

[Configuring Traffic Sampling on MX, M and T Series Routers | 420](#)

source-address (RPM)

IN THIS SECTION

- [Syntax | 1407](#)
- [Junos OS Hierarchy Level | 1407](#)
- [Junos OS Evolved Hierarchy Level | 1407](#)
- [Description | 1407](#)
- [Options | 1408](#)
- [Required Privilege Level | 1408](#)
- [Release Information | 1408](#)

Syntax

```
source-address address;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet uses the outgoing interface's address as its source. Only IPv4 addresses are supported for Junos OS Evolved.

The following addresses cannot be used for the source IP address used for probes:

- 0.0.0.0

- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

Options

address—Valid IP address. Only IPv4 addresses are supported for Junos OS Evolved.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

| [Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | 651

source-address (TWAMP)

IN THIS SECTION

- [Syntax](#) | 1409
- [Hierarchy Level](#) | 1409
- [Description](#) | 1409
- [Default](#) | 1410

- Options | 1410
- Required Privilege Level | 1410
- Release Information | 1410

Syntax

```
source-address address;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection connection-name test-session session-name
```

Description

Specify the source IP address used for probes. If the source IP address is not one of the router's or switch's assigned addresses, the packet uses the outgoing interface's address as its source.

The following IPv4 addresses cannot be used as the source IP address used for probes:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

The following IPv6 addresses cannot be used as the source IP address used for probes:

- ::/128
- ::1/128 (loopback)
- fc00::/7 (unique-local)
- 2001:db8::/32 (documentation prefix)
- 2002::/16 (6to4)
- 5f00::/8 (6bone)

- 2001:10::/28 (ORCHID)
- 2001::/32 (Teredo)
- ::/0 (default route)
- ff00::/8 (multicast)
- :ffff:0:0/96 (IPv4-mapped addresses)

Default

The source address is selected based on the available default route.

Options

address Valid IPv4 or IPv6 address. You can only specify an IPv6 address as the source address if the source address corresponds to a TWAMP Light IPv6 link-local target address.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.4R1.

source-addresses

IN THIS SECTION

- [Syntax | 1411](#)
- [Hierarchy Level | 1411](#)
- [Description | 1411](#)
- [Options | 1411](#)

- Required Privilege Level | 1411
- Release Information | 1411

Syntax

```
source-addresses [ addresses ];
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name control-source identifier]
```

Description

List the IP addresses from which the control source can send control protocol requests to the Juniper Networks router.

Options

address—Allowed IP source address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

source-id

IN THIS SECTION

- [Syntax | 1412](#)
- [Hierarchy Level | 1412](#)
- [Description | 1412](#)
- [Options | 1413](#)
- [Required Privilege Level | 1413](#)
- [Release Information | 1413](#)

Syntax

```
source-id source-id;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
```

Description

For version 9 flows, a 32-bit value that identifies the Exporter Observation Domain is called the source ID. NetFlow collectors use the combination of the source IP address and the source ID field to separate different export streams originating from the same exporter.

MX and QFX Series: Typically, this attribute is used to limit the scope of other information elements. If the observation domain is not unique, the collector cannot uniquely identify an version 9 device. If you configure the same source ID for different template types, such as for IPv4 and IPv6, it does not impact flow monitoring because the actual or the base source ID is transmitted in the flow. The actual source ID is derived from the value you configure and plus other parameters such as the slot number, lookup chip (LU) instance, and Packet Forwarding Engine instance. Such a method of computation of the observation domain ID ensures that this ID is not the same for two version 9 devices.

PTX Series: When you configure the source ID, the software attaches the ID to a particular template type. If you configure the same source ID for two different template types, such as for IPv4 and IPv6,

this does not impact flow monitoring, because the configured ID is not what is being sent. The value sent in the packets is derived from that configured value and the FPC slot value. This method ensures two version 9 devices can never have the same value of observation domain ID.

Options

source-id—Identifier for the source for version 9 flows.

- **Range:** 0 through 255

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

[Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows | 620](#)

[Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows | 625](#)

source-ip (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax | 1414](#)
- [Hierarchy Level | 1414](#)
- [Description | 1414](#)
- [Options | 1414](#)
- [Required Privilege Level | 1414](#)

Syntax

```
source-ip address;
```

Hierarchy Level

```
[edit services jflow-log collector collector-name]
```

Description

Specify the source IPv4 address of the services PIC interface to be used for generation of flow monitoring log messages in flow monitoring template format for NAT events.

Options

address—Valid IPv4 address.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272

Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275

source-ipv4-address (RFC 2544 Benchmarking)

IN THIS SECTION

- Syntax | 1415
- Junos OS Hierarchy Level | 1415
- Junos OS Evolved Hierarchy Level | 1415
- Description | 1416
- Options | 1416
- Required Privilege Level | 1416
- Release Information | 1416

Syntax

```
source-ipv4-address address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the source IPv4 address to be used in generated test frames. This parameter is optional for both `ccc` and `inet` families. If you do not configure the source IPv4 address for an `inet` family, the source address of the interface is used to transmit the test frames.

Options

address Valid IPv4 address.

- **Default:** If you do not configure the source IPv4 address for a `ccc` family, default value of 192.168.1.10 is used.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)

source-mac-address

IN THIS SECTION

● [Syntax | 1417](#)

- Junos OS Hierarchy Level | 1417
- Junos OS Evolved Hierarchy Level | 1417
- Description | 1417
- Options | 1417
- Required Privilege Level | 1418
- Release Information | 1418

Syntax

```
source-mac-address mac-address;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the source MAC address used in generated test frames. This parameter is mandatory for a bridge, ethernet-switching, or vpls family and is optional for a ccc family.

Options

mac-address Source MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*, for example, 0000:5e00:5355 or 00:00:5e:00:53:55.

Default: 0x00:0x60:0x67:0x71:0xC6:0x62

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[rfc2544-benchmarking](#) | **1347**

[rfc2544](#) | **1345**

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | **855**

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices](#) | **728**

[Configuring an RFC 2544-Based Benchmarking Test](#) | **739**

source-udp-port (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax](#) | **1419**
- [Junos OS Hierarchy Level](#) | **1419**
- [Junos OS Evolved Hierarchy Level](#) | **1419**
- [Description](#) | **1419**
- [Options](#) | **1419**
- [Required Privilege Level](#) | **1419**
- [Release Information](#) | **1419**

Syntax

```
source-udp-port port-number;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the UDP port of the source to be used in the UDP header for the generated frames. If you do not specify the UDP port, the default value of 4041 is used.

Options

port-number Source UDP port number for the test frames

- **Default:** 4041

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X52.

Statement introduced in Junos OS Release 21.1R1.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)

stamp

IN THIS SECTION

- [Syntax | 1420](#)
- [Hierarchy Level | 1420](#)
- [Description | 1420](#)
- [Options | 1420](#)
- [Required Privilege Level | 1421](#)
- [Release Information | 1421](#)

Syntax

```
(stamp | no-stamp);
```

Hierarchy Level

```
[edit forwarding-options sampling family (inet | inet6 | mpls) output file]
```

Description

Include a timestamp with each line in the output file.

Options

`no-stamp`—Do not include timestamps. This is the default.

stamp—Include a timestamp with each line of packet sampling information.

- **Default:** No timestamp is included.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

step-percent (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax](#) | 1421
- [Hierarchy Level](#) | 1422
- [Description](#) | 1422
- [Options](#) | 1422
- [Required Privilege Level](#) | 1422
- [Release Information](#) | 1422

Syntax

```
step-percent percent;
```


Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profiletest-profile profile-name]
```

Description

Specify the step percentage for frame-loss tests. This parameter is not applicable for other type of tests. If you do not configure this parameter, the default step-percent is 10 percent.

Options

- percent* Step percent for frame-loss tests.
- **Default:** 10 percent
 - **Range:** 1 through 100 percent

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[rfc2544-benchmarking](#) | 1347

store

IN THIS SECTION

- [Syntax | 1423](#)
- [Hierarchy Level | 1423](#)
- [Description | 1424](#)
- [Default | 1424](#)
- [Options | 1424](#)
- [Required Privilege Level | 1425](#)
- [Release Information | 1425](#)

Syntax

```
store {  
    database;  
    file name {  
        files number;  
        (no-world-readable | world-readable);  
        size size;  
    }  
    fwding-file name{  
        files number;  
        (no-world-readable | world-readable);  
        size size;  
    }  
}
```

Hierarchy Level

[edit system [resiliency](#)]

Description

Store exception data in a file or database.

The on-box collector is disabled by default.

To enable the on-box collector for forwarding exceptions and store the exception data in a file, configure the following:

```
user@host# set system resiliency exceptions forwarding
user@host# set system resiliency store fwding-file file-name
user@host# set system resiliency store fwding-file size size
```

For forwarding exceptions, you must also configure inline-monitoring services to create the IPFIX records used to carry the exception data; see ["Understand Juniper Resiliency Interface" on page 409](#).

To enable the on-box collector for routing and kernel exceptions and store the telemetry key-value-pair exception data in a file, configure the following:

```
user@host# set system resiliency exceptions routing
user@host# set system resiliency exceptions os
user@host# set system resiliency store file file-name
user@host# set system resiliency store size size
```

Default

database

Options

- | | |
|------------------------------|---|
| database | Store exceptions data in the on-box SQLite database, in the /var/db directory. You can copy this database to a remote server and issue SQL commands using the SQLite command line interface. |
| file <i>file-name</i> | Use the specified file to receive the telemetry exceptions data. All files are placed in the directory /var/log . The files are also archived in the same way as trace logs. |
| files <i>files</i> | (Optional) Use the specified maximum number of 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. |

- **Default:** 3 files
- **Range:** 2 through 1000 files

fwding-file file-name	Use the specified file to receive the forwarding IPFIX exceptions data. All files are placed in the directory /var/log . The files are also archived in the same way as trace logs.
no-world-readable	(Default) Disable unrestricted file access. This means the exceptions file can be accessed only by the user who configured the exceptions operation.
world-readable	(Optional) Enable unrestricted file access.
size size	(Optional) Use the specified maximum size of each 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 files with the files option. <ul style="list-style-type: none"> • Range: 10 KB through 1 GB • Default: 128 KB

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1 for MX Series routers.

Statement introduced in Junos Evolved OS Release 22.2R1 for PTX Series routers.

RELATED DOCUMENTATION

[Juniper Resiliency Interface](#) | 409

storm-control

IN THIS SECTION

- [Syntax | 1426](#)
- [Hierarchy Level | 1426](#)
- [Description | 1426](#)
- [Options | 1426](#)
- [Required Privilege Level | 1427](#)
- [Release Information | 1427](#)

Syntax

```
storm-control {  
    count number;  
    interval number;  
}
```

Hierarchy Level

```
[edit services]
```

Description

Configure the count and the interval to control the flooding of SNMP traps per flow.

Options

- count *number*** Use the specified maximum number of SNMP traps generated in the configured interval.
- interval *number*** Use the specified minimum time period, in seconds, between the generation of successive traps.

- **Default:** The default count value is 1.

The default interval is 1 second.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers](#) | 927

[alarms](#) | 973

syslog

IN THIS SECTION

- [Syntax](#) | 1427
- [Hierarchy Level](#) | 1428
- [Description](#) | 1428
- [Required Privilege Level](#) | 1428
- [Release Information](#) | 1428

Syntax

```
(syslog | no-syslog);
```

Hierarchy Level

```
[edit interfaces mo-fpc/pic/port multiservice-options]
```

Description

System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the **/var/log** directory.

- `syslog`—Enable PIC system logging.
- `no-syslog`—Disable PIC system logging.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Monitoring](#) | 5

target-address

IN THIS SECTION

- [Syntax](#) | 1429
- [Junos OS Hierarchy Levels](#) | 1429
- [Junos OS Evolved Hierarchy Level](#) | 1429
- [Description](#) | 1429

- Options | 1429
- Required Privilege Level | 1430
- Release Information | 1430

Syntax

```
target-address (address | url url) local-link IPv6-link-local-interface-name;
```

Junos OS Hierarchy Levels

```
[edit services rpm probe owner test test-name]
[edit services rpm twamp client control-connection control-client-name]
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

(Required) Specify the destination IPv4 or IPv6 address (RPM and TWAMP) or URL (RPM) used for the probes. For RPM, Junos OS Evolved only supports IPv4 addresses.

Options

address —For all RPM probe types other than the HTTP probes, and for TWAMP managed control connections and test sessions, use the specified IPv4 or IPv6 address for the target host. Only IPv4 addresses are supported for Junos OS Evolved RPM probes. IPv4 addresses are supported for each Junos OS Evolved TWAMP Light test session (where the control-type light statement is configured for the control connection), and are configured per test session, rather than on the control connection as a whole. MX Series

and PTX Series Junos OS TWAMP Light test sessions support both IPv4 and IPv6 addresses.

NOTE: Starting with Junos OS Release 14.2R2, the RPM client router (the router or switch that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address.

local-link
IPv6-link-
local-
interface-
name

—(Junos OS only) For TWAMP Light test sessions with IPv6 target addresses (where the `control-type light` statement is configured for the control connection), configure the link-local logical interface name for the egress interface._

url *url*

—(Junos OS only) For HTTP probe types, use the specified fully formed URL that includes `http://` in the URL address. You can also specify an IPv6 address of a host in the URL to denote the destination or server to which the RPM probes must be sent.

NOTE: The *url* option is for RPM only.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Packet Transport routers.

Support at the `[edit services rpm twamp client control-connection control-client-name]` hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Release 18.1 for QFX Series switches.

Statement introduced in Junos OS Evolved Release 20.1R1.

Support for IPv6 addresses for TWAMP Light test sessions introduced in Junos OS Release 21.3R1 for MX Series and PTX1000, PTX3000, and PTX5000 routers.

local-link option for IPv6 TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

Configuring the Interface for RPM Timestamping for Client/Server on a Switch (CLI Procedure)

tcp

IN THIS SECTION

- [Syntax | 1431](#)
- [Hierarchy Level | 1431](#)
- [Description | 1432](#)
- [Required Privilege Level | 1432](#)
- [Release Information | 1432](#)

Syntax

```
tcp {  
    destination-interface interface-name;  
    port port;  
}
```

Hierarchy Level

```
[edit servicesrpm probe-server]
```

Description

Specify the port information for the TCP server.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring RPM Receiver Servers](#) | 662

tcp-keepcnt

IN THIS SECTION

- [Syntax](#) | 1433
- [Hierarchy Level](#) | 1433
- [Description](#) | 1433
- [Default](#) | 1433
- [Required Privilege Level](#) | 1433
- [Release Information](#) | 1433

Syntax

```
tcp-keepcnt number;
```

Hierarchy Level

```
[edit services rpm twamp (client control-connection control-client-name [server])]
```

Description

Number of unacknowledged probes to send before considering the connection dead and notifying the application layer. The range is 1 through 50.

Default

The default number of TCP KEEPALIVES sent is 6.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

[Understand Two-Way Active Measurement Protocol](#) | 686

tcp-keepidle

IN THIS SECTION

- [Syntax | 1434](#)
- [Hierarchy Level | 1434](#)
- [Description | 1434](#)
- [Default | 1434](#)
- [Required Privilege Level | 1434](#)
- [Release Information | 1435](#)

Syntax

```
tcp-keepidle seconds;
```

Hierarchy Level

```
[edit services rpm twamp (client control-connection control-client-name |server)]
```

Description

Time interval between the last data packet sent and the first keepalive probe. The range is 1 through 600 seconds.

Default

The default value is 120 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches | 694](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

tcp-keepintvl

IN THIS SECTION

- [Syntax | 1435](#)
- [Hierarchy Level | 1435](#)
- [Description | 1436](#)
- [Default | 1436](#)
- [Required Privilege Level | 1436](#)
- [Release Information | 1436](#)

Syntax

```
tcp-keepintvl seconds;
```

Hierarchy Level

```
[edit services rpm twamp (client control-connection control-client-name |server)]
```

Description

Time interval between successive keepalive probes. The range is 1 second through 600 seconds.

Default

The default value is 5 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

[Understand Two-Way Active Measurement Protocol](#) | 686

template (Flow Monitoring IPFIX Version)

IN THIS SECTION

- [Syntax](#) | 1437
- [Hierarchy Level](#) | 1437
- [Description](#) | 1437
- [Options](#) | 1438
- [Required Privilege Level](#) | 1438
- [Release Information](#) | 1438

Syntax

```
template template-name {
  data-record-fields {
    source-prefix-as-path count;
    destination-prefix-as-path count;
    bgp-source-standard-community count;
    bgp-destination-standard-community count;
    bgp-source-extended-community count;
    bgp-destination-extended-community count;
    bgp-source-large-community count;
    bgp-destination-large-community count;
  }
  flow-active-timeout seconds;
  flow-inactive-timeout seconds;
  flow-key {
    flow-direction;
    vlan-id;
    output-interface;
  }
  (bridge-template|ipv4-template | ipv6-template | mpls-ipv4-template | mpls-ipvx-template | vpls-
  template);
  nexthop-learning (enable |disable);
  observation-domain-id;
  option-refresh-rate packets packets seconds seconds;
  options-template-id;
  template-id;
  template-refresh-rate packets packets seconds seconds;
  tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
}
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix]
```

Description

Specify the IPFIX output template properties to support flow monitoring.

Options

template-name Name of the IPFIX template.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

data-record-fields option introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

template (Flow Monitoring Version 9)

IN THIS SECTION

- [Syntax | 1439](#)
- [Hierarchy Level | 1439](#)
- [Description | 1439](#)
- [Options | 1439](#)
- [Required Privilege Level | 1439](#)
- [Release Information | 1440](#)

Syntax

```
template template-name {
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-key {
        flow-direction;
        vlan-id;
        output-interface;
    }
    (bridge-template|ipv4-template | ipv6-template | mpls-template | vpls-template|label-position
[ positions ] | mpls-ipv4-template label-position [ positions ] | mpls-ipvx-template);
    option-refresh-rate packets packets seconds seconds;
    options-template-id;
    peer-as-billing-template;
    source-id;
    template-id;
    template-refresh-rate packets packets seconds seconds;
    tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
}
```

Hierarchy Level

```
[edit services flow-monitoring version9]
```

Description

Specify the version 9 output template properties to support flow monitoring.

Options

template-name—Name of the version 9 template.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates](#) | 583

template (Forwarding Options)

IN THIS SECTION

- [Syntax](#) | 1440
- [Hierarchy Level](#) | 1440
- [Description](#) | 1441
- [Options](#) | 1441
- [Required Privilege Level](#) | 1441
- [Release Information](#) | 1441

Syntax

```
template template-name;
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server
hostname version9],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname version9]
```

Description

Specify flow monitoring version 9 template to be used for output of sampling records.

Options

template-name—Name of the version 9 template.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates](#) | 583

template (Forwarding Options Version IPFIX)

IN THIS SECTION

- [Syntax](#) | 1442
- [Hierarchy Level](#) | 1442
- [Description](#) | 1442
- [Required Privilege Level](#) | 1442
- [Release Information](#) | 1442

Syntax

```
template;
```

Hierarchy Level

```
[edit forwarding-options sampling instance family (inet | inet6 | mpls | vpls) output flow-  
server hostname version-ipfix]
```

Description

Specify flow monitoring version IPFIX properties to apply to output sampling records.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

template (Inline Monitoring)

IN THIS SECTION

● [Syntax \(Junos OS\) | 1443](#)

- Syntax (Junos OS Evolved) | 1443
- Hierarchy Level | 1444
- Description | 1444
- Options | 1444
- Required Privilege Level | 1445
- Release Information | 1445

Syntax (Junos OS)

```
template name {
    flow-inactive-timeout seconds;
    flow-monitoring {
        counter-profile profile-identifier;
        flow-rate kbps burst-size bytes;
        flow-limit number;
        sampling-rate bytes;
        sampling-profile profile-name;
        security-enable;
    }
    observation-domain-id observation-domain-id;
    option-template-id option-template-id;
    option-template-refresh-rate seconds;
    primary-data-record-fields name;
    template-id template-id;
    template-refresh-rate seconds;
    template-type (ipv4-template | ipv6-template);
}
```

Syntax (Junos OS Evolved)

```
template name {
    observation-domain-id observation-domain-id;
    option-template-refresh-rate seconds;
    primary-data-record-fields name;
```

```
template-refresh-rate seconds;  
}
```

Hierarchy Level

```
[edit services inline-monitoring]
```

Description

Configure template for inline monitoring services.

Options

<i>name</i>	Name of the template.
<i>flow-inactive-timeout</i> <i>seconds</i>	(EX switches only) Configure the inactive-timeout period for a flow, in seconds, for flow-based telemetry. Once there is no active traffic for a flow, the flow is aged out after the configured inactive-timeout period.
<i>flow-monitoring</i>	(EX switches only) Configure flow-based telemetry. See flow-monitoring (Inline Monitoring Services) for more information.
<i>observation-domain-id</i> <i>observation-domain-id</i>	Significant one byte of observation domain ID used to uniquely identify the exporting process. The other three bytes are system generated and are unique within the chassis. <ul style="list-style-type: none"> • Default: 0 • Range: 1 through 255
<i>option-template-id</i> <i>option-template-id</i>	(Junos OS only) Option template ID. For Junos OS Evolved, the system generates the option template ID. <ul style="list-style-type: none"> • Default: 640 • Range: 1024 through 65535
<i>option-template-refresh-rate</i> <i>seconds</i>	Option refresh rate in seconds. <ul style="list-style-type: none"> • Default: 600 seconds • Range: 10 through 600 seconds

primary-data-record-fields <i>name</i>	Configure which IPFIX information elements (IEs) to include in the primary data record for Juniper Resiliency Interface. See " primary-data-record-fields " on page 1310 for more information.
template-id <i>template-id</i>	(Junos OS only) Template ID. For Junos OS Evolved, the system generates the template ID. <ul style="list-style-type: none"> • Default: 384 • Range (EX4400): 1024 through 65535 • Range (EX4100 and EX4100-F): 32768 through 32831
template-refresh-rate <i>seconds</i>	Refresh rate in seconds. <ul style="list-style-type: none"> • Default: 600 seconds • Range: 10 through 600 seconds
template-type (ipv4-template ipv6-template)	(QFX5120 only, Junos OS only) For flow-based telemetry for VXLANs, specify that the template is for either IPv4 or IPv6 traffic.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 19.4R1.

flow-inactive-timeout and flow-monitoring options introduced in Junos OS Release 21.1R1.

primary-data-record-fields option introduced in Junos OS Release 21.2R1.

Statement introduced in Junos OS Evolved Release 22.2R1 for PTX Series routers.

template-type option introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

[Inline Monitoring Services Configuration](#) | 334

[Flow-Based Telemetry \(EX4100, EX4100-F, and EX4400 Series\)](#) | 349

template-id

IN THIS SECTION

- [Syntax | 1446](#)
- [Hierarchy Level | 1446](#)
- [Description | 1446](#)
- [Options | 1447](#)
- [Required Privilege Level | 1447](#)
- [Release Information | 1447](#)

Syntax

```
template-id id;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Define a template ID to be used for flow aggregation of version 9 and IPFIX flows.

MX and QFX Series

The template ID can be a value in the range of 1024 through 65535. If you do not configure a value for the template ID, a default value is assumed for this ID, which is different for the various address families. If you configure the same template ID value for different address families, such a setting is not processed properly and might cause unexpected behavior. For example, if you configure the same template ID

value for both IPv4 and IPv6, the collector validates the export data based on the template ID value that it last receives. In this case, if IPv6 is configured after IPv4, the value is effective for IPv6 and the default value is used for IPv4.

PTX Series

If you choose to configure the template ID, the range is 1024 to 65520. If you do not configure this ID, the default value that are set is in the range 256-511 and is different for each template.

You can configure the `template-id` statement for family `inet`, `inet6`, and `mpls` only.

You must not configure the same IDs for different templates (option or data template). The operating system does not check to ensure that you do not configure the same ID value for different templates. If you configure the same ID value, such a setting is not processed properly and might cause unexpected behavior.

The template ID range [configured `template-id` value + 20) is reserved and you must not configure any another ID in this range. The difference between configured template IDs across families should be at least 20; for example, if `template-id 1056` is configured for family `inet`, you should not configure template IDs in the range of 1056 to 1075 for any other families.

For Junos OS, if you change the template ID, all flows are inactively timed out. New flows are learned afresh.

For Junos OS Evolved, if you change the template ID, this change does not impact the flows.

Options

id—Unique identifier for the template to be used for version 9 or IPFIX flows. The range 0-255 is reserved and is used for template sets, options template sets, and for future expansion. If you do not configure a template ID, the software generates a template ID for you.

- **Range:** 1024 through 65,535

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows | 620](#)

[Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows | 625](#)

template-profile (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax | 1448](#)
- [Hierarchy Level | 1448](#)
- [Description | 1448](#)
- [Options | 1449](#)
- [Required Privilege Level | 1449](#)
- [Release Information | 1449](#)

Syntax

```
template-profile template-profile-name;
```

Hierarchy Level

```
[edit services jflow-log],  
[edit services service-set service-set-name jflow-log]
```

Description

Specify the name of the flow template profile to be used for generating flow monitoring format messages for NAT events and for transmitting them to the collector. You can define a template profile for the Jflow service by using this statement at the `[edit services jflow-log]` hierarchy level, and associate the template profile with a service set by using this statement at the `[edit services service-set service-set-name jflow-log]` hierarchy level.

Options

template-profile-name—Name of the flow template profile for NAT events. The name can be up to 32 alphanumeric characters in length. Allowed characters are [a-zA-Z0-9_].

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

template-refresh-rate

IN THIS SECTION

- [Syntax | 1450](#)
- [Hierarchy Level | 1450](#)
- [Description | 1450](#)
- [Options | 1450](#)
- [Required Privilege Level | 1450](#)
- [Release Information | 1450](#)

Syntax

```
template-refresh-rate packets packets seconds seconds;
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name],  
[edit services flow-monitoring version-ipfix template template-name],
```

Description

Specify the frequency at which the flow generator sends updates about template definitions to the flow collector. Specify either the number of packets or the number of seconds.

Options

packets—Refresh rate, in number of packets.

- **Range:** 1 through 480,000
- **Default:** 4800

seconds—Refresh rate, in number of seconds.

- **Range:** 10 through 600
- **Default:** 600

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the [edit services flow-monitoring version-ipfix template *template-name*] hierarchy level added in Junos OS Release 10.2.

Support at the [edit services flow-monitoring version9 template *template-name*] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

template-type (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax | 1451](#)
- [Hierarchy Level | 1451](#)
- [Description | 1452](#)
- [Options | 1452](#)
- [Required Privilege Level | 1452](#)
- [Release Information | 1452](#)

Syntax

```
template-type nat;
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the type of service for which flow template profiles, in version or IPFIX format, must be used for generating flow monitoring format messages for NAT events and for transmitting them to the collector. Currently, you can configure only NAT events or services for generation of log messages in flow monitoring format.

Options

nat—Use flow template profiles for generation of flow monitoring logs for NAT events.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

templates

IN THIS SECTION

● [Syntax | 1453](#)

- Hierarchy Level | 1454
- Description | 1454
- Options | 1454
- Required Privilege Level | 1455
- Release Information | 1455

Syntax

```

templates {
  template-name {
    interval-duration interval-duration;
    inactive-timeout inactive-timeout;
    rate {
      (layer3 layer3-packets-per-second | media media-bits-per-second);
    }
    delay-factor {
      ;
      threshold {
        (info | warning | critical) delay-factor-threshold;
      }
    }
    media-loss-rate {
      disable;
      threshold {
        (info | warning | critical) percentage mlr-percentage | packet-count mlr-packet-
count);
      }
    }
    media-rate-variation {
      disable;
      threshold {
        (info | warning | critical) mrp-variation;
      }
    }
    media-packets-count-in-layer3 media-packets-count-in-layer3;
    media-packet-size media-packet-size;
  }
}

```



```
    }  
  }  
}
```

Hierarchy Level

```
[edit services video-monitoring]
```

Description

Configure the media delivery index template containing the measurement parameters for video monitoring.

Options

delay-factor	Define delay factor syslog threshold levels.
<i>delay-factor-threshold</i>	<div>Delay factor threshold in milliseconds. When the threshold is exceeded, a syslog message is generated.</div> <ul style="list-style-type: none">• Default: 0—Do not generate syslogs.• Range: 0 though 65,535 milliseconds
disable	Disable logging for the threshold.
<i>inactive-timeout</i>	<div>Number of seconds of flow inactivity after which time media delivery index statistics collection for a flow is terminated.</div> <ul style="list-style-type: none">• Range: 30 through 300 seconds
info warning critical	Level of syslog message generated when a threshold is exceeded.
<i>interval-duration</i>	<div>Number of seconds after which time media delivery index flow monitoring statistics for the interval are reported.</div> <ul style="list-style-type: none">• Range: 1 through 50
<i>layer3-packets-per-second</i>	<div>Layer 3 packet rate in IP packets per second.</div> <ul style="list-style-type: none">• Range: 0 though 4,294,967,295 pps
<i>media-bits-per-second</i>	Media bit rate for the stream in bits per second.

media-loss-rate	Define media loss rate syslog threshold levels.
<i>media-packets-count-in-layer-3</i>	Number of media packets in an IP packet. <ul style="list-style-type: none"> • Range: 1 through 32
<i>media-packet-size</i>	Size of media packet in bits. <ul style="list-style-type: none"> • Default: 188 • Range: 1 through 2048
media-rate-variation	Define delay factor syslog threshold levels.
<i>mlr-packet-count</i>	Media loss rate threshold expressed as the number of packets dropped. When the threshold is exceeded, a syslog message is generated.
<i>mlr-percentage</i>	Media loss rate threshold expressed as the percentage of total packets dropped. When the threshold is exceeded, a syslog message is generated. <ul style="list-style-type: none"> • Range: 0 through 100
<i>mrsv-variation</i>	Media rate variation threshold. The variation is the ratio of actual media rate to the configured media rate, expressed as a percentage.
<i>template-name</i>	Name of the template containing media delivery index measurement criteria. The template can be assigned to an interface.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

| [Configuring Inline Video Monitoring on MX Series Routers](#) | 934

test

IN THIS SECTION

- [Junos OS | 1456](#)
- [Junos OS Evolved | 1457](#)
- [Junos OS Hierarchy Level | 1458](#)
- [Junos OS Evolved Hierarchy Level | 1458](#)
- [Description | 1458](#)
- [Options | 1458](#)
- [Required Privilege Level | 1459](#)
- [Release Information | 1459](#)

Junos OS

```
test test-name {  
    data-fill data;  
    data-size size;  
    destination-interface interface-name;  
    destination-port port;  
    dscp-code-points dscp-bits;  
    hardware-timestamp;  
    history-size size;  
    inet6-options {  
        source-address address;  
    }  
    moving-average-size number;  
    next-hop next-hop;  
    one-way-hardware-timestamp;  
    probe-count count;  
    probe-interval seconds;  
    probe-type type;  
    routing-instance instance-name;  
    rpm-scale {  
        destination {  
            interface interface-name.logical-unit-number;
```

```

        subunit-cnt subunit-cnt;
    }
    source {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    source-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    target {
        address-base ipv4-address-base;
        count ipv4-count;
        step ipv4-step;
    }
    target-inet6 {
        address-base ipv6-address-base;
        count ipv6-count;
        step ipv6-step;
    }
    tests-count tests-count;
}
source-address address;
target (url url | address address);
test-interval interval;
thresholds thresholds;
traps traps;
ttl hop-count
}

```

Junos OS Evolved

```

test test-name {
    data-fill data;
    data-size size;
    destination-port port;
    dscp-code-points dscp-bits;
    history-size size;
    moving-average-size number;
}

```

```

offload-type {
    none;
    pfe-timestamp;
}
probe-count count;
probe-interval seconds;
probe-type type;
routing-instance instance-name;
source-address address;
target (url url | address address);
test-interval interval;
thresholds thresholds;
traps traps;
ttl hop-count
}

```

Junos OS Hierarchy Level

```
[edit services rpm probe owner]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name]
```

Description

Specify the range of probes over which the standard deviation, average, and jitter are calculated. The test name combined with the owner name represent a single RPM configuration instance.

Options

test *name*—You can configure any name up to 32 characters in length. For Junos OS Evolved, if the owner name is one of the pre-defined names, then four pre-defined test names are available to choose from. For example, if the owner name is the pre-defined name `icmp-evo`, then these four pre-defined test names are available to configure: `icmp-evo-1`, `icmp-evo-2`, `icmp-evo-3`, and `icmp-evo-4`.

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

inet6-options option added in Junos OS Release 14.1R4 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

traps option introduced in Junos OS Evolved Release 21.2R1.

RELATED DOCUMENTATION

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches | 646](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

tests

IN THIS SECTION

- [Junos OS Syntax \(Except SRX300 and SRX550HM\) | 1460](#)
- [Junos OS Syntax \(SRX300 and SRX550HM\) | 1460](#)
- [Junos OS Evolved Syntax | 1460](#)
- [Junos OS Hierarchy Level | 1461](#)
- [Junos OS Evolved Hierarchy Level | 1461](#)
- [Description | 1461](#)
- [Required Privilege Level | 1461](#)
- [Release Information | 1462](#)

Junos OS Syntax (Except SRX300 and SRX550HM)

```
tests {
  test-name test-name {
    destination-ipv4-address address;
    destination-udp-port port-number;
    direction (egress | ingress);
    disable-signature-check;
    family (bridge | ccc | inet | vpls);
    mode reflect;
    source-ipv4-address address;
    source-udp-port port-number;
    test-interface interface-name;
    test-iterator-duration test-iterator-duration;
  }
}
```

Junos OS Syntax (SRX300 and SRX550HM)

```
tests {
  test-name test-name {
    destination-ipv4-address address;
    destination-udp-port port-number;
    disable-signature-check;
    family inet;
    mode reflect;
    source-ipv4-address address;
    source-udp-port port-number;
    test-interface interface-name;
  }
}
```

Junos OS Evolved Syntax

```
tests {
  test-name test-name {
    destination-mac-address mac-address;
    destination-ipv4-address address;
```

```

    destination-udp-port port-number;
    family (ccc | ethernet-switching | inet);
    mode reflect;
    direction (egress | ingress);
    disable-signature-check;
    in-service;
    ip-swap;
    reflect-etype;
    reflect-mode (mac-swap | no-mac-swap);
    service-type service-type;
    source-mac-address mac-address;
    source-ipv4-address address;
    source-udp-port port-number;
    test-interface interface-name;
    udp-tcp-port-swap;
  }
}

```

Junos OS Hierarchy Level

[edit services [rpm](#) [rfc2544-benchmarking](#)]

Junos OS Evolved Hierarchy Level

[edit services monitoring [rfc2544](#)]

Description

Specify the attributes of the test iteration, such as the address family (type of service, IPv4 or Ethernet), the logical interface, test duration, and test packet size, that are used for a benchmarking test to be run. The test name combined with the test profile represent a single real-time performance monitoring (RPM) configuration instance.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1.

Support for the destination-mac-address, direction, disable-signature-check, in-service, ip-swap, reflect-etype, reflect-mode, service-type, source-mac-address, and udp-tcp-port-swap configuration statements added in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)

test-count

IN THIS SECTION

- [Syntax | 1462](#)
- [Hierarchy Level | 1463](#)
- [Description | 1463](#)
- [Options | 1463](#)
- [Required Privilege Level | 1463](#)
- [Release Information | 1463](#)

Syntax

```
test-count test_count-number;
```

Hierarchy Level

```
[edit services rpm twamp client control-connection control-client-name]
```

Description

Specify the total number of test session iterations. The range is 0 through 4294967290.

Options

- **Default:** The default value is 0.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

test-finish-wait-duration (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1464](#)
- [Hierarchy Level | 1464](#)
- [Description | 1464](#)
- [Required Privilege Level | 1464](#)
- [Release Information | 1464](#)

Syntax

```
test-finish-wait-duration;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Number of seconds to wait after transmitting the last frame and before concluding that the test as complete. Once the test is complete, the frames received after that are not considered for the result computation. Use this parameter if the latency introduced by the network under test is high in the normal conditions. Default value is 1 second.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[rfc2544-benchmarking](#) | 1347

test-interface (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1465](#)
- [Junos OS Hierarchy Level | 1465](#)
- [Junos OS Evolved Hierarchy Level | 1465](#)
- [Description | 1465](#)
- [Options | 1466](#)
- [Required Privilege Level | 1466](#)
- [Release Information | 1466](#)

Syntax

```
test-interface interface-name;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Specify the logical interface on which the RFC 2544-based benchmarking test is run.

For Junos OS, if you configure an `inet` family and the test mode to initiate and terminate test frames on the same device, the interface you configure is not effective. Instead, the test is run on the egress logical interface that is determined using route lookup on the specified destination IPv4 address. If you configure an `inet` family and the test mode to reflect the frames back on the sender from the other end,

the logical interface is used as the interface to enable the reflection service (reflection is performed on the packets entering the specified interface). If you not configure the logical interface for reflection test mode, a lookup is performed on the source IPv4 address to determine the interface that hosts the address.

For Junos OS Evolved, you can configure this statement for only family `ccc` or `ethernet-switching`. You cannot configure this statement for a Layer 3 reflector (family `inet`).

Options

interface-name Name of the logical interface on which the test needs to be run.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[rfc2544-benchmarking | 1347](#)

[rfc2544 | 1345](#)

test-interval

IN THIS SECTION

- [Syntax | 1467](#)
- [Junos OS Hierarchy Levels | 1467](#)
- [Junos OS Evolved Hierarchy Level | 1467](#)
- [Description | 1467](#)
- [Options | 1468](#)
- [Required Privilege Level | 1468](#)
- [Release Information | 1468](#)

Syntax

```
test-interval seconds;
```

Junos OS Hierarchy Levels

```
[edit services rpm bgp],  
[edit services rpm probe owner test test-name],  
[edit services rpm twamp client control-connection control-client-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the time to wait between tests, in seconds. A test interval of 0 seconds causes the RPM test to stop after one iteration.

Options

seconds—Number of seconds to wait between tests.

- **Range:** [edit services rpm bgp], [edit services rpm probe *owner* test *test-name*], and [edit services monitoring rpm owner *name* test *name*] hierarchy levels: 0 through 86,400
- **Range:** [edit services rpm twamp client control-connection *control-client-name*] hierarchy level: 1 through 255
- **Default:** 1

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Release 9.3 for EX Series switches.

Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Configuring BGP Neighbor Discovery Through RPM | 671](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

test-iterator-duration (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1469](#)
- [Hierarchy Level | 1469](#)
- [Description | 1469](#)
- [Options | 1469](#)
- [Required Privilege Level | 1470](#)
- [Release Information | 1470](#)

Syntax

```
test-iterator-duration seconds;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the duration of each iteration in seconds. If you configure this value, the default value of each iteration depends on the type of test being run. For throughput, back-back-frames and frame-loss types of tests, the default value is 20 seconds. For latency tests, the default value is 120 seconds.

Options

seconds Number of seconds for which each test iteration is run.

- **Range:** 10 through 120 seconds

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

test-iterator-pass-threshold (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1470](#)
- [Hierarchy Level | 1471](#)
- [Description | 1471](#)
- [Required Privilege Level | 1471](#)
- [Release Information | 1471](#)

Syntax

```
test-iterator-pass-threshold;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the pass threshold of each iteration.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview](#) | 855

[Configuring RFC 2544-Based Benchmarking Tests](#) | 864

[rfc2544-benchmarking](#) | 1347

test-name

IN THIS SECTION

- [Junos OS Syntax \(Except SRX300 and SRX550HM\)](#) | 1472
- [Junos OS Syntax \(SRX300 and SRX550HM\)](#) | 1472
- [Junos OS Evolved Syntax](#) | 1473
- [Junos OS Hierarchy Level](#) | 1473
- [Junos OS Evolved Hierarchy Level](#) | 1473
- [Description](#) | 1473

- Options | 1473
- Required Privilege Level | 1473
- Release Information | 1474

Junos OS Syntax (Except SRX300 and SRX550HM)

```
test-name test-name {
    destination-ipv4-address address;
    destination-udp-port port-number;
    direction (egress | ingress);
    disable-signature-check;
    family (bridge | ccc | inet | vpls);
    mode reflect;
    source-ipv4-address address;
    source-udp-port port-number;
    test-interface interface-name;
    test-iterator-duration test-iterator-duration;
}
```

Junos OS Syntax (SRX300 and SRX550HM)

```
test-name test-name {
    destination-ipv4-address address;
    destination-udp-port port-number;
    disable-signature-check;
    family inet;
    mode reflect;
    source-ipv4-address address;
    source-udp-port port-number;
    test-interface interface-name;
}
```

Junos OS Evolved Syntax

```
test-name test-name {
    destination-ipv4-address address;
    destination-udp-port port-number;
    family inet;
    mode reflect;
    source-ipv4-address address;
    source-udp-port port-number;
    test-interface interface-name;
}
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests]
```

Description

Define the name of the RFC 2544-based benchmarking test. For each unique test name that you configure, you can specify a test profile, which contains the settings for a test and its type, and also a test interface, which contains the settings for test packets that are sent and received on the selected interface.

Options

test-name Test name. The name can be up to 32 characters in length.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.3.

Statement introduced in Junos OS Evolved Release 21.1R1.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)

test-profile (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax | 1474](#)
- [Hierarchy Level | 1475](#)
- [Description | 1475](#)
- [Options | 1475](#)
- [Required Privilege Level | 1475](#)
- [Release Information | 1475](#)

Syntax

```
test-profile profile-name;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking tests test-name test-name]
[edit services rpm rfc2544-benchmarking profiles]
```

Description

Specify the name of the test profile to be associated with a particular test name. This parameter is required when the test mode is configured as initiate-and-terminate. This parameter is disregarded when the test mode is configured as reflection. A reflection service does not use the parameters specified in the test profile.

Options

profile-name Name of the test profile. The name can be up to 32 characters in length. The name must start with a letter. Allowed characters are [a-zA-Z0-9_]

Required Privilege Level

system—To view this statement in the configuration.

system -control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

test-session (Junos OS)

IN THIS SECTION

- [Syntax | 1476](#)
- [Hierarchy Level | 1476](#)
- [Description | 1476](#)
- [Options | 1477](#)
- [Required Privilege Level | 1477](#)
- [Release Information | 1477](#)

Syntax

```
test-session session-name {  
    data-fill-with zeros data;  
    data-size size;  
    destination-port port;  
    dscp-code-points dscp-bits;  
    probe-count count;  
    probe-interval seconds;  
    source-address source-address;  
    target-address target-address local-link IPv6-link-local-interface-name;  
}
```

Hierarchy Level

```
[edit services rpm twamp client control-connection connection-name]
```

Description

Specify the test session details that includes the session name, the contents of the test packet, the data size, the probe details, and the target destination details.

Options

session-name Name of the session.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

destination-port option added in Junos OS Release 21.1R1.

source-address option and the local-link sub-option of the target-address option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

| [Understand Two-Way Active Measurement Protocol](#) | 686

test-session (Junos OS Evolved)

IN THIS SECTION

- [Syntax](#) | 1478
- [Hierarchy Level](#) | 1478
- [Description](#) | 1479
- [Options](#) | 1479
- [Required Privilege Level](#) | 1482
- [Release Information](#) | 1482

Syntax

```
test-session name {
    data-size data-size;
    destination-port destination-port;
    dscp-code-points dscp-code-points;
    history-size history-size;
    moving-average-size moving-average-size;
    offload-type (none | inline-timestamp | pfe-timestamp);
    probe-count probe-count;
    probe-interval seconds;
    source-address source-address;
    target target-address local-link IPv6-link-local-interface-name;
    thresholds {
        control-failure (on | off);
        successive-loss number;
        total-loss number;
        threshold-type (microseconds | average);
    }
    traps {
        egress-jitter-exceeded;
        egress-time-exceeded;
        ingress-jitter-exceeded;
        ingress-time-exceeded;
        jitter-exceeded;
        probe-failure;
        rtt-exceeded;
        test-completion;
        test-failure;
    }
    ttl hop-count;
}
```

Hierarchy Level

```
[edit services monitoring twamp client control-connection name]
```

Description

Specify the TWAMP test session details, including the session name, the contents of the test packet, the data size, the probe details, and the target destination. You must configure at least one test session.

Options

<i>name</i>	Name of the test session.
<i>data-size bytes</i>	Specify the size of the data portion of the test packet, in bytes. <ul style="list-style-type: none"> • Range: 0 through 65400 bytes • Default: 0
<i>destination-port port</i>	Specify the User Datagram Protocol (UDP) port number to which a probe is sent. You must have the control-type option for the control-connection statement set to <code>light</code> to be able to configure a destination port for a particular test session. If you set the control-type option to <code>managed</code> , the destination port is negotiated for you between the client and the server, and so you cannot configure the destination port for a particular test session. <ul style="list-style-type: none"> • Range: You can specify port 862, or any port from 49152 through 65535. • Default: 862 (IANA port for TWAMP)
<i>dscp-code-points dscp-bits</i>	Configure the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern. <ul style="list-style-type: none"> • Values: Configure a valid 6-bit pattern; for example, 001111, or one of the following DSCP aliases: <ul style="list-style-type: none"> • af11—Default: 001010 • af12—Default: 001100 • af13—Default: 001110 • af21—Default: 010010 • af22—Default: 010100 • af23 —Default: 010110

- af31 —Default: 011010
- af32 —Default: 011100
- af33 —Default: 011110
- af41 —Default: 100010
- af42 —Default:100100
- af43 —Default:100110
- be—Default: 000000
- cs1—Default: 001000
- cs2—Default: 010000
- cs3—Default: 011000
- cs4—Default: 100000
- cs5—Default: 101000
- cs6—Default: 110000
- cs7—Default: 111000
- ef—Default: 101110
- nc1—Default: 110000
- nc2—Default: 111000

**history-size
number**

Specify the number of stored history entries. The value configured for the history-size option must be equal to or larger than the value configured for the probe-count option, or the configuration will not commit.

- **Range:** 0 to 255
- **Default:** 50

**moving-
average-size
number**

Enable statistical calculations to be performed across a configurable number of the most recent samples.

- **Range:** 0 through 512
- **Default:** 0 (disable)

offload-type (none inline-timestamp pfe-timestamp)	<p>Enable timestamping of TWAMP probe messages in the Routing Engine, inline in the hardware, or Packet Forwarding Engine host processor. Prior to Junos OS Evolved 22.3R1, for IPv6 traffic, you must configure none, as only Routing Engine timestamping is supported. Starting in Junos OS Evolved 22.3R1, you can configure the pfe-timestamping option for IPv6 traffic.</p> <ul style="list-style-type: none"> • none—Timestamping in the Routing Engine. • inline-timestamping—Timestamping inline, in the hardware at the generator, and reflection in the hardware at the reflector. When this option is configured for an ACX router, the ACX can act as either a server or client, but not both. Also for an ACX router, inline-timestamping is not supported on PM50 ports if FEC74 is enabled on that port. You can configure the inline-timestamping option for up to 256 test sessions. • pfe-timestamping—Timestamping in the Packet Forwarding Engine. • Default: pfe-timestamp
probe-count <i>number</i>	<p>Specify the number of probes within a test. The value configured for the probe-count option must be smaller than the value configured for the history-size option, or the configuration will not commit.</p> <ul style="list-style-type: none"> • Range: 1 through 255 • Default: 1
probe-interval <i>seconds</i>	<p>Specify the number of seconds to wait before sending the next probe.</p> <ul style="list-style-type: none"> • Range: 1 through 255 seconds • Default: 1
source-address <i>address</i>	<p>Specify the IPv4 or IPv6 source address used for test probes. If the source address is not one of the device's assigned addresses, the probe uses the outgoing interface's address as its source.</p> <p>You cannot use IPv6 link-local addresses or the following addresses as the source address for probes:</p> <ul style="list-style-type: none"> • 0.0.0.0 • 127.0.0.0/8 (loopback) • 224.0.0.0/4 (multicast)

- 255.255.255.255 (broadcast)

When an IPv6 link-local address is configured with the target *target-address* local-link *IPv6-link-local-interface-name* option, you cannot configure the source-address *address* option. Instead, the operating system sets this option for you.

target *target-address* local-link *IPv6-link-local-interface-name*

Specify the IPv4 or IPv6 address for the target destination of the probe. Prior to Junos OS Evolved 22.3R1, you cannot use an IPv6 link-local address for the target destination of a probe. Starting in Junos OS Evolved 22.3R1, for TWAMP Light test sessions with IPv6 target addresses (where the control-type *light* statement is configured for the control connection), configure the link-local logical interface name for the egress interface. You cannot configure IPv6 link-local addresses for TWAMP managed test sessions (where the control-type *managed* statement is configured for the control connection) .

ttl *hop-count*

Specify the maximum number of hops a TWAMP probe can travel. You configure the ttl option when necessary to restrict the scope of a probe, so that a probe does not unintentionally monitor an alternative path to the destination (such as may occur after a BGP re-routing). You can also use the ttl option to monitor direct reachability by specifying a TTL of 1. Probes that exceed the number configured are discarded.

- **Range:** 1 through 255 hops
- **Default:** 64 hops

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved Release 20.3R1.

traps option introduced in Junos OS Evolved Release 21.3R1.

IPv6 address support introduced in Junos OS Evolved Release 21.4R1.

The local-link sub-option of the target option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The inline-timestamping suboption of the offload-type option introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

[Understand Two-Way Active Measurement Protocol](#) | 686

test-type (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax](#) | 1483
- [Hierarchy Level](#) | 1483
- [Description](#) | 1483
- [Options](#) | 1484
- [Required Privilege Level](#) | 1484
- [Release Information](#) | 1484

Syntax

```
test-type (throughput | latency | frame-loss | back-back-frames);
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarking profiles test-profile profile-name]
```

Description

RFC 2544 defines four main test types. You can configure and perform a test for a certain service, such as IPv4 or Ethernet, and analyze the results of the test to examine the various SLA parameters of the service. The test packets traverse through the same path as the regular service traffic.

Configure the type of RFC 2544-based benchmarking test to be performed. Because of the ability of these tests to measure throughput, bursty frames, frame loss, and latency, this mechanism is also used to diagnose and examine Ethernet-based networks.

Options

throughput	Measure the maximum rate at which none of the offered or transmitted frames are dropped by the device on which the test is performed.
latency	Measure the time interval between the arrival of the last bit of the input frame at the input port and the output of the first bit of the frame on the output port.
frame-loss	Measure the percentage of frames that must have been forwarded by a network device under steady state (constant) load conditions, but were not forwarded due to lack of resources.
back-back-frames	Measure the number of frames that are forwarded by the device on which the test is performed when a burst of frames with minimum inter-frame gaps is sent to that device from another source device.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

thresholds (Junos OS)

IN THIS SECTION

- [Syntax | 1485](#)
- [Hierarchy Level | 1485](#)
- [Description | 1485](#)
- [Options | 1486](#)
- [Required Privilege Level | 1486](#)
- [Release Information | 1486](#)

Syntax

```
thresholds thresholds;
```

Hierarchy Level

```
[edit services rpm probe owner test test-name],  
[edit services rpm twamp client control-connection control-client-name]
```

Description

Specify thresholds used for the probes. A system log message is generated when the configured threshold is exceeded. Likewise, an SNMP trap (if configured) is generated when a threshold is exceeded.

NOTE: If you configure a value of zero using the *thresholds* option for a certain probe parameter, the generation of SNMP traps is disabled for the corresponding probe attribute. For example, if you specify the set thresholds jitter-egress 0 statement, it denotes that traps are not triggered when the jitter in egress time threshold is met or exceeded.

Options

thresholds—Specify one or more threshold measurements. The following options are supported:

- *egress-time*—Measures maximum source-to-destination time per probe.
- *ingress-time*—Measures maximum destination-to-source time per probe.
- *jitter-egress*—Measures maximum source-to-destination jitter per test.
- *jitter-ingress*—Measures maximum destination-to- source jitter per test.
- *jitter-rtt*—Measures maximum jitter per test, from 0 through 60,000,000 microseconds.
- *rtt*—Measures maximum round-trip time per probe, in microseconds.
- *std-dev-egress*—Measures maximum source-to-destination standard deviation per test.
- *std-dev-ingress*—Measures maximum destination-to-source standard deviation per test.
- *std-dev-rtt*—Measures maximum standard deviation per test, in microseconds.
- *successive-loss*—Measures successive probe loss count, indicating probe failure.
- *total-loss*—Measures total probe loss count indicating test failure, from 0 through 15.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit services rpm twamp client control-connection *control-client-name*] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

thresholds (Junos OS Evolved)

IN THIS SECTION

- [Syntax | 1487](#)
- [Hierarchy Level | 1487](#)
- [Description | 1487](#)
- [Options | 1488](#)
- [Required Privilege Level | 1488](#)
- [Release Information | 1489](#)

Syntax

```
thresholds {  
    control-failure (on | off);  
    successive-loss number;  
    total-loss number;  
    threshold-type (microseconds | average);  
}
```

Hierarchy Level

```
[edit services monitoring rpm owner name test name]  
[edit services monitoring twamp client control-connection name test-session name]
```

Description

Specify the thresholds used for real-time performance monitoring (RPM) and Two-Way Active Measurement Protocol (TWAMP) probes. A system log message is generated when the configured threshold is exceeded.

Options

control-failure (on off)	Specify whether control failure should be reported as a test failure. <ul style="list-style-type: none"> • Default: on
successive-loss number	Measures the number of packets lost in succession. Set this number to specify how many packets should be lost in succession before indicating a probe failure. <ul style="list-style-type: none"> • Range: 0 through 15; setting this option to zero disables this counter. • Default: 1, which indicates that the test should fail with any packet loss.
total-loss number	Measures total number of packets lost per test. Set this number to specify how many packets should be lost in total before indicating a probe failure. <ul style="list-style-type: none"> • Range: 0 through 15; setting this option to zero disables this counter. • Default: 1, which indicates that the test should fail with any packet loss.
threshold-type (microseconds average)	Specify one or more measurement options and set thresholds for them. By default, these options set the maximum threshold. To configure the option to use the average measurement and not the maximum, configure the average option on each configured threshold type. To set the threshold type back to maximum, either delete the average option from the configuration or configure the threshold type again without the average option. <ul style="list-style-type: none"> • Range: 0 through 60,000,000 microseconds; setting this option to zero disables the configured threshold type. • Values: You can configure these threshold types: <ul style="list-style-type: none"> • egress-time—Source-to-destination time per probe. • ingress-time—Destination-to-source time per probe. • jitter-egress—Source-to-destination jitter per test. • jitter-ingress—Destination-to- source jitter per test. • jitter-rtt—Round-trip jitter per test. • rtt—Round-trip time per probe.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Evolved 20.1R1 for RPM.

Statement introduced in Junos OS Evolved 20.3R1 for TWAMP.

RELATED DOCUMENTATION

| [Understand Two-Way Active Measurement Protocol](#) | [686](#)

timestamp-format (RFC 2544 Benchmarking)

IN THIS SECTION

- [Syntax](#) | [1489](#)
- [Hierarchy Level](#) | [1489](#)
- [Description](#) | [1490](#)
- [Required Privilege Level](#) | [1490](#)
- [Release Information](#) | [1490](#)

Syntax

```
timestamp-format;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specify the time stamp format.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

traceoptions (Dynamic Flow Capture)

IN THIS SECTION

- [Syntax | 1491](#)
- [Hierarchy Level | 1491](#)
- [Description | 1491](#)
- [Options | 1491](#)
- [Required Privilege Level | 1491](#)
- [Release Information | 1492](#)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | non-world-readable>;
}
```

Hierarchy Level

```
[edit services dynamic-flow-capture]
```

Description

Enable and define tracing options for dynamic flow capture events.

Options

- | | |
|-----------------------------|--|
| file <i>filename</i> | Use the specified file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log . |
| files <i>number</i> | <p>(Optional) Use the specified maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number for files, you must also specify a maximum file size with the size option.</p> <ul style="list-style-type: none"> • Range: 2 through 1000 files. • Default: 10 files. |
| no-world-readable | (Optional) Restrict access to the file. |
| world-readable | (Optional) Enable free access to the file. |

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.2.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

traceoptions (Forwarding Options)

IN THIS SECTION

- [Syntax](#) | 1492
- [Hierarchy Level](#) | 1492
- [Description](#) | 1493
- [Required Privilege Level](#) | 1493
- [Release Information](#) | 1493

Syntax

```
traceoptions {  
    no-remote-trace;  
    file filename <files number> <size bytes> <match expression> <world-readable | no-world-readable>;  
}
```

Hierarchy Level

```
[edit forwarding-options port-mirroring],  
[edit forwarding-options sampling]
```

Description

Configure traffic sampling tracing operations.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

traceoptions (Inline Monitoring)

IN THIS SECTION

- [Syntax](#) | 1494
- [Hierarchy Level](#) | 1494
- [Description](#) | 1494
- [Options](#) | 1494
- [Additional Information](#) | 1495
- [Required Privilege Level](#) | 1495
- [Release Information](#) | 1496

Syntax

```
traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-
readable)>;
    no-remote-trace;
}
```

Hierarchy Level

```
[edit services inline-monitoring]
```

Description

(Junos OS only) Configure traceoptions for the inline monitoring process.

Options

file <i>file-name</i>	Use the specified file to receive the output of the tracing operation. All files are placed in the directory /var/log .
files <i>files</i>	<p>(Optional) Use the specified 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.</p> <ul style="list-style-type: none"> • Default: 3 files • Range: 2 through 1000 files
match <i>match</i>	(Optional) Use the specified regular expression to refine the output to include lines that contain the regular expression.
no-remote-trace	Disable remote tracing.
no-world-readable	(Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.
world-readable	(Optional) Enable unrestricted file access.

size *size* (Optional) Use the specified 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.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

Additional Information

Junos OS Evolved uses a new tracing architecture. All running applications create trace information, with multiple instances of the same application having their own trace information. In Junos OS, you enable tracing operations by configuring the `traceoptions` statement at the specific hierarchy level you want to trace. Junos OS Evolved, on the other hand, uses an application-based model, and thus trace messages are logged, viewed, and configured by application. As a result, Junos OS Evolved does not support the `traceoptions` statement at many of the hierarchy levels that Junos OS supports.

In Junos OS Evolved, you do not view trace files directly, and you should never add, edit, or remove trace files under the `/var/log/traces` directory because this can corrupt the traces. Instead, you use the `show trace application application-name node node-name` command to read and decode trace messages stored in the trace files. All running applications on Junos OS Evolved create trace information at the `info` level by default.

Inline monitoring services are governed by the `imond` application. For Junos OS Evolved, to configure traces for a severity other than `info` for the `imond` application, include the `application imond node node-name level severity` statement at the `[edit system trace]` hierarchy level.

NOTE: For general monitoring and troubleshooting of devices running Junos OS or Junos OS Evolved, we recommend using standard tools such as CLI `show` commands, system log messages, SNMP, and telemetry data. You should avoid using trace messages for general debugging purposes and long-term solutions because they are subject to change without notice.

Required Privilege Level

`trace`—To view this statement in the configuration.

`trace-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS 19.4R1.

RELATED DOCUMENTATION

[Understanding Inline Monitoring Services](#) | 334

traceoptions (Resiliency)

IN THIS SECTION

- [Syntax](#) | 1496
- [Hierarchy Level](#) | 1497
- [Description](#) | 1497
- [Options](#) | 1497
- [Additional Information](#) | 1498
- [Required Privilege Level](#) | 1498
- [Release Information](#) | 1498

Syntax

```
traceoptions {  
  file name {  
    files number;  
    match;  
    (no-world-readable | world-readable);  
    size size;  
  }  
  flag flag;  
  no-remote-trace;  
}
```

Hierarchy Level

[edit system [resiliency](#)]

Description

(Junos OS only) Configure trace options for the Juniper Resiliency Interface.

Options

- | | |
|------------------------------|--|
| file <i>file-name</i> | (Required) Use the specified file to receive the output of the tracing operation. All files are placed in the directory /var/log . |
| files <i>files</i> | <p>(Optional) Use the specified 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.</p> <ul style="list-style-type: none"> • Default: 3 files • Range: 2 through 1000 files |
| flag <i>flag</i> | <p>Use the specified tracing operation. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:</p> <ul style="list-style-type: none"> • error—Trace errored packets. • state—Trace state transitions. • telemetry—Trace telemetry state machine events. |
| match <i>match</i> | (Optional) Use the specified regular expression to refine the output to include lines that contain the regular expression. |
| no-remote-trace | (Optional) Disable remote tracing. |
| no-world-readable | (Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation. |
| world-readable | (Optional) Enable unrestricted file access. |
| size <i>size</i> | (Optional) Use the specified 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.

- **Range:** 10 KB through 1 GB
- **Default:** 128 KB

Additional Information

Junos OS Evolved uses a new tracing architecture. All running applications create trace information, with multiple instances of the same application having their own trace information. In Junos OS, you enable tracing operations by configuring the `traceoptions` statement at the specific hierarchy level you want to trace. Junos OS Evolved, on the other hand, uses an application-based model, and thus trace messages are logged, viewed, and configured by application. As a result, Junos OS Evolved does not support the `traceoptions` statement at many of the hierarchy levels that Junos OS supports.

In Junos OS Evolved, you do not view trace files directly, and you should never add, edit, or remove trace files under the `/var/log/traces` directory because this can corrupt the traces. Instead, you use the `show trace application application-name node node-name` command to read and decode trace messages stored in the trace files. All running applications on Junos OS Evolved create trace information at the `info` level by default.

The Juniper Resiliency Interface is governed by the `rpdtmd` application. For Junos OS Evolved, to configure traces for a severity other than `info` for the `rpdtmd` application, include the `application imond node node-name level severity` statement at the `[edit system trace]` hierarchy level.

NOTE: For general monitoring and troubleshooting of devices running Junos OS or Junos OS Evolved, we recommend using standard tools such as CLI `show` commands, system log messages, SNMP, and telemetry data. You should avoid using trace messages for general debugging purposes and long-term solutions because they are subject to change without notice.

Required Privilege Level

`trace`—To view this statement in the configuration.

`trace-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 21.2R1.

RELATED DOCUMENTATION

[Juniper Resiliency Interface](#) | 409

traceoptions (RPM)

IN THIS SECTION

- [Syntax](#) | 1499
- [Hierarchy Level](#) | 1499
- [Description](#) | 1499
- [Options](#) | 1500
- [Additional Information](#) | 1501
- [Required Privilege Level](#) | 1501
- [Release Information](#) | 1501

Syntax

```
traceoptions {  
    file filename <files number> <match regular-expression> <size maximum-file-size> <world-  
readable | no-world-readable>;  
    flag flag;  
}
```

Hierarchy Level

```
[edit services rpm]
```

Description

(Junos OS only) Define tracing operations for RPM processes.

Options

file <i>filename</i>	Use the specified file to receive the output of the tracing operation. All files are placed in the directory <code>/var/log</code> . <ul style="list-style-type: none"> • Default: <code>rmopd</code>
files <i>number</i>	(Optional) Use the specified 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 <code>size</code> option. <ul style="list-style-type: none"> • Range: 2 through 1000 • Default: 3 files
match <i>regular-expression</i>	(Optional) Use the specified regular expression to refine the output to include lines that contain the regular expression.
size <i>maximum-file-size</i>	(Optional) Use the specified 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 <code>files</code> option. <ul style="list-style-type: none"> • Range: 10 KB through 1 GB • Default: 128 KB
world-readable	(Optional) Enable unrestricted file access.
no-world-readable	(Default) Disable unrestricted file access. This means the log file can be accessed only by the user who configured the tracing operation.
flag <i>flag</i>	Use the specified tracing operation. To specify more than one tracing operation, include multiple <code>flag</code> statements. You can include the following flags: <ul style="list-style-type: none"> • <code>all</code>—Trace all operations. • <code>configuration</code>—Trace configuration events. • <code>error</code>—Trace events related to catastrophic errors in daemon. • <code>ipc</code>—Trace IPC events. • <code>ppm</code>—Trace ppm events. • <code>rpdp</code>—Trace rpd events.

- statistics—Trace statistics.

Additional Information

Junos OS Evolved uses a new tracing architecture. All running applications create trace information, with multiple instances of the same application having their own trace information. In Junos OS, you enable tracing operations by configuring the `traceoptions` statement at the specific hierarchy level you want to trace. Junos OS Evolved, on the other hand, uses an application-based model, and thus trace messages are logged, viewed, and configured by application. As a result, Junos OS Evolved does not support the `traceoptions` statement at many of the hierarchy levels that Junos OS supports.

In Junos OS Evolved, you do not view trace files directly, and you should never add, edit, or remove trace files under the `/var/log/traces` directory because this can corrupt the traces. Instead, you use the `show trace application application-name node node-name` command to read and decode trace messages stored in the trace files. All running applications on Junos OS Evolved create trace information at the `info` level by default.

RPM is governed by the `rmopd` application. For Junos OS Evolved, to configure traces for a severity other than `info` for the `rmopd` application, include the `application rmopd node node-name level severity` statement at the `[edit system trace]` hierarchy level.

NOTE: For general monitoring and troubleshooting of devices running Junos OS or Junos OS Evolved, we recommend using standard tools such as CLI `show` commands, system log messages, SNMP, and telemetry data. You should avoid using trace messages for general debugging purposes and long-term solutions because they are subject to change without notice.

Required Privilege Level

`trace`—To view this statement in the configuration.

`trace-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

`rp` flag added in Junos OS Release 18.4R1.

RELATED DOCUMENTATION

| [Trace RPM Operations](#) | [676](#)

transfer

IN THIS SECTION

- [Syntax](#) | [1502](#)
- [Hierarchy Level](#) | [1502](#)
- [Description](#) | [1502](#)
- [Options](#) | [1502](#)
- [Required Privilege Level](#) | [1503](#)
- [Release Information](#) | [1503](#)

Syntax

```
transfer {  
    record-level number;  
    timeout seconds;  
}
```

Hierarchy Level

```
[edit services flow-collector file-specification variant variant-number]
```

Description

Specify when to send the flow collection file. The file is sent when either of the two conditions is met.

Options

`record-level number`—Use the specified number of flow collection files collected.

`timeout seconds`—Use the specified timeout duration.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

transfer-log-archive

IN THIS SECTION

- [Syntax](#) | 1503
- [Hierarchy Level](#) | 1504
- [Description](#) | 1504
- [Required Privilege Level](#) | 1504
- [Release Information](#) | 1504

Syntax

```
transfer-log-archive {
  archive-sites {
    ftp:url {
      password "password";
      username username;
    }
  }
}
```

```

    }
    filename-prefix prefix;
    maximum-age minutes;
  }

```

Hierarchy Level

```
[edit services flow-collector]
```

Description

Configure the filename prefix, maximum age, and destination FTP server for log files containing the transfer activity history for a flow collector interface.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

transmit-failure-threshold (RFC 2544 Benchmarking)

IN THIS SECTION

● [Syntax](#) | 1505

- [Hierarchy Level | 1505](#)
- [Description | 1505](#)
- [Required Privilege Level | 1505](#)
- [Release Information | 1505](#)

Syntax

```
transmit-failure-threshold;
```

Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Description

Specifies the failure threshold value of the transmit test frames.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

RELATED DOCUMENTATION

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Configuring RFC 2544-Based Benchmarking Tests | 864](#)

[rfc2544-benchmarking | 1347](#)

traps

IN THIS SECTION

- [Syntax | 1506](#)
- [Junos OS Hierarchy Level | 1506](#)
- [Junos OS Evolved Hierarchy Level | 1506](#)
- [Description | 1507](#)
- [Options | 1507](#)
- [Additional Information | 1508](#)
- [Required Privilege Level | 1508](#)
- [Release Information | 1509](#)

Syntax

```
traps traps;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name],  
[edit services rpm twamp client control-connection control-client-name]  
[edit services rpm twamp client control-connection control-client-name test-session session-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner test test-name]  
[edit services monitoring twamp client control-connection control-client-name]  
[edit services monitoring twamp client control-connection control-client-name test-session session-name]
```

Description

Set the trap bit to generate SNMP traps for probes. Traps are sent if the configured threshold is met or exceeded. For TWAMP, you can configure both control-connection and test-session traps for managed clients, but only test-session traps for light clients. See [SNMP MIB Explorer](#) to learn more about the SNMP MIB objects that Juniper supports.

Default: no traps are generated

Options

traps Specify one or more SNMP traps. The options are configured per test session, except for the control-connection-closed and the test-iteration-done options. These options are configured only for TWAMP, and only per control connection. The following options are supported:

- control-connection-closed—(TWAMP control-connection only) Generate traps when the control connection is closed.
- egress-jitter-exceeded—Generate traps when the jitter in egress time threshold is met or exceeded. (MIB object jnxPingEgressJitterThresholdExceeded)
- egress-std-dev-exceeded—(Junos OS only) Generate traps when the egress time standard deviation threshold is met or exceeded. (MIB object jnxPingEgressStdDevThresholdExceeded)
- egress-time-exceeded—Generate traps when the maximum egress time threshold is met or exceeded. (MIB object jnxPingEgressThresholdExceeded)
- ingress-jitter-exceeded—Generate traps when the jitter in ingress time threshold is met or exceeded. (MIB object jnxPingIngressJitterThresholdExceeded)
- ingress-std-dev-exceeded—(Junos OS only) Generate traps when the ingress time standard deviation threshold is met or exceeded. (MIB object jnxPingIngressStddevThresholdExceeded)
- ingress-time-exceeded—Generate traps when the maximum ingress time threshold is met or exceeded. (MIB object jnxPingIngressThresholdExceeded)
- jitter-exceeded—Generate traps when the jitter in round-trip time threshold is met or exceeded. (MIB object jnxPingRttJitterThresholdExceeded)
- max-rtt-threshold—(Junos OS only) Generate traps when the maximum round trip time threshold at the end of the test is met or exceeded. (MIB object jnxPingMaxRttThresholdExceeded)

- **probe-failure**—Generate traps when successive probe loss thresholds are crossed. (MIB object `pingProbeFailed`)
- **rtt-exceeded**—Generate traps when the maximum round-trip time threshold is met or exceeded. (MIB object `jnxPingRttThresholdExceeded`)
- **std-dev-exceeded**—(Junos OS only) Generate traps when the round-trip time standard deviation threshold is met or exceeded. (MIB object `jnxPingRttStdDevThresholdExceeded`)
- **test-completion**—Generate traps when a test is completed. (MIB object `pingTestCompleted`)
- **test-failure**—Generate traps when the total probe loss threshold is met or exceeded. (MIB object `pingTestFailed`)
- **test-iteration-done**—(TWAMP control-connection only) Generate traps when all test sessions under control connections complete one test iteration.

NOTE: To generate RPM traps, you must configure the `remote-operations` SNMP trap category by including the `categories` statement at the `[edit snmp trap-group trap-group-name]` hierarchy level.

Additional Information

(Junos OS Evolved only) The Juniper PING MIB does not support the following notifications, because the corresponding thresholds are not configurable for Junos OS Evolved and hence the traps are never generated:

- `jnxPingRttStdDevThresholdExceeded`
- `jnxPingEgressStdDevThresholdExceeded`
- `jnxPingIngressStddevThresholdExceeded`
- `jnxPingMaxRttThresholdExceeded`

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Support at the [edit services rpm twamp client control-connection *control-client-name*] and [edit services rpm twamp client control-connection *control-client-name* test-session *session-name*] hierarchy levels introduced in Junos OS Release 15.1 for MX Series routers.

Support at the [edit services monitoring rpm owner test *test-name*] hierarchy level introduced in Junos OS Evolved Release 21.2R1 for PTX Series routers.

Support at the [edit services monitoring twamp client control-connection *control-client-name*] and [edit services monitoring twamp client control-connection *control-client-name* test-session *session-name*] hierarchy levels introduced in Junos OS Evolved Release 21.3R1 for PTX Series routers.

RELATED DOCUMENTATION

[thresholds \(Junos OS\) | 1485](#)

[thresholds \(Junos OS Evolved\) | 1487](#)

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches | 651](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

ttl

IN THIS SECTION

- [Syntax | 1510](#)
- [Hierarchy Level | 1510](#)
- [Description | 1510](#)
- [Options | 1510](#)
- [Required Privilege Level | 1510](#)
- [Release Information | 1510](#)

Syntax

```
ttl hops;
```

Hierarchy Level

```
[edit services dynamic-flow-capture capture-group client-name content-destination identifier]
```

Description

Configure the time-to-live (TTL) value for the IP-IP header.

Options

hops—TTL value.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Junos Capture Vision](#) | 289

ttl (RPM probe)

IN THIS SECTION

- [Syntax | 1511](#)
- [Junos OS Hierarchy Level | 1511](#)
- [Junos OS Evolved Hierarchy Level | 1511](#)
- [Description | 1511](#)
- [Options | 1512](#)
- [Required Privilege Level | 1513](#)
- [Release Information | 1513](#)

Syntax

```
ttl hop-count;
```

Junos OS Hierarchy Level

```
[edit services rpm probe owner test test-name ]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm owner name test name]
```

Description

Specify the maximum hop count (TTL) for all types of probes (IPv4 and IPv6) in real-time performance monitoring (RPM) and Two-Way Active Management Protocol (TWAMP). This can be useful when it is necessary to restrict the scope of the RPM probes so a probe does not unintentionally monitor an alternative path to the destination (such as may occur after a BGP re-routing). Another example is to monitor direct reachability by specifying a TTL of 1. Probes that exceed the number set for TTL are discarded.

The TTL configuration is supported on Routing Engine-based RPM, Multiservices Modular PIC Concentrator (MS-MPC) and Multiservices Modular Interfaces Card (MS-MIC)-based RPM, and Two-Way Active Management Protocol (TWAMP).

You can set a TTL value for the probe types listed below.

Software time stamping:

- icmp-ping, and icmp-ping-timestamp
- icmp6-ping (Junos OS only)
- udp-ping, and udp-ping-timestamp
- tcp-ping (Junos OS only)
- http-get, and http-metadata-get (Junos OS only)
- BGP neighbor monitoring using TCP/UDP (Junos OS only)

Hardware time stamping:

- icmp-ping and icmp-ping-timestamp
- udp-ping and udp-ping-timestamp

MS-MPC-PIC based probes (delegate):

- icmp-ping, icmp-ping-timestamp, and icmp6-ping (Junos OS only),

MS-MPC-PIC hardware timestamp:

- icmp-ping and icmp-ping-timestamp
- udp-ping and udp-ping-timestamp

TWAMP probe:

- inline TWAMP client

Options

hop-count—Prior to Junos OS Release 18.2R1, for RPM, the RPM client always sent a TTL of 64 to the RPM server under the IPv4 or IPv6 header. For TWAMP clients, the TTL was 255 sent in the IPv4 header. In Junos OS Release 18.2R1 and later, you can specify the TTL you want for RPM and TWAMP probes.

- **Range:** 1 through 255 for both RPM and TWAMP
- **Default:** 64

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.2R1 for MX Series routers.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

[Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches](#) | 651

tunnel-observation

IN THIS SECTION

- [Syntax](#) | 1513
- [Hierarchy Level](#) | 1514
- [Description](#) | 1514
- [Options](#) | 1514
- [Required Privilege Level](#) | 1515
- [Release Information](#) | 1515

Syntax

```
tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
```

Hierarchy Level

```
[edit services flow-monitoring version9 template template-name]
[edit services flow-monitoring version-ipfix template template-name]
```

Description

Specify the types of MPLS flows on which to enable inline flow monitoring. If you do not configure a `tunnel-observation` type, then plain MPLS flow records are created.

You can configure multiple values for `tunnel-observation`. Flows are created for only the deepest match. For example, if you configure both `ipv4` and `mpls-over-udp` and the traffic type is MPLS-over-UDP, flows are created for MPLS-over-UDP. If you configure `ipv4` but *do not* configure `mpls-over-udp` and the traffic type is MPLS-over-UDP, flows are created for MPLS-IPv4.

If the MPLS traffic type does *not* match any of the `tunnel-observation` values, then plain MPLS flows are created.

If you do not configure `tunnel-observation`, plain MPLS flows are created.

If the `tunnel-observation` statement is added or deleted, or if the configured value is changed, all flows related to the old template will be deleted and replaced by new flows using the changed template.

Options

- ipv4** Enable flow monitoring for MPLS-IPv4 traffic. You must also configure `mpls-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template template-name]` hierarchy level.
- ipv6** Enable flow monitoring for MPLS-IPv6 traffic. You must also configure `mpls-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template template-name]` hierarchy level. If you are running inline flow monitoring on a Lookup (LU) card on an MX Series router, you must also configure `use-extended-flow-memory` at the `[edit chassis fpc slot-number inline-services]` hierarchy level to create MPLS-IPv6 flow records.
- mpls-over-udp** (PTX Series only) Enable flow monitoring for MPLS-over-UDP traffic. Monitoring looks past the tunnel header to report the inner payload of the packets. For an MPLS-over-UDP flow that is carried between IPv4 endpoints, you must also configure `ipv4-template` at the `[edit services flow-monitoring (version9 | version-ipfix) template template-name]` hierarchy level. For an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP, you must also configure `mpls-ipvx-template` in Junos OS Release 18.1 or `mpls-template` starting in Junos OS 18.2R1 at the `[edit services flow-monitoring (version 9 | version-ipfix) template template-name]` hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 18.1R1.

ipv4 and ipv6 options added in Junos OS Release 18.2R1.

Statement introduced in Junos OS Release 18.4R1 on MX Series routers.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring on PTX Series Routers | 540](#)

[Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers | 550](#)

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

twamp

IN THIS SECTION

- [Syntax \(Junos OS\) | 1516](#)
- [Syntax \(Junos OS Evolved\) | 1518](#)
- [Hierarchy Level \(Junos OS\) | 1519](#)
- [Hierarchy Level \(Junos OS Evolved\) | 1520](#)
- [Description | 1520](#)
- [Required Privilege Level | 1520](#)
- [Release Information | 1520](#)

Syntax (Junos OS)

```

twamp {
  client {
    control-connection name {
      authentication-mode none;
      control-type (managed | light);
      destination-interface destination-interface;
      destination-port destination-port;
      history-size history-size;
      moving-average-size moving-average-size;
      persistent-results;
      routing-instance routing-instance;
      target-address target-address;
      tcp-keepcnt count;
      tcp-keepidle seconds;
      tcp-keepintvl seconds;
      test-count test-count;
      test-interval seconds;
      traps {
        control-connection-closed;
        test-iteration-done;
      }
    }
    test-session name {
      data-fill-with zeros;
      data-size data-size;
      destination-port destination-port;
      dscp-code-points (RPM) dscp-code-points;
      probe-count probe-count;
      probe-interval seconds;
      source-address source-address;
      target-address target-address local-link IPv6-link-local-interface-name;
      thresholds {
        egress-time microseconds;
        ingress-time microseconds;
        jitter-egress microseconds;
        jitter-ingress microseconds;
        jitter-rtt microseconds;
        max-rtt microseconds;
        rtt microseconds;
        std-dev-egress microseconds;
        std-dev-ingress microseconds;
      }
    }
  }
}

```

```

        std-dev-rtt microseconds;
        successive-loss successive-loss;
        total-loss total-loss;
    }
    traps {
        egress-jitter-exceeded;
        egress-std-dev-exceeded;
        egress-time-exceeded;
        ingress-jitter-exceeded;
        ingress-std-dev-exceeded;
        ingress-time-exceeded;
        jitter-exceeded;
        max-rtt-exceeded;
        probe-failure;
        rtt-exceeded;
        std-dev-exceeded;
        test-completion;
        test-failure;
    }
    ttl hop-count;
}
}
}
post-cli-implicit-firewall;
server {
    authentication-key-chain name {
        key-id name {
            secret secret;
        }
    }
    authentication-mode <authenticated> <control-only-encrypted> <encrypted> <none>;
    client-list <list-name> {
        address address <routing-instance [instance-name...]>;
    }
    light {
        port port-number;
    }
    max-connection-duration hours;
    maximum-connections maximum-connections;
    maximum-connections-per-client maximum-connections-per-client;
    maximum-sessions maximum-sessions;
    maximum-sessions-per-connection maximum-sessions-per-connection;

```



```

    port port;
    routing-instance-list name {
        port port;
    }
    server-inactivity-timeout minutes;
    tcp-keepcnt count;
    tcp-keepidle seconds;
    tcp-keepintvl seconds;
}
}

```

Syntax (Junos OS Evolved)

```

twamp {
    client {
        control-connection name {
            control-type (managed | light);
            destination-port destination-port;
            routing-instance routing-instance;
            source-address source-address;
            target target-address;
            test-start (auto | manual);
            test-interval seconds;
            traps {
                control-connection-closed;
                test-iteration-done;
            }
            test-session name {
                data-size data-size;
                destination-port destination-port;
                dscp-code-points dscp-code-points;
                history-size history-size;
                moving-average-size moving-average-size;
                offload-type (none | inline-timestamping | pfe-timestamp);
                probe-count probe-count;
                probe-interval seconds;
                source-address source-address;
                target target-address local-link IPv6-link-local-interface-name ;

                thresholds {
                    control-failure (on | off);

```

```

        successive-loss number;
        total-loss number;
        threshold-type (microseconds | average);
    }
    traps {
        egress-jitter-exceeded
        egress-time-exceeded;
        ingress-jitter-exceeded;
        ingress-time-exceeded;
        jitter-exceeded;
        probe-failure;
        rtt-exceeded;
        test-completion;
        test-failure;
    }
    ttl hop-count;
}
}
}
server {
    managed {
        client-limit limit;
        client-list list-name {
            address address <routing-instance [instance-name...]>;
        }
        control-inactivity-timeout seconds;
        control-per-client-limit limit;
        control-maximum-duration seconds;
        port port;
        test-per-client-limit limit;
    }
    light {
        port port;
    }
}
}
}

```

Hierarchy Level (Junos OS)

```
[edit services rpm]
```

Hierarchy Level (Junos OS Evolved)

[edit services monitoring]

Description

Configure the Two-Way Active Measurement Protocol (TWAMP) client or server settings on all M Series and T Series routers that support Multiservices PICs (running in either Layer 2 or Layer 3 mode), MX Series routers, ACX routers, EX4300 Switches, PTX Series routers, EX9200 Series switches, and QFX10000 Series switches.

TWAMP is an open protocol for measurement of two-way metrics. The host that initiates the TCP connection takes the roles of the control-client and (in the two-host implementation) the session-sender. Such a device is also called the TWAMP client. The host that acknowledges the TCP connection accepts the roles of a server and (in the two-host implementation) and the session-reflector. Such a device is also called the TWAMP server. The TWAMP-Test messages are exchanged between the session-sender and the session-reflector, and the TWAMP-Control messages are exchanged between the control-client and the server.

The following addresses cannot be used for the `client-list` source IP address used for probes:

- 0.0.0.0
- 127.0.0.0/8 (loopback)
- 224.0.0.0/4 (multicast)
- 255.255.255.255 (broadcast)

The remaining statements are explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

Junos OS Evolved syntax and hierarchy level introduced in Junos OS Evolved 20.3R1 for PTX10003 routers.

Support for the rest of the PTX10000 Series routers added in Junos OS Evolved Release 21.1R1.

control-type and light options added in Junos OS Release 21.1R1.

traps option added in Junos OS Evolved Release 21.3R1.

source-address option and the local-link sub-option of the target-address option for TWAMP Light test sessions introduced in Junos OS Release 21.4R1.

The local-link sub-option of the target option for TWAMP Light test sessions introduced in Junos OS Evolved Release 22.3R1.

The inline-timestamping suboption of the offload-type option introduced in Junos OS Evolved 22.4R1.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

twamp-server

IN THIS SECTION

- [Syntax](#) | 1521
- [Hierarchy Level](#) | 1522
- [Description](#) | 1522
- [Required Privilege Level](#) | 1522
- [Release Information](#) | 1522

Syntax

```
twamp-server;
```

Hierarchy Level

```
[edit interfaces sp-fpc/pic/port unit logical-unit-number]
```

Description

Specify the service PIC logical interface to provide the TWAMP service.

Required Privilege Level

system—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 9.3.

RELATED DOCUMENTATION

[Configure TWAMP on ACX, MX, M, T, and PTX Series Routers, EX4300 Series, EX9200 Series, and QFX10000 Series Switches](#) | 694

trio-flow-offload

IN THIS SECTION

- [Syntax](#) | 1523
- [Hierarchy Level](#) | 1523
- [Description](#) | 1523
- [Options](#) | 1523
- [Required Privilege Level](#) | 1523
- [Release Information](#) | 1523

Syntax

```
trio-flow-offload minimum-bytes minimum-bytes;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Description

Enable any plug-in or daemon on a PIC to generate a request to off-load flows to the Packet Forwarding Engine. This command is available on MX Series routers with Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs).

NOTE: This feature is not supported for Broadband Edge subscribers (given that service PIC off load is not available with aggregate Ethernet (AE)).

Options

minimum-bytes—Minimum number of bytes that trigger offloading. When this option is omitted, offloading is triggered when both the forward and reverse flows of the session have begun, meaning that at least one packet has flowed in each direction.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.1.

RELATED DOCUMENTATION

| [Configuring Flow Offloading on MX Series Routers](#)

udp

IN THIS SECTION

- [Junos OS | 1524](#)
- [Junos OS Evolved | 1524](#)
- [Junos OS Hierarchy Level | 1524](#)
- [Junos OS Evolved Hierarchy Level | 1525](#)
- [Description | 1525](#)
- [Required Privilege Level | 1525](#)
- [Release Information | 1525](#)

Junos OS

```
udp {  
    destination-interface interface-name;  
    port port;  
}
```

Junos OS Evolved

```
udp {  
    port port;  
}
```

Junos OS Hierarchy Level

```
[edit services rpm probe-server]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rpm probe-server]
```

Description

Enable UDP requests for the RPM probe 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.

NOTE: The destination-interface statement is not supported on PTX Series routers or for Junos OS Evolved.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

Statement introduced in Junos OS Evolved Release 20.1R1.

RELATED DOCUMENTATION

| [Configuring RPM Receiver Servers](#) | 662

udp-tcp-port-swap

IN THIS SECTION

- [Syntax | 1526](#)
- [Junos OS Hierarchy Level | 1526](#)
- [Junos OS Evolved Hierarchy Level | 1526](#)
- [Description | 1526](#)
- [Required Privilege Level | 1526](#)
- [Release Information | 1527](#)

Syntax

```
udp-tcp-port-swap;
```

Junos OS Hierarchy Level

```
[edit services rpm rfc2544-benchmarkingtests test-name test-name]
```

Junos OS Evolved Hierarchy Level

```
[edit services monitoring rfc2544 tests test-name test-name]
```

Description

Swap the source and destination UDP ports in the test packets. Only UDP port swap and UDP over IPv4 traffic is supported.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 12.3X53.

Statement introduced in Junos OS Evolved Release 22.4R1.

RELATED DOCUMENTATION

[rfc2544-benchmarking | 1347](#)

[rfc2544 | 1345](#)

[RFC 2544-Based Benchmarking Tests for ACX Routers Overview | 855](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

unit

IN THIS SECTION

- [Syntax | 1527](#)
- [Hierarchy Level | 1528](#)
- [Description | 1528](#)
- [Options | 1528](#)
- [Required Privilege Level | 1528](#)
- [Release Information | 1528](#)

Syntax

```
unit logical-unit-number {
  family inet {
    address address {
      destination destination-address;
    }
  }
}
```

```

    }
    filter {
        group filter-group-number;
        input filter-name;
        output filter-name;
    }
    sampling direction;
}
}

```

Hierarchy Level

```
[edit interfaces interface-name]
```

Description

Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options

logical-unit-number—Number of the logical unit.

- **Range:** 0 through 16,384

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Junos OS Network Interfaces Library for Routing Devices](#)

[Junos OS Network Interfaces Library for Routing Devices](#)

use-extended-flow-memory

IN THIS SECTION

- [Syntax | 1529](#)
- [Hierarchy Level | 1529](#)
- [Description | 1529](#)
- [Required Privilege Level | 1530](#)
- [Release Information | 1530](#)

Syntax

```
use-extended-flow-memory;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services]
```

Description

Configure the service to extended flow memory. This service provides more scale in flows for inline services sampling.

The new configuration set `chassis fpc slot slot-number inline-services use-extended-flow-memory` allows you to configure table to operate in side band mode with side band memory. This configuration is applicable only on a Lookup (LU) platform. It is not applicable for XL line card because XL has dedicated DMEM memory to hold 64M flow entries.

If you are configuring inline flow monitoring of MPLS-IPv6 flows on an LU platform, you must configure `use-extended-flow-memory` to get MPLS-IPv6 flow records. If you do not configure `use-extended-flow-memory` on an LU platform, plain MPLS flow records are created.

NOTE: This configuration is supported only on LU platforms. The LU platform line cards are MPC1E, MPC2E, MPC3E, MPC4E, and MPC 3D 16x10GE. Exceptions to the MPC2E and MPC3E series line card are MPC2E-NG and MPC3E-NG that use XL platform.

Required Privilege Level

`interface`—To view this statement in the configuration.

`interface-control`—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 16.1.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

[Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers | 633](#)

username (Services)

IN THIS SECTION

- [Syntax | 1531](#)
- [Hierarchy Level | 1531](#)
- [Description | 1531](#)
- [Options | 1531](#)
- [Required Privilege Level | 1531](#)
- [Release Information | 1531](#)

Syntax

```
username user-name;
```

Hierarchy Level

```
[edit services flow-collector transfer-log-archive archive-sites]
```

Description

Specify the username for the transfer log server.

Options

user-name—FTP server username.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | 226

variant

IN THIS SECTION

- [Syntax | 1532](#)
- [Hierarchy Level | 1532](#)
- [Description | 1532](#)
- [Required Privilege Level | 1532](#)
- [Release Information | 1533](#)

Syntax

```
variant variant-number {  
    data-format format;  
    name-format format;  
    transfer {  
        record-level number;  
        timeout seconds;  
    }  
}
```

Hierarchy Level

```
[edit services flow-collector file-specification]
```

Description

Configure a variant of the file format.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

| [Configuring Flow Collection](#) | [226](#)

version

IN THIS SECTION

- [Syntax](#) | [1533](#)
- [Hierarchy Level](#) | [1533](#)
- [Description](#) | [1534](#)
- [Options](#) | [1534](#)
- [Required Privilege Level](#) | [1534](#)
- [Release Information](#) | [1534](#)

Syntax

```
version format;
```

Hierarchy Level

```
[edit forwarding-options accounting name output flow-server hostname],
[edit forwarding-options sampling instance instance-name family (inet |inet6 |mpls) output flow-server
hostname],
[edit forwarding-options sampling family (inet |inet6 |mpls) output flow-server hostname]
```


Description

Specify the version format of the aggregated flows exported to a cflowd server.

Options

format—Format of the flows.

- **Values:** 5 or 8
- **Default:** 5

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[export-format](#) | 1082

[Enabling Flow Aggregation](#) | 577

version (Flow Monitoring Logs for NAT)

IN THIS SECTION

- [Syntax](#) | 1535
- [Hierarchy Level](#) | 1535
- [Description](#) | 1535
- [Options](#) | 1535
- [Required Privilege Level](#) | 1535

Syntax

```
version (ipfix | v9);
```

Hierarchy Level

```
[edit services jflow-log template-profile template-profile-name]
```

Description

Specify the flow template format, such as IPFIX or version 9, to be used for generating flow monitoring records for NAT events and for transmitting them to the collector.

Options

ipfix—Use the IPFIX flow template format for flow monitoring logs for NAT events.

v9—Use the version 9 flow template format for flow monitoring logs for NAT events.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

[Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250 | 241](#)

[Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250 | 256](#)

[Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats | 272](#)

[Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting | 275](#)

version9 (Forwarding Options)

IN THIS SECTION

- [Syntax | 1536](#)
- [Hierarchy Level | 1536](#)
- [Description | 1536](#)
- [Required Privilege Level | 1537](#)
- [Release Information | 1537](#)

Syntax

```
version9 {
    template template-name;
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (bridge | inet | inet6 | mpls vpls)
output flow-server hostname],
[edit forwarding-options sampling family (inet | inet6 | mpls | vpls | bridge) output flow-server
hostname]
```

Description

Specify flow monitoring version 9 properties to apply to output sampling records.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Support at the following hierarchy levels introduced in Junos OS Release 18.2R1: [edit forwarding-options sampling instance instance-name family bridge], [edit forwarding-options sampling instance instance-name family vpls], [edit forwarding-options sampling family bridge], and [edit forwarding-options sampling family vpls].

Statement introduced in Junos OS Release 19.2R1 for MX Series routers with MPC10E-15C-MRATE line card to define a flow record template suitable for IPv4 or IPv6 traffic only .

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

version9 (Flow Monitoring)

IN THIS SECTION

- [Syntax | 1538](#)
- [Hierarchy Level | 1538](#)
- [Description | 1538](#)
- [Required Privilege Level | 1538](#)
- [Release Information | 1539](#)

Syntax

```

version9 {
  template template-name {
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-key {
      flow-direction;
      vlan-id;
      output-interface;
    }
    (ipv4-template | ipv6-template | mpls-template label-position [ positions ] | mpls-ipv4-
template label-position [ positions ] | mpls-ipvx-template);
    option-refresh-rate packets packets seconds seconds;
    options-template-id
    peer-as-billing-template;
    source-id
    template-id
    template-refresh-rate packets packets seconds seconds;
    tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
  }
}

```

Hierarchy Level

```
[edit services flow-monitoring]
```

Description

Specify the version 9 output template properties to support flow monitoring.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 8.3.

Statement introduced in Junos OS Release 19.2R1 for MX Series routers with MPC10E-15C-MRATE line card to define a flow record template suitable for IPv4 or IPv6 traffic only.

RELATED DOCUMENTATION

[Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates | 583](#)

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

version-ipfix (Forwarding Options)

IN THIS SECTION

- [Syntax | 1539](#)
- [Hierarchy Level | 1540](#)
- [Description | 1540](#)
- [Required Privilege Level | 1540](#)
- [Release Information | 1540](#)

Syntax

```
version-ipfix {  
  template template-name;  
}
```

Hierarchy Level

```
[edit forwarding-options sampling instance instance-name family (inet | inet6 | mpls | vpls)
output flow-server hostname]
```

Description

Specify flow monitoring version IPFIX properties to apply to output sampling records.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#) | 74

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices](#) | 603

version-ipfix (Services)

IN THIS SECTION

- [Syntax](#) | 1541
- [Hierarchy Level](#) | 1542
- [Description](#) | 1542
- [Required Privilege Level](#) | 1542

Syntax

```

version-ipfix {
  template template-name {
    data-record-fields {
      source-prefix-as-path count;
      destination-prefix-as-path count;
      bgp-source-standard-community count;
      bgp-destination-standard-community count;
      bgp-source-extended-community count;
      bgp-destination-extended-community count;
      bgp-source-large-community count;
      bgp-destination-large-community count;
    }
    flow-active-timeout seconds;
    flow-inactive-timeout seconds;
    flow-key {
      flow-direction;
      vlan-id;
      output-interface;
    }
    (ipv4-template | ipv6-template | mpls-ipv4-template | mpls-ipvx-template | vpls-template);
    nexthop-learning (enable |disable);
    observation-domain-id
    option-refresh-rate packets packets seconds seconds;
    options-template-id
    template-id
    template-refresh-rate packets packets seconds seconds;
    tunnel-observation [ipv4 | ipv6 | mpls-over-udp];
  }
}

```


Hierarchy Level

```
[edit services flow-monitoring]
```

Description

Specify the IPFIX output template properties to support flow monitoring.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 10.2.

data-record-fields option introduced in Junos OS Evolved Release 21.4R1.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 | 74](#)

[Configuring Inline Active Flow Monitoring to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches, NFX Series Devices, and SRX Devices | 603](#)

[Configuring Inline Active Flow Monitoring on PTX Series Routers | 540](#)

video-monitoring

IN THIS SECTION

- [Syntax | 1543](#)
- [Hierarchy Level | 1545](#)

- [Description | 1545](#)
- [Required Privilege Level | 1545](#)
- [Release Information | 1546](#)

Syntax

```
video-monitoring {
  interfaces {
    interface-name {
      family {
        inet {
          input-flows {
            input-flow-name {
              destination-address [ address ];
              destination-port [ port ];
              source-address [ address ];
              source-port [ port ];
              template template-name;
            }
          }
          output-flows {
            output-flow-name {
              destination-address [ address ];
              destination-port [ port ];
              source-address [ address ];
              source-port [ port ];
              template template-name;
            }
          }
        }
      }
    }
    inet6 {
      input-flows {
        input-flow-name {
          destination-address [ address ];
          destination-port [ port ];
          source-address [ address ];
          source-port [ port ];
          template template-name;
        }
      }
    }
  }
}
```

```

    }
  }
  output-flows {
    output-flow-name {
      destination-address [ address ];
      destination-port [ port ];
      source-address [ address ];
      source-port [ port ];
      template template-name;
    }
  }
}
mpls {
  input-flows {
    input-flow-name {
      (destination-address [ address ] | source-address [ address ]);
      destination-port [ port ];
      payload-type (ipv4 | ipv6);
      source-port [ port ];
      template template-name;
    }
  }
  output-flows {
    output-flow-name {
      (destination-address [ address ] | source-address [ address ]);
      destination-port [ port ];
      payload-type (ipv4 | ipv6);
      source-port [ port ];
      template template-name;
    }
  }
}
}
}
}
templates {
  template-name {
    interval-duration interval-duration;
    inactive-timeout inactive-timeout;
    rate {
      (layer3 layer3-packets-per-second | media media-bits-per-second);
    }
    delay-factor {

```

```

        disable;
        threshold {
            (info | warning | critical) delay-factor-threshold;
        }
    }
    media-loss-rate {
        disable;
        threshold {
            (info | warning | critical) percentage mlr-percentage | packet-count mlr-
packet-count);
        }
    }
    media-rate-variation {
        ;
        threshold {
            (info | warning | critical) mrsv-variation;
        }
    }
    media-packets-count-in-layer3 media-packets-count-in-layer3;
    media-packet-size media-packet-size;
}
}
}

```

Hierarchy Level

[edit services]

Description

Define the options for video monitoring using media delivery index options for metrics.

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

| [Configuring Inline Video Monitoring on MX Series Routers](#) | 934

vpls-flow-table-size

IN THIS SECTION

- [Syntax](#) | 1546
- [Hierarchy Level](#) | 1546
- [Description](#) | 1546
- [Options](#) | 1547
- [Required Privilege Level](#) | 1547
- [Release Information](#) | 1547

Syntax

```
vpls-flow-table-size units;
```

Hierarchy Level

```
[edit chassis fpc slot-number inline-services flow-table-size]
```

Description

Configure the size of the VPLS flow table in units of 256K entries.

NOTE: Any change in the configured size of the flow table size initiates an automatic reboot of the FPC.

NOTE: Starting with Junos OS Release 17.3R1, the maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 256 on MPC5Es and MPC6Es with 4 GB DDR3 memory or higher. The maximum number of 256K flow entries that you can configure for IPv4 flow tables and IPv6 flow tables is 245 on MPC5Es and MPC6Es with DDR3 memory lower than 4 GB.

NOTE: The recommended flow table size is 4 so that it can scale up to 4x256K flows, which is 1M. You can configure more, however, the system will issue a warning message.

Options

units Number of 256K flow entries available for the VPLS flow table.

- **Range:** 1 through 245
- **Default:** 15 (3840K)

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13.2.

RELATED DOCUMENTATION

| [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#) | 74

vpls-template

IN THIS SECTION

- [Syntax | 1548](#)
- [Hierarchy Level | 1548](#)
- [Description | 1548](#)
- [Required Privilege Level | 1548](#)
- [Release Information | 1548](#)

Syntax

```
vpls-template;
```

Hierarchy Level

```
[edit services flow-monitoring version-ipfix template]
```

Description

Specify that the IPFIX template is used only for VPLS records. Starting in Junos OS Release 18.2R1, the `vpls-template` option is deprecated; use the `bridge-template` option instead. The `bridge-template` option supports both VPLS and bridge records.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Release Information

Statement introduced in Junos OS Release 13 .2.

RELATED DOCUMENTATION

[Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250](#) | 74

[bridge-template](#) | 990

world-readable

IN THIS SECTION

- [Syntax](#) | 1549
- [Hierarchy Level](#) | 1549
- [Description](#) | 1549
- [Options](#) | 1550
- [Required Privilege Level](#) | 1550
- [Release Information](#) | 1550

Syntax

```
(world-readable | no-world-readable);
```

Hierarchy Level

```
[edit forwarding-options port-mirroring traceoptions file],  
[edit forwarding-options sampling family (inet | inet6 | mpls) output file],  
[edit forwarding-options sampling traceoptionsfile]
```

Description

Enable unrestricted file access.

Options

no-world-readable—Restrict file access to owner. This is the default.

world-readable—Enable unrestricted file access.

- **Default:** no-world-readable

Required Privilege Level

interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Release Information

Statement introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[Configuring Port Mirroring on M, T MX, ACX, and PTX Series Routers](#)

[Configuring Traffic Sampling on MX, M and T Series Routers](#) | 420

Operational Commands

IN THIS CHAPTER

- clear passive-monitoring statistics | 1553
- clear services accounting statistics inline-jflow | 1555
- clear services dynamic-flow-capture | 1556
- clear services flow-collector statistics | 1558
- clear inband-flow-telemetry stats | 1560
- clear services inline-monitoring statistics | 1561
- clear services monitoring rfc2544 | 1562
- clear services rpm rfc2544-benchmarking | 1564
- clear services monitoring twamp server control-connection | 1565
- clear services rpm twamp server connection | 1566
- clear services service-sets statistics jflow-log | 1567
- clear services video-monitoring mdi errors fpc-slot | 1569
- clear services video-monitoring mdi statistics fpc-slot | 1571
- request services flow-collector change-destination primary interface | 1572
- request services flow-collector change-destination secondary interface | 1573
- request services flow-collector test-file-transfer | 1575
- request services monitoring twamp client | 1577
- request services rpm twamp | 1578
- show forwarding-options next-hop-group | 1580
- show forwarding-options port-mirroring | 1584
- show interfaces (Dynamic Flow Capture) | 1587
- show interfaces (Flow Collector) | 1594
- show interfaces (Flow Monitoring) | 1603
- show passive-monitoring error | 1611
- show passive-monitoring flow | 1614
- show passive-monitoring memory | 1618

- [show passive-monitoring status | 1620](#)
- [show passive-monitoring usage | 1622](#)
- [show route rpm-tracking | 1625](#)
- [show services accounting aggregation | 1630](#)
- [show services accounting aggregation template | 1636](#)
- [show services accounting errors | 1638](#)
- [show services accounting flow | 1645](#)
- [show services accounting flow-detail | 1654](#)
- [show services accounting memory | 1661](#)
- [show services accounting packet-size-distribution | 1664](#)
- [show services accounting status | 1666](#)
- [show services accounting usage | 1671](#)
- [show services dynamic-flow-capture content-destination | 1674](#)
- [show services dynamic-flow-capture control-source | 1676](#)
- [show services dynamic-flow-capture statistics | 1680](#)
- [show services flow-collector file interface | 1684](#)
- [show services flow-collector input interface | 1688](#)
- [show services flow-collector interface | 1690](#)
- [show services inband-flow-telemetry | 1701](#)
- [show services inline-monitoring feature-profile-mapping fpc-slot | 1704](#)
- [show services inline-monitoring statistics fpc-slot | 1707](#)
- [show services monitoring rfc2544 | 1710](#)
- [show services monitoring rfc2544 | 1715](#)
- [show services monitoring rpm history-results | 1719](#)
- [show services monitoring rpm probe-results | 1727](#)
- [show services monitoring twamp client control-info | 1740](#)
- [show services monitoring twamp client history-results | 1743](#)
- [show services monitoring twamp client probe-results | 1749](#)
- [show services monitoring twamp client test-info | 1759](#)
- [show services monitoring twamp server control-info | 1762](#)
- [show services monitoring twamp server test-info | 1764](#)
- [show services rpm active-servers | 1768](#)

- [show services rpm history-results | 1769](#)
- [show services rpm probe-results | 1775](#)
- [show services rpm rfc2544-benchmarking | 1791](#)
- [show services rpm rfc2544-benchmarking test-id | 1800](#)
- [show services rpm twamp client connection | 1826](#)
- [show services rpm twamp client history-results | 1828](#)
- [show services rpm twamp client probe-results | 1833](#)
- [show services rpm twamp client session | 1844](#)
- [show services rpm twamp server connection | 1847](#)
- [show services rpm twamp server session | 1849](#)
- [show services service-sets statistics jflow-log | 1852](#)
- [show services video-monitoring mdi errors fpc-slot | 1862](#)
- [show services video-monitoring mdi flows fpc-slot | 1865](#)
- [show services video-monitoring mdi stats fpc-slot | 1872](#)
- [test services monitoring rfc2544 | 1876](#)
- [test services rpm rfc2544-benchmarking test | 1880](#)

clear passive-monitoring statistics

IN THIS SECTION

- [Syntax | 1554](#)
- [Description | 1554](#)
- [Options | 1554](#)
- [Required Privilege Level | 1554](#)
- [Output Fields | 1554](#)
- [Sample Output | 1554](#)
- [Release Information | 1554](#)

Syntax

```
clear passive-monitoring statistics (all | interface interface-name)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Clear statistics for one passive monitoring interface or for all passive monitoring interfaces.

Options

all	Clear statistics for all configured passive monitoring interfaces.
interface <i>interface-name</i>	Clear statistics for the specified passive monitoring interface (mo-fpc/pic/port).

Required Privilege Level

network

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear passive-monitoring statistics

```
user@host> clear passive-monitoring statistics interface mo-5/0/0
```

Release Information

Command introduced in Junos OS Release 7.6.

clear services accounting statistics inline-jflow

IN THIS SECTION

- [Syntax | 1555](#)
- [Description | 1555](#)
- [Options | 1555](#)
- [Required Privilege Level | 1556](#)
- [Output Fields | 1556](#)
- [Sample Output | 1556](#)
- [Release Information | 1556](#)

Syntax

```
clear services accounting statistics inline-jflow  
<inline-jflow (fpc-slot slot-number)>
```

Description

Clear inline flow statistics for a specified FPC.

Options

fpc-slot *slot-number* Clear inline flow statistics for the specified FPC.

- MX80 Series routers only—Replace *slot-number* with a value from 0 through 1.
- MX104 Series routers only—Replace *slot-number* with a value from 0 through 2.
- MX240 Series routers only—Replace *slot-number* with a value from 0 through 2.
- MX480 Series routers only—Replace *slot-number* with a value from 0 through 5.
- MX960 Series routers only—Replace *slot-number* with a value from 0 through 11.

- MX2010 Series routers only—Replace *slot-number* with a value from 0 through 9.
- MX2020 Series routers only—Replace *slot-number* with a value from 0 through 19.

Required Privilege Level

view

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services accounting statistics inline-jflow

```
user@host> run clear services accounting statistics inline-jflow fpc-slot 5
Statistics Cleared
```

Release Information

Command introduced in Junos OS Release 14.2.

RELATED DOCUMENTATION

| [show services accounting flow](#) | [1645](#)

clear services dynamic-flow-capture

IN THIS SECTION

- [Syntax](#) | [1557](#)
- [Description](#) | [1557](#)

- [Options | 1557](#)
- [Required Privilege Level | 1557](#)
- [Output Fields | 1558](#)
- [Sample Output | 1558](#)
- [Release Information | 1558](#)

Syntax

```
clear services dynamic-flow-capture capture-group group-name
  <criteria-identifier identifier>
  <destination-identifier identifier>
  <force>
  <static>
```

Description

(M320 Series routers and T Series routers only) Clear dynamic flow capture information for specified capture group.

Options

capture-group <i>group-name</i>	Use the specified capture-group identifier.
criteria-identifier <i>identifier</i>	(Optional) Use the specified criteria identifier.
destination-identifier <i>identifier</i>	(Optional) Use the specified content destination identifier.
force	(Optional) Force clearing of criteria.
static	(Optional) Clear static criteria.

Required Privilege Level

network

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

clear services dynamic-flow-capture

```
user@host> clear services dynamic-flow-capture capture-group flow-a
```

Release Information

Command introduced in Junos OS Release 7.4.

clear services flow-collector statistics

IN THIS SECTION

- [Syntax | 1558](#)
- [Description | 1559](#)
- [Options | 1559](#)
- [Required Privilege Level | 1559](#)
- [Output Fields | 1559](#)
- [Sample Output | 1559](#)
- [Release Information | 1559](#)

Syntax

```
clear services flow-collector statistics (all | interface interface-name)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Clear statistics for one flow collector interface or for all flow collector interfaces.

Options

all Clear statistics for all configured flow collector interfaces.

interface *interface-name* Clear statistics for the specified flow collector interface (*cp-fpc/pic/port*).

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services flow-collector statistics

```
user@host> clear services flow-collector statistics interface cp-5/0/0
Flow collector interface: cp-5/0/0
Interface state: Collecting flows
Statistics cleared successfully
```

Release Information

Command introduced before Junos OS Release 7.4.

clear inband-flow-telemetry stats

IN THIS SECTION

- [Syntax | 1560](#)
- [Description | 1560](#)
- [Options | 1560](#)
- [Required Privilege Level | 1560](#)
- [Output Fields | 1560](#)
- [Sample Output | 1561](#)
- [Release Information | 1561](#)

Syntax

```
clear inband-flow-telemetry stats
```

Description

Clear Inband Flow Analyzer 2.0 (IFA 2.0) statistics.

You retrieve IFA statistics directly from the Packet Forwarding Engine. The Routing Engine does not maintain these statistics. Therefore, restarting the Packet Forwarding Engine process clears the statistics, whereas a Routing Engine process restart does not have any impact on the statistics.

Options

None

Required Privilege Level

network

Output Fields

This command produces no output.

Sample Output

clear inband-flow-telemetry stats (QFX5120-48Y and QFX5120-32C)

```
user@host> clear inband-flow-telemetry stats
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

[show services inband-flow-telemetry | 1701](#)

[inband-flow-telemetry | 1157](#)

[Inband Flow Analyzer \(IFA\) 2.0 Probe for Real-Time Flow Monitoring | 370](#)

clear services inline-monitoring statistics

IN THIS SECTION

- [Syntax | 1561](#)
- [Description | 1562](#)
- [Options | 1562](#)
- [Required Privilege Level | 1562](#)
- [Release Information | 1562](#)

Syntax

```
clear services inline-monitoring statistics fpc-slot fpc-slot  
collector-name collector-name  
instance-name instance-name
```

Description

Clear statistics for inline monitoring services.

Options

collector-name <i>collector-name</i>	Clear collector level statistics.
fpc-slot <i>fpc-slot</i>	Clear statistics for the specified FPC slot. <ul style="list-style-type: none"> • Range: 0 through 11
instance-name <i>instance-name</i>	Clear instance level statistics.

Required Privilege Level

view

Release Information

Command introduced in Junos OS Release 19.4R1.

Command introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

[show services inline-monitoring statistics fpc-slot | 1707](#)
[Understanding Inline Monitoring Services | 334](#)

clear services monitoring rfc2544

IN THIS SECTION

- [Syntax | 1563](#)
- [Description | 1563](#)
- [Options | 1563](#)

- [Required Privilege Level | 1563](#)
- [Output Fields | 1563](#)
- [Release Information | 1564](#)

Syntax

```
clear services monitoring rfc2544  
<active-tests | all-tests | completed-tests | terminated-tests | test-id test-id>
```

Description

Clear the indicated type of RFC 2544 benchmarking test sessions.

Options

active-tests	Clears the active RFC 2544 benchmarking test sessions.
all-tests	Clears all RFC 2544 benchmarking test sessions.
completed-tests	Clears the completed RFC 2544 benchmarking test sessions.
terminated-tests	Clears the terminated RFC 2544 benchmarking test sessions.
test-id <i>test-id</i>	Clears the particular RFC 2544 benchmarking test by specifying its test ID.

Required Privilege Level

clear

Output Fields

When you enter this command, if it is successful, there is no response. If it is not successful, the system displays an error message if the FPC for the interface specified on the `test-interface` statement reports an error.

Release Information

Command introduced in Junos OS Evolved Release 21.1R1.

all-tests and test-id options introduced in Junos OS Evolved 22.4R1.

clear services rpm rfc2544-benchmarking

IN THIS SECTION

- [Syntax | 1564](#)
- [Description | 1564](#)
- [Options | 1564](#)
- [Additional Information | 1565](#)
- [Required Privilege Level | 1565](#)
- [Output Fields | 1565](#)
- [Release Information | 1565](#)

Syntax

```
clear services rpm rfc2544-benchmarking  
<aborted-tests | active-tests | completed-tests>
```

Description

Clear the indicated type of RFC 2544 benchmarking test sessions.

Options

- | | |
|----------------------|--|
| aborted-tests | Clears the terminated RFC 2544 benchmarking test sessions. |
| active-tests | Clears the active RFC 2544 benchmarking test sessions. |

completed-tests Clears the completed RFC 2544 benchmarking test sessions.

Additional Information

Required Privilege Level

clear

Output Fields

When you enter this command, if it is successful, there is no response. If it is not successful, the system displays an error message if the FPC for the interface specified on the `test-interface` statement reports an error.

Release Information

Command introduced in Junos OS Release 12.3X52.

clear services monitoring twamp server control-connection

IN THIS SECTION

- [Syntax | 1565](#)
- [Description | 1566](#)
- [Options | 1566](#)
- [Required Privilege Level | 1566](#)
- [Release Information | 1566](#)

Syntax

```
clear services monitoring twamp server control-connection  
<connection-name>
```


Description

Clear connections established between the Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established connections are cleared, along with the sessions on those connections. To clear only a particular connection, specify the connection name when you issue the command.

Options

connection-name (Optional) Specify which connection to clear.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

[Understand Two-Way Active Measurement Protocol](#) | [686](#)

clear services rpm twamp server connection

IN THIS SECTION

- [Syntax](#) | [1567](#)
- [Description](#) | [1567](#)
- [Options](#) | [1567](#)
- [Required Privilege Level](#) | [1567](#)
- [Release Information](#) | [1567](#)

Syntax

```
clear services rpm twamp server connection  
<connection-id>
```

Description

Clear connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default all established connections are cleared (along with the sessions on those connections). To clear only a specific connection, specify the connection ID when you issue the command.

Options

connection-id (Optional) Specific connection to clear.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 9.3.

clear services service-sets statistics jflow-log

IN THIS SECTION

- [Syntax | 1568](#)
- [Description | 1568](#)
- [Options | 1568](#)
- [Required Privilege Level | 1568](#)
- [Output Fields | 1568](#)
- [Sample Output | 1569](#)

Syntax

```
clear services service-sets statistics jflow-log
<service-set service-set-name>
<interface interface-name>
```

Description

Clear flow monitoring log statistics for the logs generated in IPFIX or version 9 format for one services interface or for all services interfaces, and for one named service set or all service sets on the interface or interfaces.

Options

none	Clear flow monitoring log for all configured services interfaces and their service sets.
interface <i>interface-name</i>	(Optional) Clear flow monitoring log statistics for the specified services interface. On M Series, MX Series, and T Series routers, the <i>interface-name</i> can be <i>ms-fpc/pic/port</i> . It is supported only on MS-MICs and MS-MPCS.
service-set <i>service-set-name</i>	(Optional) Clear flow monitoring log statistics for the specified services interface.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services service-sets statistics jflow-log interface

```
user@host> clear services service-sets statistics jflow-log interface ms-5/0/0
Interface: ms-5/0/0
```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

| *show services service-sets statistics syslog*

clear services video-monitoring mdi errors fpc-slot

IN THIS SECTION

- [Syntax | 1569](#)
- [Description | 1570](#)
- [Options | 1570](#)
- [Required Privilege Level | 1570](#)
- [Output Fields | 1570](#)
- [Sample Output | 1570](#)
- [Release Information | 1570](#)

Syntax

```
clear services video-monitoring mdi errors <fpc-slot fpc-slot>
```

Description

Clear all media delivery index error counters for the specified FPC slot or for all FPC slots.

Options

none Clear error counters for all FPC slots.

fpc-slot (Optional) Clear error counters for the specified FPC slot.

Required Privilege Level

clear

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear services video-monitoring mdi errors

```
user@host> clear services video-monitoring mdi errors
Errors counters cleared
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

| [show services video-monitoring mdi stats fpc-slot](#) | [1872](#)

clear services video-monitoring mdi statistics fpc-slot

IN THIS SECTION

- [Syntax | 1571](#)
- [Description | 1571](#)
- [Options | 1571](#)
- [Required Privilege Level | 1571](#)
- [Release Information | 1571](#)

Syntax

```
clear services video-monitoring mdi statistics fpc-slot fpc-slot
```

Description

Clear all media delivery index statistics counters except for active flows.

Options

fpc-slot Number of the FPC slot.

Required Privilege Level

clear

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

| [show services video-monitoring mdi stats fpc-slot | 1872](#)

request services flow-collector change-destination primary interface

IN THIS SECTION

- [Syntax | 1572](#)
- [Description | 1572](#)
- [Options | 1572](#)
- [Required Privilege Level | 1573](#)
- [Output Fields | 1573](#)
- [Sample Output | 1573](#)
- [Release Information | 1573](#)

Syntax

```
request services flow-collector change-destination primary interface cp-fpc/pic/port
<clear-files>
<clear-logs>
<immediately | gracefully>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Switch to the primary File Transfer Protocol (FTP) server that is configured as a flow collector.

Options

none	Switch to the primary FTP server.
<i>cp-fpc/pic/port</i>	Use the specified flow collector interface name for the primary destination.
clear-files	(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.
clear-logs	(Optional) Request clearing of existing logs when the switch takes place.

**immediately |
gracefully**

(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector change-destination primary interface

```
user@host> request services flow-collector change-destination primary interface
cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

Release Information

Command introduced before Junos OS Release 7.4.

request services flow-collector change-destination secondary interface

IN THIS SECTION

- [Syntax | 1574](#)
- [Description | 1574](#)
- [Options | 1574](#)
- [Required Privilege Level | 1574](#)
- [Output Fields | 1575](#)

- [Sample Output | 1575](#)
- [Release Information | 1575](#)

Syntax

```
request services flow-collector change-destination secondary interface cp-fpc/pic/  
port  
<clear-files>  
<clear-logs>  
<immediately | gracefully>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Switch to the secondary File Transfer Protocol (FTP) server that is configured as a flow collector.

Options

none	Switch to the secondary FTP server.
<i>cp-fpc/pic/port</i>	Use the specified flow collector interface name (<i>cp-fpc/pic/port</i>) for the secondary destination.
clear-files	(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.
clear-logs	(Optional) Request clearing of existing logs when the switch takes place.
immediately gracefully	(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.

Required Privilege Level

maintenance

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector change-destination secondary interface

```
user@host> request services flow-collector change-destination secondary interface
cp-6/0/0
Flow collector interface: cp-6/0/0
Interface state: Collecting flows
Destination change successful
```

Release Information

Command introduced before Junos OS Release 7.4.

request services flow-collector test-file-transfer

IN THIS SECTION

- [Syntax | 1576](#)
- [Description | 1576](#)
- [Options | 1576](#)
- [Required Privilege Level | 1576](#)
- [Output Fields | 1576](#)
- [Sample Output | 1576](#)
- [Release Information | 1577](#)

Syntax

```
request services flow-collector test-file-transfer filename interface (all | cp-fpc/pic/port)
(channel-zero | channel-one) (primary | secondary)
```

Description

(M40e, M160, and M320 Series routers, PTX Series, and T Series routers only) Transfer a test file to the primary or secondary File Transfer Protocol (FTP) server that is configured as a flow collector. This command verifies that the output side of the flow collector interface is operating properly.

Options

<i>filename</i>	Name of the test file to transfer.
interface (all cp-fpc/pic/port)	Transfer a test file of flows from all configured flow collector interfaces or from only the specified interface.
channel-zero channel-one	Transfer a file from export channel 0 (unit 0) or channel 1 (unit 1) of the PIC.
primary secondary	Transfer a file to the primary or secondary server configured as a flow collector.

Required Privilege Level

network

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request services flow-collector test-file-transfer interface channel-one primary

```
user@host> request services flow-collector test-file-transfer test_file interface cp-7/1/0
channel-one primary
```

```
Flow collector interface: cp-7/1/0
Interface state: Collecting flows
Response: Test file transfer successfully scheduled
```

Release Information

Command introduced before Junos OS Release 7.4.

request services monitoring twamp client

IN THIS SECTION

- [Syntax | 1577](#)
- [Description | 1577](#)
- [Options | 1578](#)
- [Required Privilege Level | 1578](#)
- [Release Information | 1578](#)

Syntax

```
request services monitoring twamp client (start | stop)
<control-connection-name>
repeat <number>
```

Description

Start or stop a Two-Way Active Measurement Protocol (TWAMP) session. You can start or stop all of the sessions for all of the TWAMP clients, or start or stop a session for a specific TWAMP client control connection. When you start all the test sessions configured for a particular TWAMP client, the control-client starts all requested testing with a Start-Sessions message, and the server sends an acknowledgment. If the control connection is not active between the server and the client, the control

connection is also established and the test connections are started later. If the control-client name is not specified, all the configured test sessions are begun.

When you stop the test session, the control connection is closed only after the Stop-sessions message is sent from the TWAMP client to the TWAMP server. If the control-client name is not specified, all the configured test sessions are closed.

Options

control-connection-name (Optional) Start or stop the TWAMP session with the server for only the specified control-connection or TWAMP control-client. If you do not specify this option, all control connections are either started or stopped.

repeat number

Required Privilege Level

view

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

[control-connection \(Junos OS Evolved\) | 1019](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

request services rpm twamp

IN THIS SECTION

- [Syntax | 1579](#)
- [Description | 1579](#)
- [Options | 1579](#)
- [Required Privilege Level | 1579](#)

- [Output Fields | 1580](#)
- [Sample Output | 1580](#)
- [Release Information | 1580](#)

Syntax

```
request services rpm twamp (start | stop) client <control-connection-name>
```

Description

Start or stop a TWAMP session. You can start or stop all of the sessions for all of the TWAMP clients, or start or stop a session for a specific TWAMP client. When you start all the test session configured for a particular TWAMP client, the control-client initiates all requested testing with a Start-Sessions message, and the server sends an acknowledgment. If the control connection is not active between the server and the client, the control connection is also established and the test connections are started later. If the control-client name is not specified, all the configured test sessions are commenced.

When you stop the test session, the control connection is closed only after the Stop-sessions message is sent from the TWAMP client to the TWAMP server. If the control-client name is not specified, all the configured test sessions are closed.

Options

start client	Start the TWAMP session between the TWAMP client and the TWAMP server.
stop client	Terminate the TWAMP session between the TWAMP client and the TWAMP server.
control-connection-name	(Optional) Start or stop the TWAMP session with the server only for the specified control-connection or TWAMP control-client.

Required Privilege Level

view

Output Fields

When you enter this command, you are not provided feedback on the status of your request.

Sample Output

request services rpm twamp start client c1

```
user@host> request services rpm twamp start client c1
```

Release Information

Command introduced in Junos OS Release 15.1.

show forwarding-options next-hop-group

IN THIS SECTION

- [Syntax | 1580](#)
- [Description | 1581](#)
- [Options | 1581](#)
- [Required Privilege Level | 1581](#)
- [Output Fields | 1581](#)
- [Sample Output | 1582](#)
- [Release Information | 1584](#)

Syntax

```
show forwarding-options next-hop-group  
<terse | brief | detail>  
<group-name>
```

Description

Display current state of next-hop groups.

Options

terse | brief | detail (Optional) Display the specified level of output.

group-name (Optional) Display a single next-hop group.

Required Privilege Level

view

Output Fields

[Table 139 on page 1581](#) lists the output fields for the `show forwarding-options next-hop-group` command. Output fields are listed in the approximate order in which they appear.

Table 139: show forwarding-options next-hop-group Output Fields

Field Name	Field Description	Level of Output
Next-hop-group	Name of next-hop group.	All levels
Type	Next-hop group type, such as inet , inet6 or layer-2 .	All levels
State	Next-hop group state, either up or down .	All levels
Members Interfaces	Names of interfaces to which next-hop group members belong.	brief detail
Member Subgroup	Names of subgroups to which next-hop group members belong.	brief detail
Number of members configured	Number of next-hop group members configured.	detail

Table 139: show forwarding-options next-hop-group Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Number of members that are up	Number of next-hop group members that are up.	detail
Number of subgroups configured	Number of subgroups configured.	detail
Number of subgroups that are up	Number of subgroups that are up.	detail

Sample Output

show forwarding-options next-hop-group terse

```

user@host> show forwarding-options next-hop-group terse
Next-hop-group      Type      State
nhg                  inet      up
nhg6                 inet6     up
vpls_nhg_2          layer-2   down

```

show forwarding-options next-hop-group brief

```

user@host> show forwarding-options next-hop-group brief

Next-hop-group: nhg
  Type: inet
  State: up
  Members Interfaces:
    ge-0/2/8.0      next-hop 192.0.2.10
    ge-5/1/8.0      next-hop 198.51.100.10
    ge-5/1/9.0      next-hop 203.0.113.10

```

```

Next-hop-group: nhg6
  Type: inet6
  State: up
  Members Interfaces:
    ge-5/1/5.0          next-hop 2001:db8::1:10
    ge-5/1/6.0          next-hop 2001:db8::20:10      Member Subgroup: nhsg6
  Members Interfaces:
    ge-5/0/4.0          next-hop 2001:db8::3:1
    ge-5/1/4.0          next-hop 2001:db8::4:1

Next-hop-group: vpls_nhg_2
  Type: layer-2      State: down

```

show forwarding-options next-hop-group detail

```
user@host> show forwarding-options next-hop-group detail
```

```

Next-hop-group: nhg
Type: inet
State: up
Number of members configured      : 3
Number of members that are up    : 3
Number of subgroups configured   : 0
Number of subgroups that are up  : 0
Members Interfaces:
  ge-0/2/8.0      next-hop 192.0.2.10      up
  ge-5/1/8.0      next-hop 203.0.113.10     up
  ge-5/1/9.0      next-hop 198.51.100.10.10 up

Next-hop-group: nhg6
Type: inet6
State: up
Number of members configured      : 2
Number of members that are up    : 2
Number of subgroups configured   : 1
Number of subgroups that are up  : 1
Members Interfaces:
  ge-5/1/5.0      next-hop 2001:db8::1:10    up
  ge-5/1/6.0      next-hop 2001:db8::20:10   up
Member Subgroup: nhsg6
Number of members configured      : 2

```

```

Number of members that are up    : 2
Members Interfaces:
    ge-5/0/4.0      next-hop 2001:db8::3:1  up
    ge-5/1/4.0      next-hop 2001:db8::4:1  up

Next-hop-group: vpls_nhg_2
Number of members configured    : 2
Number of members that are up   : 0
Number of subgroups configured  : 0
Number of subgroups that are up : 0
Type: layer-2      State: down
Members Interfaces:      State
    ge-2/2/1.100        down
    ge-2/3/9.0          down

```

Release Information

Command introduced in Junos OS Release 9.6.

Support for IPv6 introduced in Junos OS Release 14.2 for the MX Series routers.

RELATED DOCUMENTATION

[show forwarding-options port-mirroring](#) | 1584

show forwarding-options port-mirroring

IN THIS SECTION

- [Syntax](#) | 1585
- [Description](#) | 1585
- [Options](#) | 1585
- [Required Privilege Level](#) | 1585
- [Output Fields](#) | 1585
- [Sample Output](#) | 1586

Syntax

```
show forwarding-options port-mirroring
<terse | detail>
<instance-name>
```

Description

Display current state of port-mirroring instances.

Options

- terse | detail** (Optional) Display the specified level of output.
- instance-name*** (Optional) Display a single port-mirroring instance.

Required Privilege Level

view

Output Fields

Table 140 on page 1585 lists the output fields for the show forwarding-options port-mirroring command. Output fields are listed in the approximate order in which they appear.

Table 140: show forwarding-options port-mirroring Output Fields

Field Name	Field Description	Level of Output
Instance Name	Name of port-mirroring instance.	All levels
Instance Id	Instance identification number.	All levels

Table 140: show forwarding-options port-mirroring Output Fields (Continued)

Field Name	Field Description	Level of Output
State	Instance state, either up or down.	All levels
Input parameters		
Rate	Rate (ratio of packets sampled).	detail
Run-length	Run length (number of consecutive packets sampled).	detail
Maximum-packet-length	Maximum packet length.	detail
Output parameters		
Family	Protocol family.	detail
State	Instance state, either up or down.	detail
Destination	Destination (next-hop group name).	detail
Next-hop	IP address of the next hop to the destination.	detail

Sample Output

show forwarding-options port-mirroring terse

```

user@host> show forwarding-options port-mirroring terse
Instance Name      Instance Id  State
&global_instance  1           up
inst1              2           up

```

show forwarding-options port-mirroring detail

```

user@host> show forwarding-options port-mirroring detail
Instance Name: pm1
Instance Id: 2
Input parameters:
  Rate           : 2
  Run-length     : 0
  Maximum-packet-length : 0
Output parameters:
  Family      State   Destination      Next-hop
  inet        up      ge-0/0/0.0      10.1.1.2
  inet6       up      ge-0/0/0.0      2001:db8::2
  any         up      ge-0/0/1.0      NA

```

Release Information

Command introduced in Junos OS Release 9.6.

show interfaces (Dynamic Flow Capture)

IN THIS SECTION

- [Syntax | 1588](#)
- [Description | 1588](#)
- [Options | 1588](#)
- [Required Privilege Level | 1588](#)
- [Output Fields | 1588](#)
- [Sample Output | 1593](#)
- [Release Information | 1594](#)

Syntax

```
show interfaces dfc-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M320 and M120 Series routers and T Series routers only) Display status information about the specified dynamic flow capture interface.

Options

<i>dfc-fpc/pic/port:channel</i>	Display standard status information about the specified dynamic flow capture interface.
brief detail extensive terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index <i>snmp-index</i>	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

[Table 141 on page 1589](#) lists the output fields for the `show interfaces` (Dynamic Flow Capture) command. Output fields are listed in the approximate order in which they appear.

Table 141: Dynamic Flow Capture show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	Sate of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link type	Data transmission type.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .	All levels

Table 141: Dynamic Flow Capture show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour.minute.second timezone (hour.minute.second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input rate, Output rate—Number of bits per second (packets per second) received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Input errors	<ul style="list-style-type: none"> • Errors—Input errors on the interface. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed Discards—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that the Junos OS does not support. • Resource errors—Sum of transmit drops. 	extensive

Table 141: Dynamic Flow Capture show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Output errors	<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 then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of outgoing frame terminates 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 dropped by the ASIC RED mechanism. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Flags	Information about the logical interface; values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Input packets	Number of packets received on the logical interface.	None specified
Output packets	Number of packets transmitted on the logical interface.	None specified

Table 141: Dynamic Flow Capture show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Protocol	Protocol family configured on the logical interface (such as iso or inet6).	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Addresses associated with the logical interface and information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none

Sample Output

show interfaces (Dynamic Flow Capture)

```

user@host> show interfaces dfc-0/0/0
Physical interface: dfc-0/0/0, Enabled, Physical link is Up
  Interface index: 146, SNMP ifIndex: 36
  Type: Adaptive-Services, Link-level type: Dynamic-Flow-Capture, MTU: 9192, Speed: 2488320kbps
  Device flags : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type : Full-Duplex
  Link flags : None
  Last flapped : 2005-08-26 15:08:36 PDT (01:18:42 ago)
  Input rate : 0 bps (0 pps)
  Output rate : 44800440 bps (100000 pps)

Logical interface dfc-0/0/0.0 (Index 67) (SNMP ifIndex 43)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 74
  Output packets: 132
  Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.36.100.1, Local: 10.36.100.2

Logical interface dfc-0/0/0.1 (Index 68) (SNMP ifIndex 49)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 0
  Output packets: 402927263
  Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded

Logical interface dfc-0/0/0.2 (Index 69) (SNMP ifIndex 50)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 0
  Output packets: 0
  Protocol inet, MTU: 9192
  Flags: Receive-options, Receive-TTL-Exceeded

Logical interface dfc-0/0/0.16383 (Index 70) (SNMP ifIndex 44)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Dynamic-Flow-Capture
  Input packets : 1427

```

```

Output packets: 98
Protocol inet, MTU: 9192
Flags: Receive-options, Receive-TTL-Exceeded
Addresses, Flags: Is-Preferred Is-Primary
Destination: 10.0.0.16, Local: 10.0.0.1

```

Release Information

Command introduced in Junos OS Release 7.4.

show interfaces (Flow Collector)

IN THIS SECTION

- [Syntax | 1594](#)
- [Description | 1595](#)
- [Options | 1595](#)
- [Required Privilege Level | 1595](#)
- [Output Fields | 1595](#)
- [Sample Output | 1601](#)
- [Release Information | 1603](#)

Syntax

```

show interfaces cp-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>

```

Description

(M Series and T Series routers only) Display status information about the specified flow collector interface.

Options

<i>cp-fpc/pic/port:channel</i>	Display standard status information about the specified flow collector interface.
brief detail extensive terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index <i>snmp-index</i>	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

[Table 142 on page 1595](#) lists the output fields for the `show interfaces (Flow Collector)` command. Output fields are listed in the approximate order in which they appear.

Table 142: Flow Collector Show interfaces Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical Interface	Name of the physical interface type.	All levels
Link	Status of the link: up or down .	All levels

Table 142: Flow Collector Show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Enabled	State of the interface type. Possible values are described in the “Enabled Devices” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link type	Data transmission type.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .	All levels

Table 142: Flow Collector Show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Physical info	Information about the physical interface.	All levels
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	detail extensive none
Current address	Configured MAC address.	detail extensive none
Hardware address	Media access control (MAC) address of the interface.	detail extensive none
Alternate link address	Backup link address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour.minute.second timezone (hour.minute.second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .	detail extensive
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive

Table 142: Flow Collector Show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Input errors	<ul style="list-style-type: none"> • Errors—Input errors on the interface. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed Discards—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that Junos does not support. • Resource errors—Sum of transmit drops. 	extensive
Output errors	<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 then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of outgoing frame terminates 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 dropped by the ASIC RED mechanism. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
Logical interface	Name of the logical interface	All levels

Table 142: Flow Collector Show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface; values are described in the "Logical Interface Flags" section under <i>Common Output Fields Description</i> .	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Traffic statistics	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Local statistics	Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	detail extensive
Transit statistics	Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.	detail extensive

Table 142: Flow Collector Show interfaces Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Protocol	Protocol family configured on the logical interface (such as iso or inet6).	detail extensive none
MTU	MTU size on the logical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Route table	Route table in which this address exists; for example, Route table:0 refers to inet.0.	detail extensive
Flags	Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Addresses, Flags	Information about the address flags. Possible values are described in the “Addresses Flags” section under <i>Common Output Fields Description</i> .	detail extensive none
Destination	IP address of the remote side of the connection.	detail extensive none
Local	IP address of the logical interface.	detail extensive none
Broadcast	Broadcast address.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive

Sample Output

show interfaces extensive (Flow Collector)

```

user@host> show interfaces extensive cp-5/0/0
Physical interface: cp-5/0/0, Enabled, Physical link is Up
  Interface index: 145, SNMP ifIndex: 52, Generation: 29
  Type: Flow-collector, Link-level type: Flow-collection, MTU: 9192,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-05-24 16:48:11 PDT (00:12:04 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes :          2041661287          0 bps
    Output bytes :          3795049544      43816664 bps
    Input  packets:          1365534          0 pps
    Output packets:          3865644      3670 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface cp-5/0/0.0 (Index 74) (SNMP ifIndex 53) (Generation 28)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection
  Traffic statistics:
    Input  bytes :          1064651568
    Output bytes :          37144290
    Input  packets:          711324
    Output packets:          713672
  Local statistics:
    Input  bytes :          0
    Output bytes :          0
    Input  packets:          0

```

```

Output packets:          0
Transit statistics:
Input  bytes  :          1064651568          0 bps
Output bytes  :          37144290          0 bps
Input  packets:          711324          0 pps
Output packets:          713672          0 pps
Protocol inet, MTU: 9192, Generation: 39, Route table: 0
  Flags: Receive-options, Receive-TTL-Exceeded
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 192.0.2.2, Local: 192.0.2.1, Broadcast: Unspecified,
    Generation: 40

```

Logical interface cp-5/0/0.1 (Index 75) (SNMP ifIndex 54) (Generation 29)

Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection

Traffic statistics:

```

Input  bytes  :          976793823
Output bytes  :          34099481
Input  packets:          652729
Output packets:          655127

```

Local statistics:

```

Input  bytes  :          0
Output bytes  :          0
Input  packets:          0
Output packets:          0

```

Transit statistics:

```

Input  bytes  :          976793823          0 bps
Output bytes  :          34099481          0 bps
Input  packets:          652729          0 pps
Output packets:          655127          0 pps

```

Protocol inet, MTU: 9192, Generation: 40, Route table: 0

Flags: Receive-options, Receive-TTL-Exceeded

Addresses, Flags: Is-Preferred Is-Primary

Destination: 198.51.100.2, Local: 198.51.100.1, Broadcast: Unspecified,
Generation: 42

Logical interface cp-5/0/0.2 (Index 80) (SNMP ifIndex 55) (Generation 30)

Flags: Point-To-Point SNMP-Traps Encapsulation: Flow-collection

Traffic statistics:

```

Input  bytes  :          0
Output bytes  :          3723079376
Input  packets:          0
Output packets:          2495372

```

Local statistics:

```

Input bytes :          0
Output bytes :          0
Input packets:          0
Output packets:         0
Transit statistics:
Input bytes :          0          0 bps
Output bytes :    3723079376    43816664 bps
Input packets:          0          0 pps
Output packets:    2495372    3670 pps
Protocol inet, MTU: 9192, Generation: 41, Route table: 0
  Flags: Receive-options, Receive-TTL-Exceeded
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 203.0.113.2, Local: 203.0.113.1, Broadcast: Unspecified,
    Generation: 44

Logical interface cp-5/0/0.16383 (Index 81) (SNMP ifIndex 56) (Generation 31)
...
```

Release Information

Command introduced before Junos OS Release 7.4.

show interfaces (Flow Monitoring)

IN THIS SECTION

- [Syntax | 1604](#)
- [Description | 1604](#)
- [Options | 1604](#)
- [Required Privilege Level | 1604](#)
- [Output Fields | 1604](#)
- [Sample Output | 1610](#)
- [Release Information | 1611](#)

Syntax

```
show interfaces mo-fpc/pic/port:channel
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index snmp-index>
<statistics>
```

Description

(M Series and T Series routers only) Display status information about the specified flow monitoring interface.

Options

<i>mo-fpc/pic/port:channel</i>	Display standard status information about the specified flow monitoring interface.
brief detail extensive terse	(Optional) Display the specified level of output.
descriptions	(Optional) Display interface description strings.
media	(Optional) Display media-specific information about network interfaces.
snmp-index <i>snmp-index</i>	(Optional) Display information for the specified SNMP index of the interface.
statistics	(Optional) Display static interface statistics.

Required Privilege Level

view

Output Fields

[Table 143 on page 1605](#) lists the output fields for the `show interfaces` (Flow Monitoring) command. Output fields are listed in the approximate order in which they appear.

Table 143: show interfaces Output Fields (Flow Monitoring)

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Link	Status of the link: up or down .	All levels
Enabled	State of the interface. Possible values are described in the “Enabled Field” section under <i>Common Output Fields Description</i> .	All levels
Interface index	Physical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	SNMP index number for the physical interface.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Description	Description and name of the interface.	All levels
Type	Type of interface.	All levels
Link-level type	Encapsulation type used on the physical interface.	All levels
MTU	Maximum Transmit Unit (MTU). Size of the largest packet to be transmitted.	All levels
Clocking	Reference clock source of the interface.	All levels
Speed	Network speed on the interface.	All levels

Table 143: show interfaces Output Fields (Flow Monitoring) (Continued)

Field Name	Field Description	Level of Output
Device flags	Information about the physical device. Possible values are described in the “Device Flags” section under <i>Common Output Fields Description</i> .	All levels
Interface flags	Information about the interface. Possible values are described in the “Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Link type	Data transmission type.	All levels
Link flags	Information about the link. Possible values are described in the “Link Flags” section under <i>Common Output Fields Description</i> .	All levels
Physical info	Information about the physical interface.	All levels
Hold-times	Current interface hold-time up and hold-time down. Value is in milliseconds.	detail extensive
Current address	Configured MAC address.	detail extensive none
Hardware address	Media access control (MAC) address of the interface.	detail extensive none
Alternate link address	Backup link address.	detail extensive none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: <i>year-month-day hour.minute.second timezone (hour.minute.second ago)</i> . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)	detail extensive
Statistics last cleared	Time when the statistics for the interface were last set to zero.	detail extensive

Table 143: show interfaces Output Fields (Flow Monitoring) *(Continued)*

Field Name	Field Description	Level of Output
Traffic statistics	<p>Number and rate of bytes and packets received and transmitted on the physical interface.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Input errors	<ul style="list-style-type: none"> • Errors—Input errors on the interface. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Frames received smaller than the runt threshold. • Giants—Frames received larger than the giant threshold. • Policed Discards—Frames that the incoming packet match code discarded because the frames did not recognize them or were not of interest. Usually, this field reports protocols that Junos does not support. • Resource errors—Sum of transmit drops. 	extensive

Table 143: show interfaces Output Fields (Flow Monitoring) (Continued)

Field Name	Field Description	Level of Output
Output errors	<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 then up, or another problem occurs. If the number of carrier transitions increments quickly, possibly once every 10 seconds, the cable, the remote system, or the interface is malfunctioning. • Errors—Sum of outgoing frame terminates 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 dropped by the ASIC Red mechanism. • Resource errors—Sum of transmit drops. 	extensive
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Logical interface index number, which reflects its initialization sequence.	detail extensive none
SNMP ifIndex	Logical interface SNMP interface index number.	detail extensive none
Generation	Unique number for use by Juniper Networks technical support only.	detail extensive
Flags	Information about the logical interface; values are described in the “Logical Interface Flags” section under <i>Common Output Fields Description</i> .	All levels
Encapsulation	Encapsulation on the logical interface.	All levels

Table 143: show interfaces Output Fields (Flow Monitoring) (Continued)

Field Name	Field Description	Level of Output
Traffic statistics	<p>Total number of bytes and packets received and transmitted on the logical interface. These statistics are the sum of the local and transit statistics. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p> <ul style="list-style-type: none"> • Input bytes, Output bytes—Number of bytes received and transmitted on the interface. • Input packets, Output packets—Number of packets received and transmitted on the interface. 	detail extensive
Local statistics	<p>Statistics for traffic received from and transmitted to the Routing Engine. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p>	detail extensive
Transit statistics	<p>Statistics for traffic transiting the router. When a burst of traffic is received, the value in the output packet rate field might briefly exceed the peak cell rate. It takes awhile (generally, less than 1 second) for this counter to stabilize.</p>	detail extensive
Protocol	<p>Protocol family configured on the logical interface (such as iso or inet6).</p>	detail extensive none
MTU	<p>MTU size on the logical interface.</p>	detail extensive none
Generation	<p>Unique number for use by Juniper Networks technical support only.</p>	detail extensive
Route table	<p>Route table in which this address exists; for example, Route table:0 refers to inet.0.</p>	detail extensive
Flags	<p>Information about the protocol family flags. Possible values are described in the “Family Flags” section under <i>Common Output Fields Description</i>.</p>	detail extensive none

Sample Output

show interfaces extensive (Flow Monitoring)

```

user@host> show interfaces mo-4/0/0    extensive
Physical interface: mo-4/0/0, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 42, Generation: 28
  Description: monitor pic 2
  Type: Adaptive-Services, Link-level type: Adaptive-Services, MTU: Unlimited,
  Clocking: Unspecified, Speed: 800mbps
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-05-24 16:43:12 PDT (00:17:46 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :          756824218          8328536 bps
    Output bytes  :          872916185          8400160 bps
    Input packets :          508452          697 pps
    Output packets:          15577196          18750 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0,
    Policed discards: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 2, Errors: 0, Drops: 0, MTU errors: 0,
    Resource errors: 0

Logical interface mo-4/0/0.0 (Index 83) (SNMP ifIndex 43) (Generation 26)
  Flags: Point-To-Point SNMP-Traps Encapsulation: Adaptive-Services
  Traffic statistics:
    Input bytes   :          756781796
    Output bytes  :          872255328
    Input packets :          507233
    Output packets:          15575988
  Local statistics:
    Input bytes   :          0
    Output bytes  :          0

```

```

Input  packets:          0
Output packets:          0
Transit statistics:
Input  bytes   :          756781796          8328536 bps
Output bytes   :          872255328          8400160 bps
Input  packets:          507233             697 pps
Output packets:          15575988          18750 pps
Protocol inet, MTU: Unlimited, Generation: 38, Route table: 0
Flags: None

Logical interface mo-4/0/0.16383 (Index 84) (SNMP ifIndex 58) (Generation 27)
...

```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring error

IN THIS SECTION

- [Syntax | 1611](#)
- [Description | 1612](#)
- [Options | 1612](#)
- [Required Privilege Level | 1612](#)
- [Output Fields | 1612](#)
- [Sample Output | 1614](#)
- [Release Information | 1614](#)

Syntax

```
show passive-monitoring error (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring error statistics.

Options

*** | all | mo-*fpc/pic/port*** Display error statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

[Table 144 on page 1612](#) lists the output fields for the `show passive-monitoring error` command. Output fields are listed in the approximate order in which they appear.

Table 144: show passive-monitoring error Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Interface state	<p>State of the passive monitoring interface:</p> <ul style="list-style-type: none"> • Monitoring—Specified interface is actively monitoring. • Disabled—Specified interface has been disabled from the CLI. • Not monitoring—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • Unknown—Unknown state. • Error—An error occurred during the process of determining the state of the interface.

Table 144: show passive-monitoring error Output Fields *(Continued)*

Field Name	Field Description
Error information	
Packets dropped (no memory)	Number of packets dropped because of memory shortage.
Packets dropped (not IP)	Number of non-IP packets dropped.
Packets dropped (not IPv4)	Number of packets dropped because they failed the IPv4 version check.
Packets dropped (header too small)	Number of packets dropped because the packet length or IP header length was too small.
Memory allocation failures	Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space.
Memory free failures	Number of flow record memory free failures.
Memory free list failures	Number of flow records received from free list that failed. Memory is nearly exhausted or too many new flows greater than 128 KB are being created per second.
Memory warning	Whether the flows have exceeded 1 million packets per second (Mpps) on a Monitoring Services PIC or 2 Mpps on a Monitoring Services II PIC. The response can be Yes or No .
Memory overload	Whether the memory has been overloaded. The response can be Yes or No .
PPS overload	Whether the PIC is receiving more packets per second than the configured threshold. The response can be Yes or No .
BPS overload	Whether the PIC is receiving more bits per second than the configured threshold. The response can be Yes or No .

Sample Output

show passive-monitoring error all

```
user@host> show passive-monitoring error all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
Error information
  Packets dropped (no memory): 0, Packets dropped (not IP): 0
  Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
  Memory allocation failures: 0, Memory free failures: 0
  Memory free list failures: 0
  Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Not monitoring
Error information
  Packets dropped (no memory): 0, Packets dropped (not IP): 0
  Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
  Memory allocation failures: 0, Memory free failures: 0
  Memory free list failures: 0
  Memory warning: No, Memory overload: No, PPS overload: No, BPS overload: No
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring flow

IN THIS SECTION

- [Syntax | 1615](#)
- [Description | 1615](#)
- [Options | 1615](#)
- [Required Privilege Level | 1615](#)

- [Output Fields | 1615](#)
- [Sample Output | 1617](#)
- [Release Information | 1617](#)

Syntax

```
show passive-monitoring flow (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive flow statistics.

Options

*** | all | mo-fpc/pic/port** Display passive flow statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

[Table 145 on page 1615](#) lists the output fields for the `show passive-monitoring flow` command. Output fields are listed in the approximate order in which they appear.

Table 145: show passive-monitoring flow Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.

Table 145: show passive-monitoring flow Output Fields (*Continued*)

Field Name	Field Description
Interface state	<p>State of the passive monitoring interface:</p> <ul style="list-style-type: none"> • Monitoring—Specified interface is actively monitoring. • Disabled—Specified interface has been disabled from the CLI. • Not monitoring—The interface is operational, but not monitoring. This condition occurs when an interface first comes online, or when the interface is operational, but no logical unit has been configured under the physical interface. • Unknown—Unknown state. • Error—An error occurred during the process of determining the state of the interface.
Flow information	
Flow packets	Number of packets received by an operational PIC.
Flow bytes	Number of bytes received by an operational PIC.
Flow packets 10-second rate	Number of packets per second handled by the PIC and displayed as a 10-second average.
Flow bytes 10-second rate	Number of bytes per second handled by the PIC and displayed as a 10-second average.
Active flows	Number of currently active flows tracked by the PIC.
Total flows	Total number of flows received by an operational PIC.
Flows exported	Total number of flows exported by an operational PIC.
Flows packets exported	Total number of cflowd packets exported by an operational PIC.

Table 145: show passive-monitoring flow Output Fields (*Continued*)

Field Name	Field Description
Flows inactive timed out	Total number of flows that are exported because of inactivity.
Flows active timed out	Total number of long-lived flows that are exported because of an active timeout.

Sample Output

show passive-monitoring flow all

```

user@host> show passive-monitoring flow all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
  Flow information
    Flow packets: 6533434, Flow bytes: 653343400
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 1599
    Flows exported: 1599, Flows packets exported: 55
    Flows inactive timed out: 1599, Flows active timed out: 0

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Monitoring
  Flow information
    Flow packets: 6537780, Flow bytes: 653778000
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 1601
    Flows exported: 1601, Flows packets exported: 55
    Flows inactive timed out: 1601, Flows active timed out: 0

```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring memory

IN THIS SECTION

- [Syntax | 1618](#)
- [Description | 1618](#)
- [Options | 1618](#)
- [Required Privilege Level | 1618](#)
- [Output Fields | 1618](#)
- [Sample Output | 1619](#)
- [Release Information | 1620](#)

Syntax

```
show passive-monitoring memory (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring memory and flow record statistics

Options

*** | all | mo-*fpc/pic/port*** Display memory and flow record statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

[Table 146 on page 1619](#) lists the output fields for the `show passive-monitoring memory` command. Output fields are listed in the approximate order in which they appear.

Table 146: show passive-monitoring memory Output Fields

Field Name	Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Memory utilization	
Allocation count	Number of flow records allocated.
Free count	Number of flow records freed.
Maximum allocated	Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time.
Allocations per second	Flow records allocated per second during the last statistics interval on the PIC.
Frees per second	Flow records freed per second during the last statistics interval on the PIC.
Total memory used, Total memory free	Total memory currently used and total amount of memory currently free (in bytes).

Sample Output

show passive-monitoring memory all

```

user@host> show passive-monitoring memory all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Memory utilization
Allocation count: 1600, Free count: 1599, Maximum allocated: 1600
Allocations per second: 3200, Frees per second: 1438
Total memory used (in bytes): 103579176, Total memory free (in bytes): 163914184

```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring status

IN THIS SECTION

- [Syntax | 1620](#)
- [Description | 1620](#)
- [Options | 1620](#)
- [Required Privilege Level | 1620](#)
- [Output Fields | 1621](#)
- [Sample Output | 1622](#)
- [Release Information | 1622](#)

Syntax

```
show passive-monitoring status (*| all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring status.

Options

| all | mo-*fpc/pic/port Display status for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

Table 147 on page 1621 lists the output fields for the `show passive-monitoring status` command. Output fields are listed in the approximate order in which they appear.

Table 147: show passive-monitoring status Output Fields

Output Field	Output Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
Interface state	Monitoring state of the passive monitoring interface. <ul style="list-style-type: none"> • Monitoring—PIC is actively monitoring. • Disabled—PIC has been disabled using the CLI. • Not monitoring—PIC is operational, but not monitoring. This condition can happen while the PIC is coming online, or when the PIC is operational but has no logical unit configured under the physical interface. • Unknown
Group index	Integer that represents the monitoring group of which the PIC is a member. Group index is a mapping from the group name to an index. It is not related to the number of monitoring groups.
Export interval	Configured export interval for cflowd records, in seconds.
Export format	Configured export format (only cflowd version 5 is supported).
Protocol	Protocol the PIC is configured to monitor (only IPv4 is supported).
Engine type	Configured engine type that is inserted in output cflowd packets.
Engine ID	Configured engine ID that is inserted in output cflowd packets.

Sample Output

show passive-monitoring status all

```
user@host> show passive-monitoring status all
Passive monitoring interface: mo-4/0/0, Local interface index: 44
Interface state: Monitoring
  Group index: 0
  Export interval: 15 secs, Export format: cflowd v5
  Protocol: IPv4, Engine type: 1, Engine ID: 1

Passive monitoring interface: mo-4/1/0, Local interface index: 45
Interface state: Disabled

Passive monitoring interface: mo-4/2/0, Local interface index: 46
Interface state: Not monitoring
```

Release Information

Command introduced before Junos OS Release 7.4.

show passive-monitoring usage

IN THIS SECTION

- [Syntax | 1623](#)
- [Description | 1623](#)
- [Options | 1623](#)
- [Required Privilege Level | 1623](#)
- [Output Fields | 1623](#)
- [Sample Output | 1624](#)
- [Release Information | 1624](#)

Syntax

```
show passive-monitoring usage (* | all | mo-fpc/pic/port)
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring usage statistics.

Options

*** | all | mo-fpc/pic/port** Display usage statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.

Required Privilege Level

view

Output Fields

[Table 148 on page 1623](#) lists the output fields for the `show passive-monitoring usage` command. Output fields are listed in the approximate order in which they appear.

Table 148: show passive-monitoring usage Output Fields

Output Field	Output Field Description
Passive monitoring interface	Name of the passive monitoring interface.
Local interface index	Index counter of the local interface.
CPU utilization	
Uptime	Time, in milliseconds, that the PIC has been operational.
Interrupt time	Total time that the PIC has spent processing packets since the last PIC reset.

Table 148: show passive-monitoring usage Output Fields (Continued)

Output Field	Output Field Description
Load (5 second)	CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.
Load (1 minute)	CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.

Sample Output

show passive-monitoring usage all

```

user@host> show passive-monitoring usage
Passive monitoring interface: mo-4/0/0, Local interface index: 44
  CPU utilization
    Uptime: 653155 milliseconds, Interrupt time: 40213754 microseconds
    Load (5 second): 20%, Load (1 minute): 17%

Passive monitoring interface: mo-4/1/0, Local interface index: 45
  CPU utilization
    Uptime: 652292 milliseconds, Interrupt time: 40223178 microseconds
    Load (5 second): 22%, Load (1 minute): 15%

Passive monitoring interface: mo-4/2/0, Local interface index: 46
  CPU utilization
    Uptime: 649491 milliseconds, Interrupt time: 40173645 microseconds
    Load (5 second): 22%, Load (1 minute): 10098862%

```

Release Information

Command introduced before Junos OS Release 7.4.

show route rpm-tracking

IN THIS SECTION

- [Syntax | 1625](#)
- [Description | 1625](#)
- [Options | 1625](#)
- [Required Privilege Level | 1626](#)
- [Output Fields | 1626](#)
- [Sample Output | 1627](#)
- [Release Information | 1630](#)

Syntax

```
show route rpm-tracking
<destination IP-address>
<owner name>
<test test-name>
```

Description

Display a brief summary of state of RPM-tracked routes along with the current state for a given test.

Options

destination <i>IP-address</i>	(Optional) Show RPM-tracked routes for a particular IPv4 or IPv6 destination prefix.
owner <i>name</i>	The probe owner as configured at the [edit services rpm probe <i>owner</i>] hierarchy level.
test <i>test-name</i>	The test name as configured at the [edit services rpm probe <i>owner</i> test <i>test-name</i>] hierarchy level.

Required Privilege Level

view

Output Fields

Table 149 on page 1626 lists the output fields for the `show route rpm-tracking` command. You can filter on routing table name, destination prefix, RPM probe owner, and RPM test name. If no filter is present all RPM-tracked routes are displayed. Output fields are listed in the approximate order in which they appear.

Table 149: shows route rpm-tracking Output Fields

Field Name	Field Description
Destination	Displays the IPv4 or IPv6 address and optional prefix length of the configured target address.
Next-Hop <ul style="list-style-type: none"> ucst ulst 	Specifies the IPv4 or IPv6 next-hop address of the route to be injected during failover. When there are multiple next-hop entries, a type attribute is shown to indicate whether it is a single unicast next-hop, ucst, or a list of unicast next-hops, ulst.
Metric	Specifies a number that is associated with route preference; when multiple routes have the same preference, the route with lowest metric is made active in the routing table.
Preference	Specifies a number that is used to select routes to destinations in external autonomous systems (ASs) or routing domains. The route with lowest preference value is selected by the routing protocol process as the active one. This column in the output appears only if you configured the preference statement at the [edit routing-options rpm-tracking route <i>destination-prefix</i>] hierarchy level.

Table 149: shows route rpm-tracking Output Fields (*Continued*)

Field Name	Field Description
Route-tag	Specifies a number that is used to represent and advertise a group of routes throughout the routing domain. The route with the lowest tag value is selected by the routing protocol process as the active one. This column in the output appears only if you configured the tag statement at the [edit routing-options rpm-tracking route <i>destination-prefix</i>] hierarchy level.
Owner	Name of the probe owner.
Test Name	Name of the test.
State	Display the state of the route injection action. Routes added to the routing table appear as active in RPM. The initial state of an RPM-tracked route, that is, before the first completion of its associated RPM test, is inactive. Routes removed from the routing table appear as inactive.

Sample Output

show route rpm-tracking

```
user@host> show route rpm-tracking
```

```
Routing table: inet.0
```

Destination	Next-Hop	Metric	Owner	Test Name	State
10.10.10.0/24	10.10.10.11	1	probe1	test1	Active
10.10.20.0/24		1			
	10.10.10.22		probe1	test2	Active
	10.10.10.33		probe1	test3	Inactive
10.10.30.0/24		1			
	10.1.010.11		probe1	test1	Active

```
Routing table: inet6.0
```

Destination	Next-Hop	Metric	Owner	Test Name	State
2001:db8:10::/64	2001:db8:10::11	1	probe1	test1	Active
2001:db8:20::/64		1			
	2001:db8:10::22		probe1	test2	Active
	2001:db8:10::33		probe1	test3	Inactive

```

10.10.20.0/24          1
                      2001:db8:10::11      probe1  test1      Active

```

show route rpm-tracking when preference and tag are configured

```
user@host> show route rpm-tracking
```

Routing table: inet.0

Destination	Next-Hop	Metric	preference	route-tag	Owner	Test Name	State
10.10.10.0/0	-	5	50	12	-	-	
-							

	10.10.10.11	-	-	-	RPM-OWNER	RPM-TEST	Active
Destination	Next-Hop	Metric	preference	route-tag	Owner	Test Name	State
10.10.10.0/16	-	2	18	15	-	-	
-							
	10.10.10.22	-	-	-	RPM-OWNER	RPM-TEST1	
Inactive							
	10.10.10.33	-	-	-	RPM-OWNER	RPM-TEST	Active

Routing table: inet6.0

Destination	Next-Hop	Metric	preference	route-tag	Owner	Test Name	State
2001:db8:10::/64	-	2	11	4	-	-	
-							
	2001:db8:10::11	-	-	-	RPM-OWNER	RPM-TEST	
Active							

show route rpm-tracking destination *destination-prefix*

```
user@host> show route rpm-tracking destination 10.39.0.0/16
```

Routing table: inet.0

Destination	Next-Hop	Metric	Owner	Test Name	State
10.39.0.0/16	10.20.21.2	2	probe-delegate	test7984	
Active					
10.39.1.0/16	10.20.21.3	2	probe-delegate	test7985	

```

Active
10.39.2.0/16          10.20.21.4          2          probe-delegate test7986
Active
10.39.3.0/16          10.20.21.5          2          probe-delegate test7987
Active
10.39.4.0/16          10.20.21.6          2          probe-delegate test7988
Active

```

show route rpm-tracking destination *destination-prefix* when preference and tag are configured

```

user@router> show route rpm-tracking destination 10.10.10.0/32

Routing table: inet.0
Destination      Next-Hop      Metric  preference  route-tag  Owner      Test Name  State
10.10.10.0/32    -              1       14          20         -          -          -
-                10.10.10.2    -       -           -          probe-1    test-1-1   Active

```

show route rpm-tracking destination *destination-prefix* owner *name* test *test-name*

```

user@host> show route rpm-tracking destination 10.39.0.0/16 owner probe-delegate test test7998

Destination      Next-Hop      Metric  Owner      Test Name  State
10.39.14.0/24     10.20.21.2    2       probe-delegate test7998
Active            inet.0

```

show route rpm-tracking destination *destination-prefix* owner *name* test *test-name* when preference and tag are configured

```

user@router> show route rpm-tracking destination 10.10.10.0/32 owner probe-1 test test-1-1

Destination      Next-Hop      Metric  preference  route-tag  Owner      Test Name  State
10.10.10.0/32    -              1       14          20         -          -          -
-                inet.0
                10.10.10.2    -       -           -          probe-1    test-1-1   Active

```


Release Information

Command introduced in Junos OS Release 18.4 R1.

Output showing multiple next hops added in Junos OS Release 19.1R1.

Output showing new preference and route-tag fields added in Junos OS Release 20.4R1.

RELATED DOCUMENTATION

[rpm-tracking](#) | [1375](#)

[Understanding Using Probes for Real-Time Performance Monitoring on M, T, ACX, MX, and PTX Series Routers, EX and QFX Switches](#) | [646](#)

show services accounting aggregation

IN THIS SECTION

- [Syntax](#) | [1630](#)
- [Description](#) | [1631](#)
- [Options](#) | [1631](#)
- [Additional Information](#) | [1632](#)
- [Required Privilege Level](#) | [1632](#)
- [Output Fields](#) | [1632](#)
- [Sample Output](#) | [1634](#)
- [Release Information](#) | [1636](#)

Syntax

```
show services accounting aggregation aggregation-type <aggregation-value>
<detail | extensive | terse>
<limit limit-value>
```

```
< name service-name>
<order (bytes | packets)>
```

Description

Display information about the aggregated active flows being processed by the accounting service.

Options

<i>aggregation-type</i> <i><aggregation-value></i>	<p>Display information for the specified aggregation type and optional value:</p> <ul style="list-style-type: none"> • as <i><source-as-value destination-as-value input-snmp-interface-index-value output-snmp-interface-index-value></i>—Aggregate by autonomous system (AS). • destination-prefix <i><destination-prefix-value destination-as-value output-snmp-interface-index-value></i>—Aggregate by destination prefix. • protocol-port <i><protocol-value source-port-value destination-port-value></i>—Aggregate by protocol and port. • source-destination-prefix <i><source-prefix-value destination-prefix-value destination-as-value source-as-value input-snmp-interface-index-value output-snmp-interface-index-value></i>—Aggregate by source and destination prefix. • source-prefix <i><source-prefix-value source-as-value input-snmp-interface-index-value></i>—Aggregate by source prefix.
detail extensive terse	(Optional) Display the specified level of output.
limit <i>limit-value</i>	(Optional) Limit the display output to the specified number of flows. The default is no limit.
name <i>service-name</i>	(Optional) Display information about the aggregated flows for a specified service name.
order (bytes packets)	(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.

Additional Information

For information about aggregation configuration options, see the [Junos OS Services Interfaces Library for Routing Devices](#).

Required Privilege Level

view

Output Fields

Table 150 on page 1632 lists the output fields for the `show services accounting aggregation` command. Output fields are listed in the approximate order in which they appear.

Table 150: `show services accounting aggregation` Output Fields

Field Name	Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index corresponding to the service accounting interface.
Service name	Name of a service that was configured at the <code>[edit forwarding-options accounting]</code> hierarchy level. The default display, (default sampling) , indicates the service was configured at the [edit forwarding-options sampling-level] hierarchy level.
Protocol	Protocol identifier and number.
Source Port	Source port identifier and number.
Destination Port	Destination port identifier and number.
Source-AS	Source autonomous system (AS) number.
Destination-AS	Destination AS number.

Table 150: show services accounting aggregation Output Fields (*Continued*)

Field Name	Field Description
Source Prefix	Source prefix.
Destination Prefix	Destination prefix.
Source address	Source address.
Source prefix length	Source prefix length.
Destination address	Destination address.
Destination prefix length	Destination prefix length.
Input SNMP interface index	SNMP index of the interface the packet came in on.
Output SNMP interface index	SNMP index of the interface the packet went out on.
Start time	Actual time when the packet in this aggregation was first seen.
End time	Actual time when the packet in this aggregation was last seen.
Flow count	Number of flows in the aggregation.
Packet count	Number of packets in the aggregation.
Byte count	Number of bytes in the aggregation.

Sample Output

show services accounting aggregation protocol-port detail

```

user@host> show service accounting aggregation protocol-port detail
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: (default sampling)
  Protocol: 6, Source port: 20, Destination port: 20
  Start time: 442349, End time: 6425714
  Flow count: 194, Packet count: 4294964388, Byte count: 4294781184

  Protocol: 0, Source port: 0, Destination port: 0
  Start time: 442349, End time: 6425749
  Flow count: 204, Packet count: 4294964324, Byte count: 4294777088

  Protocol: 17, Source port: 123, Destination port: 123
  Start time: 442364, End time: 6425784
  Flow count: 186, Packet count: 4294964152, Byte count: 4294766080

```

show services accounting aggregation source-destination-prefix

```

user@host> show service accounting aggregation source-destination-prefix
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting

```

Source prefix	Destination prefix	Input interface	Output interface	Flow count	Packet count	Byte count
192.0.2.0/20	198.51.100.0/24	ge-5/0/1.0	ge-5/0/0.0	256	491761	31472704
192.0.2.0/20	203.0.113.36/32	ge-5/0/1.0	ge-5/0/0.0	1	1926	123264
192.0.2.0/20	203.0.113.59/32	ge-5/0/1.0	ge-5/0/0.0	1	1926	123264
192.0.2.0/20	192.168.0.63/32	ge-5/0/1.0	ge-5/0/0.0	1	1925	123200
192.0.2.0/20	192.168.0.32/32	ge-5/0/1.0	ge-5/0/0.0	1	1925	

show services accounting aggregation source-destination- prefix order packet detail

```

user@host> show service accounting aggregation source-destination-prefix order packet detail
name t2 input-snmp-interface-index 538
Service Accounting interface: mo-2/0/0, Local interface index: 468

```

Service name: t2

Source	Destination	Input SNMP	Output	SNMP	Flow	Packet Byte
Prefix	Prefix	Index	Index	Count	Count	Count
10.1.1.2/20	192.168.167.1/0	538	432	1	60	46483
10.1.1.2/20	192.168.168.1/0	538	432	1	60	5191
10.1.1.2/20	192.168.154.1/0	538	432	2	60	45504
10.1.1.2/20	192.168.76.1/0	538	432	1	60	42177
10.1.1.2/20	192.168.149.1/0	538	432	1	60	49184
10.1.1.2/20	192.168.113.1/0	538	432	2	60	48757

show services accounting aggregation source-destination- prefix extensive limit

```
user@host> show service accounting aggregation source-destination-prefix name t2 extensive limit
3
```

Service Accounting interface: mo-2/0/0, Local interface index: 542

Service name: t2

Source address: 10.1.1.2, Source prefix length: 20

Destination address: 192.168.200.176.1, Destination prefix length: 0

Input SNMP interface index: 24, Output SNMP interface index: 26

Source-AS: 69, Destination-AS: 69

Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003

Flow count: 0, Packet count: 6, Byte count: 5340

Source address: 10.1.1.2, Source prefix length: 20

Destination address: 192.168.160.1, Destination prefix length: 0

Input SNMP interface index: 24, Output SNMP interface index: 26

Source-AS: 69, Destination-AS: 69

Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003

Flow count: 0, Packet count: 6, Byte count: 5490

Source address: 10.1.1.2, Source prefix length: 20

Destination address: 192.168.160.1, Destination prefix length: 0

Input SNMP interface index: 24, Output SNMP interface index: 26

Source-AS: 69, Destination-AS: 69

Start time: Fri Feb 21 14:16:57 2003, End time: Fri Feb 21 14:22:50 2003

Flow count: 0, Packet count: 6, Byte count: 4079

show services accounting aggregation source-destination-prefix name terse

```
user@host> show service accounting aggregation source-destination-prefix name T3 terse
Service Accounting interface: rsp0, Local interface index: 171
Service name: T3
Interface state: Accounting
```

Source prefix	Destination prefix	Input interface	Output interface	Flow count	Packet count	Byte count
10.1.0.0/20	192.168.3.0/24	ge-5/0/1.0	ge-5/0/0.0	256	639822	40948608
10.1.0.0/20	192.168.2.67/32	ge-5/0/1.0	ge-5/0/0.0	1	2485	159040
10.1.0.0/20	192.168.2.92/32	ge-5/0/1.0	ge-5/0/0.0	1	2485	

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting aggregation template

IN THIS SECTION

- [Syntax | 1636](#)
- [Description | 1637](#)
- [Options | 1637](#)
- [Required Privilege Level | 1637](#)
- [Output Fields | 1637](#)
- [Sample Output | 1638](#)
- [Release Information | 1638](#)

Syntax

```
show services accounting aggregation template
<template-name template-name>
```

Description

Display information for flow aggregation version 9 templates.

Options

- none

Display information for all flow aggregation version version 9 templates.
- template-name *template-name*

(Optional) Display information for the specified template only.

Required Privilege Level

view

Output Fields

[Table 151 on page 1637](#) lists the output fields for the show services accounting aggregation template command. Output fields are listed in the approximate order in which they appear.

Table 151: show services accounting aggregation template Output Fields

Field Name	Field Description
MPLS Label 1	Position of first MPLS label.
MPLS Label 2	Position of second MPLS label.
MPLS Label 3	Position of third MPLS label.
MPLS Top Level Address	Outer top label FEC IP address.
Packet Count	Number of packets sent.

Sample Output

show services accounting aggregation template template-name

```
user@host> show services accounting aggregation template template-name mpls
MPLS label 1: 299808, MPLS label 2: 0, MPLS label 3: 0
Source address: 192.0.2.2, Destination address: 10.255.15.22, Top Label Address: 198.51.100.10
Source port: 0, Destination port: 0
Protocol: 61, TOS: 0, TCP flags: 0
Source mask: 24, Destination mask: 32
Input SNMP interface index: 503, Output SNMP interface index: 505
Start time: 40780, End time: 157330
Packet count: 3949198, Byte count: 181663062
```

Release Information

Command introduced in Junos OS Release 8.3.

show services accounting errors

IN THIS SECTION

- [Syntax | 1639](#)
- [Description | 1639](#)
- [Options | 1639](#)
- [Required Privilege Level | 1639](#)
- [Output Fields | 1639](#)
- [Sample Output | 1641](#)
- [Sample Output | 1642](#)
- [Release Information | 1645](#)

Syntax

```
show services accounting errors
<inline-jflow | name (* | all | service-name)>
```

Description

Display active flow error statistics.

Options

- none** Display error statistics for all services accounting instances.
- inline-jflow fpc-slot *slot-number*** (Optional) Display error statistics for inline jflow.
- name (* | all | *service-name*)** (Optional) Display active flow error statistics. Use a wildcard character, specify all services, or provide a specific service name.

Required Privilege Level

view

Output Fields

[Table 152 on page 1639](#) lists the output fields for the `show services accounting errors` command. Output fields are listed in the approximate order in which they appear.

Table 152: show services accounting errors Output Fields

Field	Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index counter of the local interface.

Table 152: show services accounting errors Output Fields (*Continued*)

Field	Field Description
FPC slot	Slot number of the FPC for which the flow information is displayed. (Available only when the <code>inline-jflow fpc-slot slot-number</code> option is used.)
Service name	Name of a service that was configured at the [edit forwarding-options accounting] hierarchy level. The default display, (default sampling) , indicates the service was configured at the [edit forwarding-options sampling-level] hierarchy level.
Error Information	
Packets dropped (no memory)	Number of packets dropped because of memory shortage.
Packets dropped (not IP)	Number of non-IP packets dropped.
Packets dropped (not IPv4)	Number of packets dropped because they failed the IPv4 version check.
Packets dropped (header too small)	Number of packets dropped because the packet length or IP header length was too small.
Memory allocation failures	Number of flow record memory allocation failures. A small number reflects failures to replenish the free list. A large number indicates the monitoring station is almost out of memory space.
Memory free failures	Number of flow record memory free failures.
Memory free list failures	Number of flow records received from the free list that failed. Memory is nearly exhausted, or too many new flows greater than 128 KB are being created per second.
Memory overload	Whether the memory has been overloaded. The response can be Yes or No .
PPS overload	Whether the PIC is receiving more packets per second than the configured threshold. The response can be Yes or No .

Table 152: show services accounting errors Output Fields (*Continued*)

Field	Field Description
BPS overload	Whether the PIC is receiving more bits per second than the configured threshold. The response can be Yes or No .
Flow Creation Failures	Number of times flow creation failed.
Route Record Lookup Failures	Number of times the route record lookup failed.
AS Lookup Failures	Number of times autonomous system lookup failed.
Export Packet Failures	Number of times packet export failed.

Sample Output

show services accounting errors (Monitoring PIC interface)

```

user@host> show services accounting errors
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
Error information
  Packets dropped (no memory): 0, Packets dropped (not IP): 0
  Packets dropped (not IPv4): 0, Packets dropped (header too small): 0
  Memory allocation failures: 0, Memory free failures: 0
  Memory free list failures: 0
  Memory overload: No, PPS overload: No, BPS overload: No

```

Sample Output

show services accounting errors (Service PIC interface)

```
user@host> show services accounting errors
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
Error information
  Service sets dropped: 0, Active timeout failures: 0
  Export packet failures: 0, Flow creation failures: 0
  Memory overload: No
```

```
Service Accounting interface: sp-1/0/0
Service name: (default sampling)
Error information
  Service sets dropped: 0, Active timeout failures: 0
  Export packet failures: 0, Flow creation failures: 0
  Memory overload: No
```

show services accounting errors inline-jflow fpc-slot (When Only IPv6 Is Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 5
Error information
  FPC Slot: 5
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0, AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No, Memory Alloc Fail Count: 0
```

show services accounting errors inline-jflow fpc-slot (When IPv4, IPv6, VPLS, and Bridge Are Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 5
Error information
  FPC Slot: 5
  Flow Creation Failures: 0
```

```

Route Record Lookup Failures: 0, AS Lookup Failures: 0
Export Packet Failures: 0
Memory Overload: No, Memory Alloc Fail Count: 0

IPv4:
IPv4 Flow Creation Failures: 0
IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
IPv4 Export Packet Failures: 0

IPv6:
IPv6 Flow Creation Failures: 0
IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
IPv6 Export Packet Failures: 0

VPLS:
VPLS Flow Creation Failures: 0
VPLS Export Packet Failures: 0

BRIDGE:
BRIDGE Flow Creation Failures: 0
BRIDGE Route Record Lookup Failures: 0, BRIDGE AS Lookup Failures: 0
BRIDGE Export Packet Failures: 0

```

show services accounting errors inline-jflow (MX80 Router When Both IPv4 and IPv6 Are Configured)

```

user@host> show services accounting errors inline-jflow
Error information
  TFEB Slot: 0
  Flow Creation Failures: 0
  Route Record Lookup Failures: 0, AS Lookup Failures: 0
  Export Packet Failures: 0
  Memory Overload: No

  IPv4:
  IPv4 Flow Creation Failures: 0
  IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
  IPv4 Export Packet Failures: 0

  IPv6:

```

```
IPv6 Flow Creation Failures: 0
IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
IPv6 Export Packet Failures: 0
```

show services accounting errors inline-jflow fpc-slot(PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow fpc-slot 0
Error information
FPC Slot: 0
Flow Creation Failures: 0
Route Record Lookup Failures: 0, AS Lookup Failures: 0
Export Packet Failures: 0
Memory Overload: No, Memory Alloc Fail Count: 0

IPv4:
IPv4 Flow Creation Failures: 0
IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
IPv4 Export Packet Failures: 0

IPv6:
IPv6 Flow Creation Failures: 0
IPv6 Route Record Lookup Failures: 0, IPv6 AS Lookup Failures: 0
IPv6 Export Packet Failures: 0
```

show services accounting errors inline-jflow (SRX Series Devices When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting errors inline-jflow
Error information
FPC Slot: 0
Flow Creation Failures: 0
Route Record Lookup Failures: 0, AS Lookup Failures: 0
Export Packet Failures: 0
Memory Overload: No, Memory Alloc Fail Count: 0

IPv4:
IPv4 Flow Creation Failures: 0
```

IPv4 Route Record Lookup Failures: 0, IPv4 AS Lookup Failures: 0
IPv4 Export Packet Failures: 0

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[show services accounting flow](#) | [1645](#)

show services accounting flow

IN THIS SECTION

- [Syntax](#) | [1645](#)
- [Description](#) | [1645](#)
- [Options](#) | [1646](#)
- [Required Privilege Level](#) | [1646](#)
- [Output Fields](#) | [1646](#)
- [Sample Output](#) | [1648](#)
- [Release Information](#) | [1654](#)

Syntax

```
show services accounting flow  
<inline-jflow fpc-slot slot-number | logical-system (all | logical-system) | name (* | all |  
service-name)>
```

Description

Display active flow statistics.

Options

none	Display active flow statistics for all service instances.
logical-system (all <i>logical-system</i>)	(Optional) Display active flow statistics for the specified logical system or all logical systems on the device.
inline-jflow (fpc-slot <i>slot-number</i>)	<p>(Optional) Display inline flow statistics for the specified FPC.</p> <p>For PTX Series, starting in Junos OS Evolved Release 21.2R1 and Junos OS Release 21.3R1, every sampled packet is considered to be a flow. When the sampled packet is received, the flow is created and immediately timed out as inactive, and the software exports a record to the collector. Therefore, when you specify this option, the command now displays 0 for all of the various Active Flows and Timed Out fields. The values of the various Total Flows fields are now equal to their respective Flow Packets field values. The values of the various Flows Inactive Timed Out fields are now equal to their respective Flow Packets field values.</p>
name (* all <i>service-name</i>)	(Optional) Display services accounting active flow statistics. Use a wildcard character, specify all services, or provide a specific service name.

Required Privilege Level

view

Output Fields

Table 153 on page 1646 lists the output fields for the `show services accounting flow` command. Output fields are listed in the approximate order in which they appear.

Table 153: show services accounting flow Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Local interface index	Index counter of the local interface.

Table 153: show services accounting flow Output Fields *(Continued)*

Output Field	Output Field Description
Service name	Name of a service that was configured at the [edit forwarding-options accounting] hierarchy level. The default display, (default sampling) , indicates the service was configured at the [edit forwarding-options sampling-level] hierarchy level.
Flow Information	
FPC Slot	Slot number of the FPC for which the flow information is displayed. (Available only when the inline-jflow fpc-slot <i>slot-number</i> option is used.)
Flow packets	Number of packets received by an operational PIC.
Flow bytes	Number of bytes received by an operational PIC.
Flow packets 10-second rate	Number of packets per second handled by the PIC and displayed as a 10-second average.
Flow bytes 10-second rate	Number of bytes per second handled by the PIC and displayed as a 10-second average.
Active flows	Number of currently active flows tracked by the PIC.
Total flows	Total number of flows received by an operational PIC.
Flows exported	Total number of flows exported by an operational PIC.
Flows packets exported	Total number of cflowd packets exported by an operational PIC.
Flows inactive timed out	Total number of flows that are exported because of inactivity.
Flows active timed out	Total number of long-lived flows that are exported because of an active timeout.

Sample Output

show services accounting flow (Flow Aggregation v5/v8 Configuration)

```
user@host> show services accounting flow
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Flow information
  Flow packets: 87168293, Flow bytes: 5578770752
  Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928962
  Active flows: 1000, Total flows: 2000
  Flows exported: 19960, Flows packets exported: 582
  Flows inactive timed out: 1000, Flows active timed out: 29000
```

show services accounting flow (Flow Aggregation v9 Configuration)

```
user@host> show services accounting flow
Flow information
  Service Accounting interface: sp-7/1/0, Local interface index: 149
  Flow packets: 0, Flow bytes: 0
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 0
  Flows exported: 0, Flows packets exported: 1
  Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow name

```
user@host> show services accounting flow name count2
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: count2
Flow information
  Flow packets: 0, Flow bytes: 0
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 0
  Flows exported: 0, Flows packets exported: 0
  Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow name all

```

user@host> show services accounting flow name all
Service Accounting interface: rsp0, Local interface index: 171
Service name: T2
Interface state: Accounting
  Flow information
    Flow packets: 37609891, Flow bytes: 2407033024
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928953
    Active flows: 1000, Total flows: 1000
    Flows exported: 6705, Flows packets exported: 198
    Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T3
Interface state: Accounting
  Flow information
    Flow packets: 37750807, Flow bytes: 2416051712
    Flow packets 10-second rate: 45762, Flow bytes 10-second rate: 2928940
    Active flows: 1000, Total flows: 1000
    Flows exported: 13437, Flows packets exported: 378
    Flows inactive timed out: 0, Flows active timed out: 13000

Service Accounting interface: rsp0, Local interface index: 171
Service name: T4
Interface state: Accounting
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0
    Flows exported: 0, Flows packets exported: 0
    Flows inactive timed out: 0, Flows active timed out: 0

Service Accounting interface: rsp0, Local interface index: 171
Service name: count1
Interface state: Accounting
  Flow information
    Flow packets: 0, Flow bytes: 0
    Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
    Active flows: 0, Total flows: 0

```

```
Flows exported: 0, Flows packets exported: 0
Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow (Multiple Sampling Instances)

```
user@host> show services accounting flow
Flow information
  Service Accounting interface: sp-2/0/0, Local interface index: 215
  Flow packets: 9867, Flow bytes: 631488
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 628
  Active flows: 2, Total flows: 10
  Flows exported: 4028, Flows packets exported: 6150
  Flows inactive timed out: 8, Flows active timed out: 4026

  Service Accounting interface: sp-2/1/0, Local interface index: 223
  Flow packets: 0, Flow bytes: 0
  Flow packets 10-second rate: 0, Flow bytes 10-second rate: 0
  Active flows: 0, Total flows: 0
  Flows exported: 0, Flows packets exported: 1
  Flows inactive timed out: 0, Flows active timed out: 0
```

show services accounting flow inline-jflow fpc-slot (for IPv4 Flow)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
Flow information
  FPC Slot: 5
  Flow Packets: 0, Flow Bytes: 0
  Active Flows: 0, Total Flows: 0
  Flows Exported: 0, Flow Packets Exported: 0
  Flows Inactive Timed Out: 0, Flows Active Timed Out: 0
```

show services accounting flow inline-jflow fpc-slot (with IPv4, IPv6, VPLS, and Bridge Configuration)

```
user@host> show services accounting flow inline-jflow fpc-slot 5
Flow information
  FPC Slot: 5
  Flow Packets: 0, Flow Bytes: 0
```

```

Active Flows: 0, Total Flows: 0
Flows Exported: 0, Flow Packets Exported: 0
Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

IPv4 Flows:
IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
IPv4 Active Flows: 0, IPv4 Total Flows: 0
IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

VPLS Flows:
VPLS Flow Packets: 0, VPLS Flow Bytes: 0
VPLS Active Flows: 0, VPLS Total Flows: 0
VPLS Flows Exported: 0, VPLS Flow Packets Exported: 0
VPLS Flows Inactive Timed Out: 0, VPLS Flows Active Timed Out: 0

BRIDGE Flows:
BRIDGE Flow Packets: 0, BRIDGE Flow Bytes: 0
BRIDGE Active Flows: 0, BRIDGE Total Flows: 0
BRIDGE Flows Exported: 0, BRIDGE Flow Packets Exported: 0
BRIDGE Flows Inactive Timed Out: 0, BRIDGE Flows Active Timed Out: 0
BRIDGE Flow Insert Count: 0

```

show services accounting flow inline-jflow (MX80 Router with IPv4 and IPv6 Configuration)

```

user@host> show services accounting flow inline-jflow
Flow information
  TFEB Slot: 0
  Flow Packets: 0, Flow Bytes: 0
  Active Flows: 0, Total Flows: 0
  Flows Exported: 0, Flow Packets Exported: 0
  Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

  IPv4 Flows:
  IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0

```

```

IPv4 Active Flows: 0, IPv4 Total Flows: 0
IPv4 Flows Exported: 0, IPv4 Flow Packets exported: 0
IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

```

show services accounting flow inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```

user@host> show services accounting flow inline-jflow fpc-slot 0
Flow information
FPC Slot: 0
Flow Packets: 47427946, Flow Bytes: 5217074060
Active Flows: 0, Total Flows: 2
Flows Exported: 194, Flow Packets Exported: 7045
Flows Inactive Timed Out: 2, Flows Active Timed Out: 192

IPv4 Flows:
IPv4 Flow Packets: 47427946, IPv4 Flow Bytes: 5217074060
IPv4 Active Flows: 0, IPv4 Total Flows: 2
IPv4 Flows Exported: 194, IPv4 Flow Packets exported: 7045
IPv4 Flows Inactive Timed Out: 2, IPv4 Flows Active Timed Out: 192

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0, IPv6 Flow Packets Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

```

show services accounting flow inline-jflow (Junos OS Evolved 21.2R1 and later, for a PTX10003 router)

```

user@host> show services accounting flow inline-jflow fpc-slot 0
Flow information
FPC Slot: 0

```

```

Flow Packets: 0, Flow Bytes: 0
Active Flows: 0, Total Flows: 0
Flows Exported: 0, Flow Packets Exported: 0
Flows Inactive Timed Out: 0, Flows Active Timed Out: 0

IPv4 Flows:
IPv4 Flow Packets: 0, IPv4 Flow Bytes: 0
IPv4 Active Flows: 0, IPv4 Total Flows: 0
IPv4 Flows Exported: 0
IPv4 Flows Inactive Timed Out: 0, IPv4 Flows Active Timed Out: 0

IPv6 Flows:
IPv6 Flow Packets: 0, IPv6 Flow Bytes: 0
IPv6 Active Flows: 0, IPv6 Total Flows: 0
IPv6 Flows Exported: 0
IPv6 Flows Inactive Timed Out: 0, IPv6 Flows Active Timed Out: 0

MPLS Flows:
MPLS Flow Packets: 0, MPLS Flow Bytes: 0
MPLS Active Flows: 0, MPLS Total Flows: 0
MPLS Flows Exported: 0
MPLS Flows Inactive Timed Out: 0, MPLS Flows Active Timed Out: 0

```

show services accounting flow inline-jflow (SRX Series When IPv4 is configured)

```

user@host> show services accounting flow inline-jflow
Flow information
  FPC Slot: 0
  Flow Packets: 462680, Flow Bytes: 45433206
  Active Flows: 34, Total Flows: 61093
  Flows Exported: 138936, Flow Packets Exported: 96649
  Flows Inactive Timed Out: 61083, Flows Active Timed Out: 138936
  Total Flow Insert Count: 0

IPv4 Flows:
IPv4 Flow Packets: 462680, IPv4 Flow Bytes: 45433206
IPv4 Active Flows: 34, IPv4 Total Flows: 61093
IPv4 Flows Exported: 138936, IPv4 Flow Packets exported: 96649
IPv4 Flows Inactive Timed Out: 61083, IPv4 Flows Active Timed Out: 138936
IPv4 Flow Insert Count: 0

```


Release Information

Command introduced before Junos OS Release 7.4.

Junos OS Release 10.0 added the capability to display output from multiple sampling instances.

RELATED DOCUMENTATION

[show services accounting status](#) | [1666](#)

show services accounting flow-detail

IN THIS SECTION

- [Syntax](#) | [1654](#)
- [Description](#) | [1655](#)
- [Options](#) | [1655](#)
- [Additional Information](#) | [1656](#)
- [Required Privilege Level](#) | [1656](#)
- [Output Fields](#) | [1656](#)
- [Sample Output](#) | [1658](#)
- [Release Information](#) | [1661](#)

Syntax

```
show services accounting flow-detail
<detail | extensive | terse>
<filters>
<limit limit-value>
<name (* | all | service-name)>
<order (bytes | packets)>
```

Description

Display information about the flows being processed by the accounting service.

Options

none	Display information about all flows.
detail extensive terse	(Optional) Display the specified level of output.
<i>filters</i>	<p>(Optional) Filter the display output of the currently active flow records. The following filters query actively changing data structures and result in different results for multiple invocations:</p> <ul style="list-style-type: none"> • destination-as—Display flow records filtered by destination autonomous system information. • destination-port—Display flow records filtered by destination port information. • destination-prefix—Display flow records filtered by destination prefix information. • input-snmp-interface-index—Display flow records filtered by SNMP input interface index information. • output-snmp-interface-index—Display flow records filtered by SNMP output interface index information. • proto—Display flow records filtered by protocol type. • source-as—Display flow records filtered by source autonomous system information. • source-port—Display flow records filtered by source port information. • source-prefix—Display flow records filtered by source prefix information. • tos—Display flow records filtered by type of service classification.
limit <i>limit-value</i>	(Optional) Limit the display output to the specified number of flows. The default is no limit.
name (* all <i>service-name</i>)	(Optional) Display information about the flows being processed. Use a wildcard character, specify all services, or provide a specific services name.

order (bytes | packets) (Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.

Additional Information

When no PIC is active, or when no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled. This command displays information about two concurrent sessions only. If a third session is attempted, the command pauses with no output until one of the previous sessions is completed.

Required Privilege Level

view

Output Fields

[Table 154 on page 1656](#) lists the output fields for the `show services accounting flow-detail` command. Output fields are listed in the approximate order in which they appear.

Table 154: show services accounting flow-detail Output Fields

Field Name	Field Description	Output Level
Service Accounting interface	Name of the service accounting interface.	All levels
Service name	Name of a service that was configured at the [edit forwarding-options accounting] hierarchy level. The default display, (default sampling) , indicates the service was configured at the [edit forwarding-options sampling] hierarchy level.	All levels
Local interface index	Index counter of the local interface.	All levels
TOS	Type-of-service value from the IP header.	extensive

Table 154: show services accounting flow-detail Output Fields *(Continued)*

Field Name	Field Description	Output Level
Input SNMP interface index	SNMP index of the interface on which the packet came in.	extensive
Output SNMP interface index	SNMP index of the interface on which the packet went out.	extensive
Source-AS	Source AS number.	extensive
Destination-AS	Destination AS number.	extensive
Protocol	Name of the protocol used for the packet flow from the corresponding source address.	All levels
Input interface	Interface on which the packets were received.	All levels
Output interface	Interface on which the packets were transmitted.	All levels
TCP flags	Number of TCP header flags detected in the flow.	extensive
Source address	Address where the flow originated.	All levels
Source port	Name of the source port.	All levels
Source prefix length	Source prefix length.	extensive
Destination address	Address where the flow is sent.	All levels
Destination prefix length	Destination prefix length.	extensive

Table 154: show services accounting flow-detail Output Fields *(Continued)*

Field Name	Field Description	Output Level
Destination port	Name of the destination port.	All levels
Start time	Actual time when the packet in this aggregation was first seen.	detail extensive
End time	Actual time when the packet in this aggregation was last seen.	detail extensive
Packet count	Number of packets in the aggregation.	All levels
Byte count	Number of bytes in the aggregation.	All levels
Time since last active timeout	Amount of time elapsed since the last active timeout, in the format <i>hh:mm:ss</i> .	None specified
Packet count for last active timeout	Number of packets in the aggregation since the last active timeout.	None specified
Byte count for last active timeout	Number of bytes in the aggregation since the last active timeout.	None specified

Sample Output

show services accounting flow-detail

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```

user@host> show services accounting flow-detail
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Protocol   Input           Source           Source   Output

```

```

      interface      address      port  interface...
tcp(6)  ge-5/0/1.0    192.0.2.2      0    ge-5/0/0.0
tcp(6)  ge-5/0/1.0    192.0.2.2      0    ge-5/0/0.0

Destination      Destination      Packet      Byte  Time since last
address          port          count      count  active timeout...
198.51.100.149      0          2660      170240      00:00:58
198.51.100.138      0          2660      170240      00:00:58

Packet count for      Byte count for
last active timeout    last active timeout
2805                  179520
2805                  179520

```

show services accounting flow-detail limit

In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```

user@host> show services accounting flow-detail limit 1
Service Accounting interface: rsp0, Local interface index: 171
Service name: (default sampling)
Interface state: Accounting
Protocol  Input      Source      Source  Output
         interface  address      port    interface...
tcp(6)   ge-5/0/1.0    192.0.2.2      0    ge-5/0/0.0

Destination      Destination      Packet      Byte  Time since last
address          port          count      count  active timeout...
198.51.100.149      0          2158      138112      00:00:47

Packet count for      Byte count for
last active timeout    last active timeout
2827                  180928

```

show services accounting flow-detail name extensive

```

user@host> show services accounting flow-detail name cf-2 extensive
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2

```

```

TOS: 0, Protocol: udp(17), TCP flags: 0
Source address: 10.10.10.1, Source prefix length: 0, Destination address: 203.0.113.20,
Destination prefix length: 0, Source port: 1173, Destination port: 69
Input SNMP interface index: 65, Output SNMP interface index: 0, Source-AS: 0,
Destination-AS: 0
Start time: 62425, End time: 635265, Packet count: 165845, Byte count: 9453165

```

show services accounting flow-detail limit order bytes

The output of the following command is displayed over 141 columns, not the standard 80 columns. In this sample, the output is split into three sections, with ellipses (...) indicating where the sections are continued.

```

user@host> show services accounting flow-detail limit 5 order bytes
Service Accounting interface: mo-2/0/0, Local interface index: 356
Service name: (default sampling)

```

Protocol	Input interface	Source address	Source port	Output interface...
icmp(1)	ge-2/3/0.0	192.0.2.2	0	.local.
icmp(1)	ge-2/3/0.0	192.0.2.2	0	.local.
icmp(1)	ge-2/3/0.0	192.0.2.2	0	.local.
icmp(1)	ge-2/3/0.0	192.0.2.2	0	.local.
icmp(1)	ge-2/3/0.0	192.0.2.2	0	.local.

Destination address	Destination port	Packet count	Byte count	Time since last active timeout...
192.168.128.2	0	16	12148	Not applicable
192.168.144.2	0	16	15229	Not applicable
192.168.192.2	0	16	13296	Not applicable
192.168.16.2	0	16	13924	Not applicable
192.168.48.2	0	16	13428	Not applicable

Packet count for last active timeout	Byte count for last active timeout
Not applicable	Not applicable
Not applicable	Not applicable
Not applicable	Not applicable
Not applicable	Not applicable
Not applicable	Not applicable

show services accounting flow-detail name detail source-port

```
user@host> show services accounting flow-detail name cf-2 detail source-port 1173
Service Accounting interface: mo-0/2/0, Local interface index: 145
Service name: cf-2
  Protocol: udp(17), Source address: 10.10.10.1, Source port: 1173, Destination address:
203.0.113.20, Destination port: 69
  Start time: 62425, End time: 811115, Packet count: 142438, Byte count: 8118966
```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting memory

IN THIS SECTION

- [Syntax | 1661](#)
- [Description | 1661](#)
- [Options | 1662](#)
- [Required Privilege Level | 1662](#)
- [Output Fields | 1662](#)
- [Sample Output | 1663](#)
- [Release Information | 1664](#)

Syntax

```
show services accounting memory
```

Description

Display memory and flow record statistics.

Options

This command has no options.

Required Privilege Level

view

Output Fields

Table 155 on page 1662 lists the output fields for the show services accounting memory command. Output fields are listed in the approximate order in which they appear.

Table 155: show services accounting memory Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.
Memory Utilization	
Local interface index	Index counter of the local interface.
Allocation count	Number of flow records allocated.
Free count	Number of flow records freed.
Maximum allocated	Maximum number of flow records allocated since the monitoring station booted. This number represents the peak number of flow records allocated at a time.
Allocations per second	Flow records allocated per second during the last statistics interval on the PIC.
Frees per second	Flow records freed per second during the last statistics interval on the PIC.
Total memory used	Total amount of memory currently used (in bytes).

Table 155: show services accounting memory Output Fields (Continued)

Output Field	Output Field Description
Total memory free	Total amount of memory currently free (in bytes).

Sample Output

show services accounting memory (Monitoring PIC Interface)

```

user@host> show                services accounting                memory
Service Accounting interface: mo-2/0/0, Local interface index: 468
Memory utilization
Allocation count: 437340, Free count: 433699, Maximum allocated: 6782
Allocations per second: 3366, Frees per second: 6412
Total memory used (in bytes): 133460320,
Total memory free (in bytes): 133918352

```

show services accounting memory (Service PIC Interface)

```

user@host> show                services accounting                memory
Service Accounting interface: sp-0/1/0
Memory utilization
Allocation count: 1000, Free count: 0
Allocations per second: 0, Frees per second: 0
Total memory used (in bytes): 218158272
Total memory free (in bytes): 587147696

```

```

Service Accounting interface: sp-1/0/0
Memory utilization
Allocation count: 1000, Free count: 0
Allocations per second: 0, Frees per second: 0
Total memory used (in bytes): 218157592
Total memory free (in bytes): 587148376

```

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting packet-size-distribution

IN THIS SECTION

- [Syntax | 1664](#)
- [Description | 1664](#)
- [Options | 1664](#)
- [Required Privilege Level | 1665](#)
- [Output Fields | 1665](#)
- [Sample Output | 1665](#)
- [Release Information | 1666](#)

Syntax

```
show services accounting packet-size-distribution  
<name (* | all | service-name)>
```

Description

Display a packet size distribution histogram.

Options

- | | |
|---|---|
| none | Display a packet size distribution histogram of all accounting services. |
| name (* all <i>service-name</i>) | (Optional) Display a packet size distribution histogram. Use a wildcard character, specify all services, or provide a specific services name. |

Required Privilege Level

view

Output Fields

Table 156 on page 1665 lists the output fields for the `show services accounting packet-size-distribution` command. Output fields are listed in the approximate order in which they appear.

Table 156: show services accounting packet-size-distribution Output Fields

Field Name	Field Description
Service Accounting interface	Name of the service accounting interface.
Service name	Name of a service that was configured at the <code>[edit-forwarding-options accounting]</code> hierarchy level. The default display, <code>(default sampling)</code> , indicates the service was configured at the <code>[edit-forwarding-options sampling-level]</code> hierarchy level.
Local interface index	Index counter of the local interface.
Range start	Smallest packet length (in bytes) to count.
Range end	Largest packet length (in bytes) to count.
Number of packets	Count of packets detected in the size between Range start and Range end .
Percentage packets	Percentage of the total number of packets that are in this size range.

Sample Output

`show services accounting packet-size-distribution name`

```
user@host> show services accounting packet-size-distribution name test3
Service Accounting interface: mo-0/2/0, Local interface index: 163
Service name: test3
```

Range start	Range end	Number of packets	Percentage packets
32	64	2924	100

Release Information

Command introduced before Junos OS Release 7.4.

show services accounting status

IN THIS SECTION

- [Syntax | 1666](#)
- [Description | 1666](#)
- [Options | 1666](#)
- [Required Privilege Level | 1667](#)
- [Output Fields | 1667](#)
- [Sample Output | 1669](#)
- [Sample Output | 1669](#)
- [Release Information | 1671](#)

Syntax

```
show services accounting status
<inline-jflow fpc-slot slot-number | name (* | all | service-name)>
```

Description

Display available Physical Interface Cards (PICs) for accounting services.

Options

none Display available PICs for all accounting services.

- inline-jflow fpc-slot *slot-number*** (Optional) Display inline flow accounting status for the specified FPC. For a two-member MX Series Virtual Chassis or EX9200 Virtual Chassis, the primary router or switch uses FPC slot numbers 0 through 11 with no offset; the backup router or switch uses FPC slot numbers 12 through 23, with an offset of 12.
- name (* | all | *service-name*)** (Optional) Display available PICs. Use a wildcard character, specify all services, or provide a specific services name.

Required Privilege Level

view

Output Fields

[Table 157 on page 1667](#) lists the output fields for the `show services accounting status` command. Output fields are listed in the approximate order in which they appear.

Table 157: show services accounting status Output Fields

Field	Field Description
Service Accounting interface	Name of the service accounting interface.
Service name	Name of a service that was configured at the [edit-forwarding-options accounting] hierarchy level. The default display,(default sampling), indicates the service was configured at the [edit-forwarding-options sampling-level] hierarchy level.
FPC Slot	Slot number of the FPC for which the flow information is displayed.
Local interface index	Index counter of the local interface.

Table 157: show services accounting status Output Fields (Continued)

Field	Field Description
Interface state	Accounting state of the passive monitoring interface. <ul style="list-style-type: none"> Accounting—PIC is actively accounting. Disabled—PIC has been disabled from the CLI. Not accounting—PIC is up but not accounting. This can happen while the PIC is coming online, or when the PIC is up but has no logical unit configured under the physical interface. Unknown
Group index	Integer that represents the monitoring group of which the PIC is a member. Group index is a mapping from the group name to an index. It is not related to the number of monitoring groups.
Export interval (in seconds)	Configured export interval for cflowd records, in seconds.
Export format	Configured export format.
Protocol	Protocol the PIC is configured to monitor.
Engine type	Configured engine type that is inserted in output cflowd packets.
Engine ID	Configured engine ID that is inserted in output cflowd packets.
Route Record Count	Number of routes recorded.
AS Record Count	Number of autonomous systems recorded.
Route Records Set	Status of route recording; whether routes are recorded or not.
Configuration Set	Status of monitoring configuration; whether monitoring configuration is set or not.

Sample Output

show services accounting status name (Monitoring PIC Interface)

```
user@host> show services accounting status name count1
Service Accounting interface: mo-2/0/0, Local interface index: 468
Service name: count1
Interface state: Accounting
  Group index: 0
  Export interval (in seconds): 60, Export format: cflowd v8
  Protocol: IPv4, Engine type: 55, Engine ID: 5
```

Sample Output

show services accounting status name (Service PIC Interface)

```
user@host> show services accounting status name
Service Accounting interface: sp-0/1/0
Interface state: Accounting
  Export format: 9, Route record count: 0
  IFL to SNMP index count: 7, AS count: 0
  Configuration set: Yes, Route record set: No, IFL SNMP map set: Yes
```

```
Service Accounting interface: sp-1/0/0
Interface state: Accounting
  Export format: 9, Route record count: 33
  IFL to SNMP index count: 7, AS count: 1
  Configuration set: Yes, Route record set: Yes, IFL SNMP map set: Yes
```

show services accounting status inline-jflow fpc-slot (When IPv4, IPv6 and Bridge Family Are Configured)

```
user@host> show services accounting status inline-jflow fpc-slot 0
FPC Slot: 0
  IPV4 export format: Version-IPFIX, IPV6 export format: Not set
  BRIDGE export format: Version-IPFIX, MPLS export format: Version-IPFIX
```



```

IPv4 Route Record Count: 31, IPv6 Route Record Count: 0, MPLS Route Record Count: 13
Route Record Count: 44, AS Record Count: 1
Route-Records Set: Yes, Config Set: Yes
Service Status: PFE-0: Steady PFE-1: Steady
Using Extended Flow Memory?: PFE-0: No PFE-1: No
Flex Flow Sizing ENABLED?: PFE-0: No PFE-1: No
IPv4 MAX FLOW Count: 1024, IPv6 MAX FLOW Count: 512
BRIDGE MAX FLOW Count: 1024, MPLS MAX FLOW Count: 1024
MAX Flow Table size: 15

```

show services accounting status inline-jflow (MX80 Router When Both IPv4 and IPv6)

```
user@host> show services accounting status inline-jflow
```

```

Status information
    TFEB Slot: 0
    Export format: IP-FIX
    IPv4 Route Record Count: 6, IPv6 Route Record Count: 8
    Route Record Count: 14, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes

```

show services accounting status inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured)

```
user@host> show services accounting status inline-jflow fpc-slot 0
```

```

Status information
    FPC Slot: 0
    IPV4 export format: Version-IPFIX, IPV6 export format: Version-IPFIX
    MPLS export format: Not set
    IPv4 Route Record Count: 23, IPv6 Route Record Count: 3, MPLS Route Record Count: 0
    Route Record Count: 26, AS Record Count: 1
    Route-Records Set: Yes, Config Set: Yes

```

show services accounting status inline-jflow (SRX Series Devices When Both IPv4 and IPv6 Are Configured)

```

user@host> show services accounting status inline-jflow
Status information

```

```

FPC Slot: 0
IPv4 export format: Version9, IPv6 export format: Version9
BRIDGE export format: Not set, MPLS export format: Not set
IPv4 Route Record Count: 24, IPv6 Route Record Count: 0, MPLS Route Record Count: 0
Route Record Count: 24, AS Record Count: 1
Route-Records Set: Yes, Config Set: Yes
Service Status: PFE-0: Steady
Using Extended Flow Memory?: PFE-0: No
Flex Flow Sizing ENABLED?: PFE-0: No
IPv4 MAX FLOW Count: 0, IPv6 MAX FLOW Count: 0
BRIDGE MAX FLOW Count: 0, MPLS MAX FLOW Count: 0

```

Release Information

Command introduced before Junos OS Release 7.4.

RELATED DOCUMENTATION

[show services accounting flow](#) | 1645

[Inline Flow Monitoring for Virtual Chassis Overview](#)

show services accounting usage

IN THIS SECTION

- [Syntax](#) | 1672
- [Description](#) | 1672
- [Options](#) | 1672
- [Additional Information](#) | 1672
- [Required Privilege Level](#) | 1672
- [Output Fields](#) | 1672
- [Sample Output](#) | 1673
- [Release Information](#) | 1674

Syntax

```
show services accounting usage
<name service-name>
```

Description

Display the CPU usage of PIC used for active flow monitoring.

Options

- none** Display CPU usage for all service names.
- name *service-name*** (Optional) Display CPU usage for the specified service name.

Additional Information

When no route record has been downloaded from the PIC, this command reports no flows, even though packets are being sampled.

Required Privilege Level

view

Output Fields

[Table 158 on page 1672](#) lists the output fields for the `show services accounting usage` command. Output fields are listed in the approximate order in which they appear.

Table 158: show services accounting usage Output Fields

Output Field	Output Field Description
Service Accounting interface	Name of the service accounting interface.

Table 158: show services accounting usage Output Fields (*Continued*)

Output Field	Output Field Description
Service name	Name of a service that was configured at the [edit-forwarding-options accounting] hierarchy level. The default display, (default sampling), indicates the service was configured at the [edit-forwarding-options sampling-level] hierarchy level.
Local interface index	Index counter of the local interface.
Uptime	Time that the PIC has been operational (in milliseconds).
Interrupt time	Total time that the PIC has spent processing packets since the last PIC reset (in microseconds).
Load (5 second)	CPU load on the PIC, averaged more than 5 seconds. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.
Load (1 minute)	CPU load on the PIC, averaged more than 1 minute. The number is a percentage obtained by dividing the time spent on active tasks by the total elapsed time.

Sample Output

show services accounting usage (Monitoring PIC Interface)

```

user@host> show services accounting usage
Service Accounting interface: mo-1/1/0, Local interface index: 15
Service name: (default sampling)
CPU utilization
  Uptime: 600413856 milliseconds, Interrupt time: 2403 microseconds
  Load (5 second): 43%, Load (1 minute): 24%

```

show services accounting usage (Service PIC Interface)

```

user@host> show services accounting usage
Service Accounting interface: sp-0/1/0

```

```
Service name: (default sampling)
CPU utilization
  Uptime: 7853940 milliseconds, Interrupt time: 0 microseconds
  Load (5 second): 2%, Load (1 minute): 0%
```

```
Service Accounting interface: sp-0/1/0
Service name: (default sampling)
CPU utilization
  Uptime: 331160 milliseconds, Interrupt time: 0 microseconds
  Load (5 second): 2%, Load (1 minute): 0%
```

Release Information

Command introduced before Junos OS Release 7.4.

show services dynamic-flow-capture content-destination

IN THIS SECTION

- [Syntax | 1675](#)
- [Description | 1675](#)
- [Options | 1675](#)
- [Required Privilege Level | 1675](#)
- [Output Fields | 1675](#)
- [Sample Output | 1676](#)
- [Release Information | 1676](#)

Syntax

```
show services dynamic-flow-capture content-destination capture-group group-name destination-  
identifier identifier  
<terse>
```

Description

(M320 Series routers and T Series routers only) Display information about the content destination that receives packets from the dynamic flow capture (DFC) interface.

Options

- capture-group *group-name***Display information for the specified capture-group identifier.
- destination-identifier *identifier***Display information for the specified content destination identifier.
- terse**(Optional) Display summary information.

Required Privilege Level

view

Output Fields

[Table 159 on page 1675](#) lists the output fields for the `show services dynamic-flow-capture content-destination` command. Output fields are listed in the approximate order in which they appear.

Table 159: show services dynamic-flow-capture content-destination Output Fields

Output Field	Output Field Description
Capture group	Name of the capture group.
Content destination	Name of the content destination.
Criteria	Number of criteria specified.

Table 159: show services dynamic-flow-capture content-destination Output Fields (*Continued*)

Output Field	Output Field Description
Bandwidth	Bandwidth used by the matched traffic.
Matched packets	Number of matched packets sent to the content destination.
Matched bytes	Number of matched bytes sent to the content destination.
Congestion notifications	Number of notification messages sent.

Sample Output

show services dynamic-flow-capture content-destination capture-group

```
user@host> show services dynamic-flow-capture content-destination capture-group g1 destination-
identifier cd1 terse
  Capture group: g1, Content destination: cd1, Criteria: 0, Bandwidth: 0, Matched packets: 0,
  Matched bytes: 0, Congestion notifications: 0
```

Release Information

Command introduced in Junos OS Release 7.4.

show services dynamic-flow-capture control-source

IN THIS SECTION

- [Syntax | 1677](#)
- [Description | 1677](#)
- [Options | 1677](#)

- [Required Privilege Level | 1677](#)
- [Output Fields | 1677](#)
- [Sample Output | 1679](#)
- [Release Information | 1679](#)

Syntax

```
show services dynamic-flow-capture control-source capture-group group-name control-source
source-identifier identifier
<detail | terse>
```

Description

(M320 Series routers and T Series routers only) Display information about the control source that makes dynamic flow capture requests to the dynamic flow capture interface.

Options

capture-group <i>group-name</i>	Capture group identifier.
source-identifier <i>identifier</i>	Control source identifier.
detail terse	(Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

[Table 160 on page 1678](#) lists the output fields for the `show services dynamic-flow-capture control-source` scommand. Output fields are listed in the approximate order in which they appear.

Table 160: show services dynamic-flow-capture control-source Output Fields

Output Field	Output Field Description
Capture group	Name of the capture group.
Control source	Name of the control source.
Criteria added, Criteria add failed	Number of criteria added or added and failed.
Active criteria	Number of active criteria.
Static criteria, Dynamic criteria	Number of static or dynamic criteria.
Control protocol requests	Total number of control protocol requests.
Requests	Number of Add , Delete , List , Refresh , and No-op control protocol requests.
Failed	Number of Add , Delete , List , Refresh , and No-op failed control protocol requests.
Add request rate	Rate of add requests.
Add request peak rate	Peak rate of add requests.
Bandwidth across all criteria	Bandwidth used by all the requests.
Total notifications	Total number of notifications sent and the number of notifications by category: Restart , Rollover , Timeout , Congestion , Congestion delete , and Dups (duplicates) dropped.
Criteria deleted	Total number of criteria deleted and the number of deleted criteria by category: Timeout idle , Timeout total , Packets , and Bytes .
Sequence number	Sequence number.

Sample Output

show services dynamic-flow-capture control-source source-identifier capture-group

```
user@host> show services dynamic-flow-capture control-source source-identifier cs0_cg0 capture-
group cg_0
Capture group: cg_0, Control source: cs0_cg0
  Criteria added: 28, Criteria add failed: 0, Active criteria: 0, Control protocol requests:
28,    Add request rate: 0,
    Add request peak rate: 1, Bandwidth across all criteria: 0, Total notifications: 1, Criteria
deleted: 28,    Sequence number: 0
```

show services dynamic-flow-capture control-source ource-identifier capture-group detail

```
user@host> show services dynamic-flow-capture control-source source-identifier cs0_cg0 capture-
group cg_0 detail
Capture group: cg_0, Control source: cs0_cg0
  Criteria added: 28, Criteria add failed: 0
  Active criteria: 0
    Static criteria: 0, Dynamic criteria: 0
  Control protocol requests: 28
    Add      Delete      List      Refresh      No-op
  Requests   28         0         0           0           0
  Failed      0          0         0           0           0
  Add request rate: 0
  Add request peak rate: 1
  Bandwidth across all criteria: 0
  Total notifications: 1
    Restart: 1, Rollover: 0, No-op: 0, Timeout: 0, Congestion: 0, Congestion delete: 0, Dups
dropped: 0
  Criteria deleted: 28
    Timeout idle: 0, Timeout total: 0, Packets: 0, Bytes: 0
  Sequence number: 0
```

Release Information

Command introduced in Junos OS Release 7.4.

show services dynamic-flow-capture statistics

IN THIS SECTION

- [Syntax | 1680](#)
- [Description | 1680](#)
- [Options | 1680](#)
- [Required Privilege Level | 1680](#)
- [Output Fields | 1680](#)
- [Sample Output | 1683](#)
- [Release Information | 1684](#)

Syntax

```
show services dynamic-flow-capture statistics capture-group group-name
```

Description

(M320 Series routers and T Series routers only) Display statistics information about the capture group specified for dynamic flow capture.

Options

capture-group *group-name* Display information for the specified capture group identifier.

Required Privilege Level

view

Output Fields

[Table 161 on page 1681](#) lists the output fields for the show services dynamic-flow-capture statistics command. Output fields are listed in the approximate order in which they appear.

Table 161: show services dynamic-flow-capture statistics Output Fields

Output Field	Output Field Description
Input	<p>Incoming dynamic flow capture packet statistics:</p> <ul style="list-style-type: none"> • Control protocol packets—Number of control protocol packets received. • Captured data packets—Number of data packets captured. • Control IRI packets—Number of control IRI packets received.
Control protocol drops	<p>Control protocol packets dropped for the following reasons:</p> <ul style="list-style-type: none"> • Not IP packets—Dropped packets were not IP packets. • Not UDP packets—Dropped packets were not User Datagram Protocol (UDP) packets. • Invalid destination address—Dropped packets had invalid destination addresses. • No memory—Packets dropped because of insufficient memory. • Unauthorized control source—Packets dropped because the control source was not authenticated. • Bad request—Packets dropped because the request was invalid. • Unknown control source—Packets dropped because the control source was not known. • Not DTCP—Dropped packets did not adhere to the control protocol format. • Bad command line—Packets dropped because of a version mismatch. • Bandwidth exceeded—Packets dropped because the bandwidth was exceeded. • Drop rate due to exceeded bandwidth—Rate of traffic dropped because the bandwidth was exceeded. • Other—Packets dropped for other reasons or undetermined causes.

Table 161: show services dynamic-flow-capture statistics Output Fields (*Continued*)

Output Field	Output Field Description
Input drops	<p>Incoming dynamic flow capture packets dropped for the following reasons:</p> <ul style="list-style-type: none"> • Unknown packets—Packets dropped because the packet type was not recognized. • Captured data not IPv4—Packets dropped because they were not IPv4 packets. • Captured data too small—Packets dropped because they were smaller than the size reported in their headers. • Captured data drops—Data packets dropped because of undetermined causes. • Captured data not matched—Packets dropped because they did not match filter criteria. • Bandwidth exceeded—Packets dropped because the bandwidth was exceeded. • Drop rate due to exceeded bandwidth—Rate of traffic dropped because the bandwidth was exceeded.
Output	<p>Outgoing dynamic flow capture packet statistics:</p> <ul style="list-style-type: none"> • Control protocol packets—Number of control protocol packets sent. • Captured data packets—Number of captured data packets sent.
Output drops	<p>Outgoing packets dropped:</p> <ul style="list-style-type: none"> • Control protocol drops—Number of control protocol packets dropped. • Captured data drops—Number of captured data packets dropped.

Table 161: show services dynamic-flow-capture statistics Output Fields (*Continued*)

Output Field	Output Field Description
Flow Statistics	<p>DFC flow statistics:</p> <ul style="list-style-type: none"> • Active flow cache entries • Active flow cache usage percentage • Flow cache entries allocated • Number of control sources • Number of content destinations • Number of criteria • Maximum criteria matching one flow • Cached flows purged for memory • Maximum filters matching one packet

Sample Output

show services dynamic-flow-capture statistics capture-group

```

user@host> show services dynamic-flow-capture statistics capture-group g1
Input:

Control protocol packets: 643, Captured data packets: 69977, Control IRI packets: 337

Control protocol drops:

Not IP packets: 0, Not UDP packets: 3, Invalid destination address: 0, No memory: 0,
Unauthorized control source: 0,

Bad request: 0, Unknown control source: 0, Not DTCP: 0, Bad command line: 0, Bandwidth
exceeded: 0,

Drop rate due to exceeded bandwidth: 0, Other: 0

```

Input drops:

Unknown packets: 0, Captured data not IPv4: 0, Captured data too small: 0, Captured data drops: 0, Captured data not matched: 0,

Bandwidth exceeded: 0, Drop rate due to exceeded bandwidth: 0

Output:

Control protocol packets: 644, Captured data packets: 1119624

Output drops:

Control protocol drops: 0, Captured data drops: 0

Flow Statistics:

Active flow cache entries: 40, Active flow cache usage percentage: 0, Flow cache entries allocated: 40,

Number of control sources: 4, Number of content destinations: 64, Number of criteria: 640,

Maximum criteria matching one flow: 16, Cached flows purged for memory: 0, Maximum filters matching one packet: 16

Release Information

Command introduced in Junos OS Release 7.4.

show services flow-collector file interface**IN THIS SECTION**

- [Syntax | 1685](#)
- [Description | 1685](#)
- [Options | 1685](#)

- [Additional Information | 1685](#)
- [Required Privilege Level | 1685](#)
- [Output Fields | 1686](#)
- [Sample Output | 1687](#)
- [Release Information | 1688](#)

Syntax

```
show services flow-collector file interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display information about flow collector files.

Options

- | | |
|-------------------------------------|---|
| none | Display file information for all configured flow collector interfaces. |
| all cp-<i>fpc/pic/port</i> | Display file information for all configured flow collector interfaces or for the specified interface. |
| detail extensive terse | (Optional) Display the specified level of output. |

Additional Information

No entries are displayed for files that have been successfully transferred.

Required Privilege Level

view

Output Fields

Table 162 on page 1686 lists the output fields for the `show services flow-collector file interface` command. Output fields are listed in the approximate order in which they appear.

Table 162: show services flow-collector file interface Output Fields

Output Field	Output Field Description	Level of Output
Filename	Name of the file created on the flow collector interface.	All levels
Flows	Total number of collector flows for which records are present in the file.	none specified
Throughput	Throughput statistics: <ul style="list-style-type: none"> • Flow records—Number of flow records in the file. <ul style="list-style-type: none"> • per second—Average number of flow records per second. • peak per second—Peak number of flow records per second. • Uncompressed bytes—Total file size before compression. <ul style="list-style-type: none"> • per second—Average number of uncompressed bytes per second. • peak per second—Peak number of uncompressed bytes per second. • Compressed bytes—Total file size after compression. <ul style="list-style-type: none"> • per second—Average number of compressed bytes per second. • peak per second—Peak number of compressed bytes per second. 	extensive

Table 162: show services flow-collector file interface Output Fields *(Continued)*

Output Field	Output Field Description	Level of Output
Status	<p>File statistics:</p> <ul style="list-style-type: none"> • Compressed blocks—(extensive output only) Data blocks in the file that have been compressed. The file is exported only when the compressed block count and block count become the same. • Block count—(extensive output only) Total number of data blocks in the file. • State—Processing state of the file. <ul style="list-style-type: none"> • Active—The flow collector interface is writing to the file. • Export 1—File export is in progress to the primary server. • Export 2—File export is in progress to the secondary server. • Wait—File is pending export. • Transfer attempts 0.—Number of attempts made to transfer the file. If the file is successfully transferred in the first attempt, this field is 0. 	All levels

Sample Output

show services flow-collector file interface extensive

```

user@host> show services flow-collector file interface cp-3/2/0 extensive
Filename: cFlowd-py69Ni69-0-20031112_014301-so_3_0_0_0.bcp.bi.gz
Throughput:
  Flow records: 188365, per second: 238, peak per second: 287
  Uncompressed bytes: 21267756, per second: 27007, peak per second: 32526
  Compressed bytes: 2965643, per second: 0, peak per second: 22999
Status:
  Compressed blocks: 156, Block count: 156
  State: Active, Transfer attempts: 0

```

Release Information

Command introduced before Junos OS Release 7.4.

show services flow-collector input interface

IN THIS SECTION

- [Syntax | 1688](#)
- [Description | 1688](#)
- [Options | 1688](#)
- [Required Privilege Level | 1689](#)
- [Output Fields | 1689](#)
- [Sample Output | 1689](#)
- [Release Information | 1690](#)

Syntax

```
show services flow-collector input interface (all | cp-fpc/pic/port)  
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display the number of packets received by collector interfaces from monitoring interfaces.

Options

- | | |
|-------------------------------------|---|
| none | Display packets received by all configured flow collector interfaces. |
| all cp-<i>fpc/pic/port</i> | Display packets received by all configured flow collector interfaces or by the specified interface. |

detail | extensive | terse (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

[Table 163 on page 1689](#) lists the output fields for the `show services flow-collector input interface` command. Output fields are listed in the approximate order in which they appear.

Table 163: show services flow-collector input interface Output Fields

Output Field	Output Field Description
Interface	Name of the monitoring interface.
Packets	Number of packets traveling from the monitoring interface to the flow collector interface.
Bytes	Number of bytes traveling from the monitoring interface to the flow collector interface.

Sample Output

`show services flow-collector input interface`

```

user@host> show services flow-collector input interface cp-3/2/0
Interface                Packets    Bytes
mo-3/0/0.0              21706     32328568
mo-3/1/0.0              21706     32329096

```

`show services flow-collector input interface all`

```

user@host> show services flow-collector input interface all
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Interface                Packets    Bytes

```

```

mo-3/0/0.0          274      416232
mo-3/3/0.0          274      416184
mo-1/0/0.0          274      416232
mo-1/1/0.0          274      416232
mo-1/2/0.0          274      416232
mo-1/3/0.0          274      416232
mo-3/1/0.0          274      416232
mo-4/0/0.0          274      416232
mo-4/1/0.0          274      416232
mo-4/2/0.0          274      416184
mo-4/3/0.0          274      416232
mo-5/0/0.0          274      416232
mo-5/1/0.0          274      416232
mo-5/2/0.0          274      416232
mo-5/3/0.0          274      416232
mo-6/0/0.0          274      416232

```

Flow collector interface: cp-6/3/0

Interface state: Collecting flows

Release Information

Command introduced before Junos OS Release 7.4.

show services flow-collector interface

IN THIS SECTION

- [Syntax | 1691](#)
- [Description | 1691](#)
- [Options | 1691](#)
- [Required Privilege Level | 1691](#)
- [Output Fields | 1691](#)
- [Sample Output | 1696](#)
- [Release Information | 1701](#)

Syntax

```
show services flow-collector interface (all | cp-fpc/pic/port)
<detail | extensive | terse>
```

Description

(M40e, M160, and M320 Series routers and T Series routers only) Display overall statistics for the flow collector application.

Options

- none** Display statistics for flow collector applications on all interfaces.
- all | cp-fpc/pic/port** Display statistics for flow collector applications on all interfaces or for the specified interface.
- detail | extensive | terse** (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

[Table 164 on page 1691](#) lists the output fields for the `show services flow-collector interface` command. Output fields are listed in the approximate order in which they appear.

Table 164: show services flow-collector interface Output Fields

Output Field	Output Field Description	Level of Output
Flow collector interface	Name of the flow collector interface.	All levels
Interface state	Collecting flow state for the interface.	All levels

Table 164: show services flow-collector interface Output Fields (*Continued*)

Output Field	Output Field Description	Level of Output
Packets	Total number of packets received.	none specified
Flows Uncompressed Bytes	Total uncompressed data size for all files created on this PIC.	none specified
Compressed Bytes	Total compressed data size for all files created on this PIC.	none specified
FTP bytes	Total number of bytes transferred to the FTP server, including those dropped during transfer.	none specified
FTP files	Total number of FTP transfers attempted by the server.	none specified
Memory	Bytes used on the PIC and bytes free.	detail extensive
Input	<p>Incoming flow collector packet statistics:</p> <ul style="list-style-type: none"> • Packets—Number of packets received on the unit. <ul style="list-style-type: none"> • per second—Average number of packets per second. • peak per second—Peak number of packets per second. • Bytes—Number of bytes received on the unit. <ul style="list-style-type: none"> • per second—Average number of bytes per second. • peak per second—Peak number of bytes per second. • Flow records processed—Number of records in the flow collector packets that were processed by the flow-collector interface. <ul style="list-style-type: none"> • per second—Average number of flow records processed per second. • peak per second—Peak number of flow records per second. 	detail extensive

Table 164: show services flow-collector interface Output Fields (*Continued*)

Output Field	Output Field Description	Level of Output
Allocation	<p>Data block statistics:</p> <ul style="list-style-type: none"> • Blocks allocated—Total number of data blocks (containing flow records) allocated to the files created on this PIC. <ul style="list-style-type: none"> • per second—Average number of blocks allocated per second. • peak per second—Peak number of blocks allocated per second. • Blocks freed—Total number of data blocks freed. <ul style="list-style-type: none"> • per second—Average number of blocks freed per second. • peak per second—Peak number of blocks freed per second. • Blocks unavailable—Total number of data block requests denied, typically because of a memory shortage. <ul style="list-style-type: none"> • per second—Average number of blocks unavailable per second. • peak per second—Peak number of blocks unavailable per second. 	extensive
Files	<p>File statistics, incremented since the PIC last booted:</p> <ul style="list-style-type: none"> • Files created—Total number of files created on this PIC. • Files exported— Number of files successfully created and exported. • Files destroyed—(extensive output only) Number of files successfully exported and files dropped by the flow collection interface. 	detail extensive

Table 164: show services flow-collector interface Output Fields (*Continued*)

Output Field	Output Field Description	Level of Output
Throughput	<p>Throughput statistics:</p> <ul style="list-style-type: none"> • Uncompressed bytes—Total uncompressed data size for all files created on this PIC. • per second—Average number of uncompressed bytes per second. • peak per second—Peak number of uncompressed bytes per second. • Compressed bytes—Total compressed data size for all files created on this PIC. • per second—Average number of compressed bytes per second. • peak per second—Peak number of compressed bytes per second. 	detail extensive
Packet drops	<p>Number of packets dropped for the following causes:</p> <ul style="list-style-type: none"> • No memory—Packets dropped because of insufficient memory. • Not IP—Packets dropped because they are not IP packets. • Not IPv4—Packets dropped because they are not IP version 4 packets. • Too small—Packets dropped because each packet was smaller than the size reported in its header. • Fragments—Packets dropped because of fragmentation. Fragments are not reassembled. • ICMP—Packets dropped because they are not ICMP packets. • TCP—Packets dropped because they are not TCP packets. • Unknown—Packets dropped because of undetermined causes. • Not Junos flow—Packets dropped because they are not interpreted by Junos OS. Junos OS interprets only IPv4, UDP cflowd version 5 packets. 	extensive

Table 164: show services flow-collector interface Output Fields (*Continued*)

Output Field	Output Field Description	Level of Output
File transfer	<p>File transfer statistics:</p> <ul style="list-style-type: none"> • FTP bytes—Total number of bytes transferred to the FTP server, including those dropped during transfer. • FTP files—Total number of FTP transfers attempted by the server. • FTP failure—Total number of FTP failures encountered by the server. 	detail extensive
Flow collector interface	Physical interface acting as a flow collector.	detail

Table 164: show services flow-collector interface Output Fields (*Continued*)

Output Field	Output Field Description	Level of Output
Export channel	<p>Export channel 0 is unit 0. Export channel 1 is unit 1. Flow receive channel is unit 2. Server status statistics are the following:</p> <ul style="list-style-type: none"> • Current server Primary or Secondary—Current FTP server being used. Value is • Primary server state—State of the server: <ul style="list-style-type: none"> • OK—Server is operating without problems. • FTP error—Server encountered an FTP protocol error while sending files. • Network error—Flow-collector interface has errors when contacting the primary FTP server. • Unknown—First file transfer has not been sent to the primary server. • Secondary server state—State of the server: <ul style="list-style-type: none"> • OK—Server is operating without errors. • FTP error—Server encountered an FTP protocol error while sending files. • Network error—Flow-collector interface has errors when contacting the secondary FTP server. • Unknown—First file transfer has not been sent to the secondary server. • Not configured—Secondary server is not configured. 	detail extensive

Sample Output

show services flow-collector interface all detail

```

user@host> show services flow-collector interface all detail
Flow collector interface: cp-6/1/0
Interface state: Collecting flows

```

Memory:

Used: 51452732, Free: 440329088

Input:

Packets: 4384, per second: 0, peak per second: 156

Bytes: 6659616, per second: 0, peak per second: 249695

Flow records processed: 131070, per second: 0, peak per second: 4914

Files:

Files created: 1, per second: 0, peak per second: 0

Files exported: 1, per second: 0, peak per second: 0

Throughput:

Uncompressed bytes: 13742307, per second: 0, peak per second: 593564

Compressed bytes: 3786177, per second: 0, peak per second: 162826

File Transfer:

FTP bytes: 3786247, per second: 0, peak per second: 378620

FTP files: 1, per second: 0, peak per second: 0

FTP failure: 0

Export channel: 0

Current server: Primary

Primary server state: OK, Secondary server state: OK

Export channel: 1

Current server: Primary

Primary server state: Unknown, Secondary server state: OK

Flow collector interface: cp-6/3/0

Interface state: Collecting flows

Memory:

Used: 51452732, Free: 440329088

Input:

Packets: 0, per second: 0, peak per second: 0

Bytes: 0, per second: 0, peak per second: 0

Flow records processed: 0, per second: 0, peak per second: 0

Files:

Files created: 0, per second: 0, peak per second: 0

Files exported: 0, per second: 0, peak per second: 0

Throughput:

Uncompressed bytes: 0, per second: 0, peak per second: 0

Compressed bytes: 0, per second: 0, peak per second: 0

File Transfer:

FTP bytes: 70, per second: 0, peak per second: 6

FTP files: 0, per second: 0, peak per second: 0

FTP failure: 0

Export channel: 0

Current server: Primary

```

Primary server state: Unknown, Secondary server state: OK
Export channel: 1
Current server: Primary
Primary server state: Unknown, Secondary server state: OK

```

show services flow-collector interface all extensive

```

user@host> show services flow-collector interface all extensive
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
Memory:
  Used: 51452732, Free: 440329088
Input:
  Packets: 4384, per second: 0, peak per second: 156
  Bytes: 6659616, per second: 0, peak per second: 249695
  Flow records processed: 131070, per second: 0, peak per second: 4914
Allocation:
  Blocks allocated: 108, per second: 0, peak per second: 0
  Blocks freed: 108, per second: 0, peak per second: 10
  Blocks unavailable: 0, per second: 0, peak per second: 0
Files:
  Files created: 1, per second: 0, peak per second: 0
  Files exported: 1, per second: 0, peak per second: 0
  Files destroyed: 1, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 13742307, per second: 0, peak per second: 593564
  Compressed bytes: 3786177, per second: 0, peak per second: 162826
Packet drops:
  No memory: 0, Not IP: 0
  Not IPv4: 0, Too small: 0
  Fragments: 0, ICMP: 0
  TCP: 0, Unknown: 0
  Not JUNOS flow: 0
File Transfer:
  FTP bytes: 3786247, per second: 0, peak per second: 378620
  FTP files: 1, per second: 0, peak per second: 0
  FTP failure: 0
Export channel: 0
Current server: Primary
Primary server state: OK, Secondary server state: OK
Export channel: 1

```

Current server: Primary

Primary server state: Unknown, Secondary server state: OK

Flow collector interface: cp-6/3/0

Interface state: Collecting flows

Memory:

Used: 51452732, Free: 440329088

Input:

Packets: 0, per second: 0, peak per second: 0

Bytes: 0, per second: 0, peak per second: 0

Flow records processed: 0, per second: 0, peak per second: 0

Allocation:

Blocks allocated: 0, per second: 0, peak per second: 0

Blocks freed: 0, per second: 0, peak per second: 0

Blocks unavailable: 0, per second: 0, peak per second: 0

Files:

Files created: 0, per second: 0, peak per second: 0

Files exported: 0, per second: 0, peak per second: 0

Files destroyed: 0, per second: 0, peak per second: 0

Throughput:

Uncompressed bytes: 0, per second: 0, peak per second: 0

Compressed bytes: 0, per second: 0, peak per second: 0

Packet drops:

No memory: 0, Not IP: 0

Not IPv4: 0, Too small: 0

Fragments: 0, ICMP: 0

TCP: 0, Unknown: 0

Not JUNOS flow: 0

File Transfer:

FTP bytes: 70, per second: 0, peak per second: 6

FTP files: 0, per second: 0, peak per second: 0

FTP failure: 0

Export channel: 0

Current server: Primary

Primary server state: Unknown, Secondary server state: OK

Export channel: 1

Current server: Primary

Primary server state: Unknown, Secondary server state: OK

show services flow-collector interface all terse

```

user@host> show services flow-collector interface all terse
Flow collector interface: cp-6/1/0
Interface state: Collecting flows
  Packets    Bytes    Flows Uncompressed  Compressed  FTP bytes FTP files
           Bytes      Bytes
    4384  6659616   131070   13742307   3786177    3786247        1

Flow collector interface: cp-6/3/0
Interface state: Collecting flows
  Packets    Bytes    Flows Uncompressed  Compressed  FTP bytes FTP files
           Bytes      Bytes
         0         0         0         0         0         70         0

```

show services flow-collector interface extensive

```

user@host> show services flow-collector interface cp-5/2/0 extensive
Flow collector interface: cp-5/2/0
Interface state: Collecting flows
Memory:
  Used: 458311860, Free: 40810008
Input:
  Packets: 922629, per second: 2069, peak per second: 3266
  Bytes: 1376559252, per second: 3096940, peak per second: 4880051
  Flow records processed: 25764957, per second: 42564, peak per second: 98124
Allocation:
  Blocks allocated: 20862, per second: 31, peak per second: 72
  Blocks freed: 17161, per second: 40, peak per second: 202
  Blocks unavailable: 58786, per second: 652, peak per second: 1120
Files:
  Files created: 52, per second: 0, peak per second: 0
  Files exported: 42, per second: 0, peak per second: 0
  Files destroyed: 42, per second: 0, peak per second: 0
Throughput:
  Uncompressed bytes: 2592070401, per second: 7297307,
  peak per second: 8630023
  Compressed bytes: 659600068, per second: 1858458, peak per second: 2198471
Packet drops:
  No memory: 58786, Not IP: 0

```

```

Not IPv4: 0, Too small: 0
Fragments: 0, ICMP: 0
TCP: 0, Unknown: 0
Not JUNOS flow: 0
File Transfer:
  FTP bytes: 585981447, per second: 1313320, peak per second: 4857798
  FTP files: 48, per second: 0, peak per second: 0
  FTP failure: 8
Export channel: 0
  Current server: Primary
  Primary server state: FTP error, Secondary server state: Not configured Export channel: 1
  Current server: Primary
  Primary server state: OK, Secondary server state: Not configured

```

Release Information

Command introduced before Junos OS Release 7.4.

show services inband-flow-telemetry

IN THIS SECTION

- [Syntax | 1702](#)
- [Description | 1702](#)
- [Options | 1702](#)
- [Required Privilege Level | 1702](#)
- [Output Fields | 1702](#)
- [Sample Output | 1703](#)
- [Release Information | 1704](#)

Syntax

```
show services inband-flow-telemetry (global | profile | stats)
```

Description

Display Inband Flow Analyzer (IFA) 2.0 statistics.

Options

- global** Display information about IFA global parameters configured in the IFA node. The IFA global parameters include the global device ID, metadata stack length, and hop limit.
- profile** Display information about the IFA profile configured in the IFA node. The IFA profile configuration parameters include the profile name, sample rate, and so on.
- stats** Display information about the number of IFA probe copies that traversed the IFA initiator, transit, and terminating nodes.

Required Privilege Level

view

Output Fields

[Table 165 on page 1702](#) lists the output fields for the `show services inband-flow-telemetry` command.

Table 165: show services inband-flow-telemetry Output Fields

Field Name	Field Description	Level of Output
Global Device ID	Global device identifier, which is a unique identifier configured in the IFA node.	global
Meta-data Stack Length	Metadata stack length configured in the IFA node.	global
Hop Limit	Maximum number of hops in an IFA zone.	global

Table 165: show services inband-flow-telemetry Output Fields (Continued)

Field Name	Field Description	Level of Output
Profile Name	Name configured for the IFA profile. A profile name uniquely identifies each profile.	profile
Sample rate	Sample rate configured on the IFA node.	profile
Source Address	IFA terminating node's IPv4 source address.	profile
Destination Address	Collector's IPv4 destination address.	profile
Destination Port	Collector's destination port number.	profile
IFA Init Packets	Number of IFA probes initiated from the device.	stats
IFA Transit Packets	Number of IFA probes transited.	stats
IFA Terminate Rx Packets	Number of IFA probes received at the terminating node.	stats
IFA Terminate Tx Packets	Number of IFA copies transmitted to the collector application.	stats

Sample Output

show services inband-flow-telemetry global (QFX5120-48Y and QFX5120-32C)

```

user@host> show services inband-flow-telemetry global
Global Device ID       : 20
Meta-data Stack Length : 160
Hop Limit              : 255

```

show services inband-flow-telemetry profile (QFX5120-48Y and QFX5120-32C)

```
user@host> show services inband-flow-telemetry profile
Profile Name       : p1
Sample rate       : 1
Source Address    : 10.80.1.1
Destination Address : 10.80.1.2
Destination Port   : 2055
```

show services inband-flow-telemetry stats (QFX5120-48Y and QFX5120-32C)

```
user@host> show services inband-flow-telemetry stats
IFA Init Packets      : 10
IFA Transit Packets   : 20
IFA Terminate Rx Packets : 30
IFA Terminate Tx Packets : 30
```

Release Information

Command introduced in Junos OS Release 21.4R1.

RELATED DOCUMENTATION

[clear inband-flow-telemetry stats | 1560](#)

[inband-flow-telemetry | 1157](#)

[Inband Flow Analyzer \(IFA\) 2.0 Probe for Real-Time Flow Monitoring | 370](#)

show services inline-monitoring feature-profile-mapping fpc-slot

IN THIS SECTION

- [Syntax | 1705](#)
- [Description | 1705](#)

- Options | 1705
- Required Privilege Level | 1705
- Output Fields | 1705
- Sample Output | 1706
- Release Information | 1707

Syntax

```
show services inline-monitoring feature-profile-mapping fpc-slot slot-number
```

Description

(EX4100 and EX4100-F switches only) Show what features you configured for flow-based telemetry and how they are ordered in the flow.

Options

fpc-slot *slot-number* Specify the slot number to view the feature mapping for that Flexible PIC Concentrator (FPC). You must have already configured a feature profile for inline monitoring for this command to contain any information.

Required Privilege Level

view

Output Fields

[Table 166 on page 1706](#) lists the output fields for the `show services inline-monitoring feature-profile-mapping` command.

Table 166: show services inline-monitoring feature-profile-mapping Output Fields

Field Name	Field Description
FPC Slot	Slot configured for flow-based telemetry.
Flow counter Instance Name	Name of the inline-monitoring instance you configured for flow-based telemetry.
EnterpriseElemId	Information Element number for a feature you configured for the inline-monitoring feature profile.
Feature/Counter Name	Name of a feature you configured in an inline-monitoring feature profile.

Sample Output

show services inline-monitoring feature-profile-mapping fpc-slot 0

```
user@host> show services inline-monitoring feature-profile-mapping fpc-slot 0
```

```
FPC Slot: 0, Flow counter Instance Name: i1
```

```
-----+
EnterpriseElemId | Feature/Counter Name
-----+
34               Inter-arrival-time(4bytes)
35               Inter-departure-delay (4bytes)
36               Chip-delay (4 bytes)
37               shared-pool-congestion-level(1 byte)
38               Queue-congestion-level (1 byte)
16               Dos-attack( 4 bytes)
40               Ingress-drop-reason(2 bytes)
41               Egress-drop-reason(2 bytes)
42               aggregate-intf-member-id(1 byte)
```

Release Information

Command introduced in Junos OS Release 22.2R1.

RELATED DOCUMENTATION

[Flow-Based Telemetry \(EX4100, EX4100-F, and EX4400 Series\) | 349](#)

show services inline-monitoring statistics fpc-slot

IN THIS SECTION

- [Syntax | 1707](#)
- [Description | 1707](#)
- [Options | 1708](#)
- [Required Privilege Level | 1708](#)
- [Output Fields | 1708](#)
- [Sample Output | 1709](#)
- [Release Information | 1709](#)

Syntax

```
show services inline-monitoring statistics fpc-slot fpc-slot
```

Description

Display inline-monitoring services statistical information, including for features that use inline-monitoring services, such as flow-based telemetry.

Options

fpc-slot Display inline-monitoring statistical information for the specified Flexible PIC Concentrator (FPC) number.

- **Range:** 0 through 11

Required Privilege Level

view

Output Fields

Table 167 on page 1708 lists the output fields for the `show services inline-monitoring statistics fpc-slot` command. Output fields are listed in the approximate order in which they appear.

Table 167: `show services inline-monitoring statistics fpc-slot` Output Fields

Field Name	Field Description
FPC Slot	Flexible PIC Concentrator (FPC) slot level counters: <ul style="list-style-type: none">• Packets—Number of packets serviced under FPC slot.• Bytes—Number of bytes serviced under FPC slot.
Instance Name	Instance level counters: <ul style="list-style-type: none">• Packets—Number of packets serviced under instance number under the FPC slot.• Bytes—Number of bytes serviced under instance number under the FPC slot.
Collector Name	Collector level counters: <ul style="list-style-type: none">• Packets—Number of packets serviced under collector number of an instance under the FPC slot.• Bytes—Number of bytes serviced under collector number of an instance under the FPC slot.

Sample Output

show services inline-monitoring services fpc-slot

```

user@host> show services inline-monitoring services fpc-slot fpc-slot
IMON Statistics
  FPC Slot      : <FPC_SLOT>
  Packets       : <F_P>      Bytes : <F_B>

  Instance Name : <I1>
  Packets       : <I1_P>      Bytes : <I1_B>

    Collector Name : <I1_C1>
    Packets       : <I1_C1_P> Bytes : <I1_C1_B>

    Collector Name : <I1_C2>
    Packets       : <I1_C2_P> Bytes : <I1_C2_B>

  Instance Name : <I2>
  Packets       : <I2_P>      Bytes : <I2_B>

    Collector Name : <I2_C1>
    Packets       : <I2_C1_P> Bytes : <I2_C1_B>

```

show services inline-monitoring services fpc-slot (flow-based telemetry)

```

user@host> show services inline-monitoring statistics fpc-slot 0
IMON Statistics
  FPC Slot: 0
    Instance Name: i1
    Flow exceeded count: 0
    Flow missed count: 0
    Flow aged out count: 1
    Flow learnt count: 3
    Flow meter exceeded count: 0

```

Release Information

Command introduced in Junos OS Release 19.4R1.

Command introduced in Junos OS Evolved Release 22.2R1.

RELATED DOCUMENTATION

[Understanding Inline Monitoring Services | 334](#)

[Flow-Based Telemetry \(EX4100, EX4100-F, and EX4400 Series\) | 349](#)

show services monitoring rfc2544

IN THIS SECTION

- [Syntax | 1710](#)
- [Description | 1710](#)
- [Options | 1711](#)
- [Additional Information | 1711](#)
- [Required Privilege Level | 1711](#)
- [Output Fields | 1711](#)
- [Sample Output | 1712](#)
- [Release Information | 1715](#)

Syntax

```
show services monitoring rfc2544
(active-tests <extensive>| terminated-tests <extensive> | summary)
```

Description

Display information about the results of each category or state of the RFC 2544-based benchmarking tests, such as terminated tests and active tests. You can also display summary statistics about the total number of tests of each state for a high-level, quick analysis. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

Options

active-tests <extensive>	Display the results of the set of tests that are currently running.
terminated-tests <extensive>	Display the list of tests that were terminated or stopped. This list includes tests that failed due to various error conditions and tests that you terminated by entering the test services monitoring rfc2544 test <i>test-name</i> stop command. The Status field in the output specifies the reason for the termination of the test.
summary	Display summary output.

Additional Information

Required Privilege Level

view

Output Fields

[Table 168 on page 1711](#) lists the output fields for the show services monitoring rfc2544 command.

Table 168: show services monitoring rfc2544 Output Fields

Field Name	Field Description
Counters last cleared	Date, time, and how long ago the statistics for the test were cleared. The format is <i>year-month-day hour:minute:second timezone</i> . For example, 2021-02-11 07:51:28 PDT. If you did not clear the statistics previously at any point, a dash (-) is displayed.
Status	Indicates whether the test is currently in progress or has been terminated. This field is displayed for tests that are in progress or were terminated by entering the test services monitoring rfc2544 test <i><test-name / test-id></i> stop command.
Test family	The family type configured for the test; for example, INET.
Test id	Unique identifier configured for the test.

Table 168: show services monitoring rfc2544 Output Fields (*Continued*)

Field Name	Field Description
Test mode	Mode configured for the test on the router. Test mode is: <ul style="list-style-type: none"> • Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4.
Test name	Name configured for the test.
Test start time and Test finish time	Time at which the test started and finished in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). A dash (-) is displayed if the test is still running.
Test state	State of the test that is in progress or active when the output is displayed.

Sample Output

show services monitoring rfc2544 summary

```
user@host> show services monitoring rfc2544 summary
```

```
Tests summary :
  Number of active tests: 1, Number of terminated tests: 2
```

This output indicates that 1 test iteration is currently in progress (at the time of issue of the command) and 2 tests were halted.

show services monitoring rfc2544 active-tests

```
user@host> show services monitoring rfc2544 active-tests
```

```
Active tests:

Test id: 3, Test name: mytest
```

```

Test mode: Reflect
Test family: INET
Test state: Active
Status: Running
Test start time: 2021-02-11 04:39:10 PST
Test finish time: -
Counters last cleared: 2021-02-11 04:39:12 PST

```

show services monitoring rfc2544 active-tests extensive

```
user@host> show services monitoring rfc2544 active-tests extensive
```

Active tests:

Test id: 3, Test name: mytest

```

Test mode: Reflect
Test family: INET
Test state: Active
Status: Running
Test start time: 2021-02-11 04:39:10 PST
Test finish time: -
Counters last cleared: 2021-02-11 04:39:12 PST

```

INET family Configuration:

```

Destination IPv4 address: 192.0.2.2
Destination UDP port: 7890
Source IPv4 address: 192.0.2.1

```

Elapsed time	Reflected Packets	Reflected Bytes
271	0	0

show services monitoring rfc2544 terminated-tests

```
user@host> show services monitoring rfc2544 terminated-tests
```

Terminated tests:

Test id: 1, Test name: mytest

Test mode: Reflect

Test family: INET

Test state: Terminated

Status: Stopped from CLI

Test start time: 2021-02-11 03:14:19 PST

Test finish time: 2021-02-11 03:17:26 PST

Counters last cleared: -

show services monitoring rfc2544 terminated-tests extensive

```
user@host> show services monitoring rfc2544 terminated-tests extensive
```

Terminated tests:

Test id: 1, Test name: mytest

Test mode: Reflect

Test family: INET

Test state: Terminated

Status: Stopped from CLI

Test start time: 2021-02-11 03:14:19 PST

Test finish time: 2021-02-11 03:17:26 PST

Counters last cleared: -

INET family Configuration:

Destination IPv4 address: 192.0.2.2

Destination UDP port: 7890

Source IPv4 address: 192.0.2.1

Elapsed	Reflected	Reflected
---------	-----------	-----------

time	Packets	Bytes
187	0	0

Release Information

Command introduced in Junos OS Evolved Release 21.1R1.

show services monitoring rfc2544

IN THIS SECTION

- [Syntax | 1715](#)
- [Description | 1716](#)
- [Options | 1716](#)
- [Additional Information | 1716](#)
- [Required Privilege Level | 1716](#)
- [Output Fields | 1716](#)
- [Sample Output | 1717](#)
- [Release Information | 1719](#)

Syntax

```
show services monitoring rfc2544
<active-tests>
<completed-tests>
<detail>
<extensive>
<summary>
<terminated-tests>
<test-id test-id>
```

Description

Display information about the results of the RFC 2544-based benchmarking test for a specific test ID or for a category of tests. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

Options

- active-tests** (Optional) Display results for all active RFC 2544 benchmarking tests.
- completed-tests** (Optional) Display results for all completed RFC 2544 benchmarking tests.
- detail** (Optional) When specified with the `test-id` option, display detailed results. When specified on its own, display detailed results of all RFC 2544 benchmarking tests (active, completed, and terminated).
- extensive** (Optional) When specified with the `test-id` option, display extensively detailed results. When specified on its own, displays extensive details of all RFC 2544 benchmarking tests (active, completed, and terminated).
- summary** (Optional) When specified with the `test-id` option, display summarized results. When specified on its own, display a summary of all RFC 2544 benchmarking tests (active, completed, and terminated).
- terminated-tests** (Optional) Display results for all terminated RFC 2544 benchmarking tests.
- test-id *test-id*** Display test results for the specified unique identifier.

Additional Information

Required Privilege Level

view

Output Fields

[Table 168 on page 1711](#) lists the output fields for the `show services monitoring rfc2544 test-id` command.

Table 169: show services monitoring rfc2544 test-id Output Fields

Field Name	Field Description
Counters last cleared	Date, time, and how long ago the statistics for the test were cleared. The format is <i>year-month-day hour:minute:second timezone</i> . For example, 2021-02-11 07:51:28 PDT. If you did not clear the statistics previously at any point, a dash (-) is displayed.
Status	Indicates whether the test is currently in progress or has been terminated. This field is displayed for tests that are in progress or were terminated by entering the test services monitoring rfc2544 test <test-name / test-id> stop command.
Test family	The family type configured for the test; for example, INET.
Test id	Unique identifier configured for the test.
Test mode	Mode configured for the test on the router. Test mode is: <ul style="list-style-type: none"> • Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4.
Test name	Name configured for the test.
Test start time and Test finish time	Time at which the test started and finished in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS). A dash (-) is displayed if the test is still running.
Test state	State of the test that is in progress or active when the output is displayed.

Sample Output

show services monitoring rfc2544 test-id *test-id*

```
user@host> show services monitoring rfc2544 test-id 2
```

```
Test id: 2, Test name: mytest
  Test mode: Reflect
  Test family: INET
```



```

Test state: Terminated
Status: Stopped from CLI
Test start time: 2021-02-11 03:38:25 PST
Test finish time: 2021-02-11 03:38:43 PST
Counters last cleared: -

```

show services monitoring rfc2544 test-id *test-id* detail

```
user@host> show services monitoring rfc2544 test-id 1 detail
```

```

Test id: 1, Test name: t1
  Test mode: reflect
  Test family: ethernet-switching
  Test state: Active
  Status: Running
  Test start time: 2022-05-30 13:23:21 PDT
  Test finish time: -
  Counters last cleared: -

```

Bridge family Configuration:

```

  Interface: et-0/0/3.0
  Direction: egress
  Service type: elan
  Destination MAC address: 00:00:00:08:00:08
  Source MAC address: 00:00:00:07:00:07

```

Elapsed time	Reflected Packets	Reflected Bytes
1016	5196751848	2868607020096

show services monitoring rfc2544 test-id *test-id* extensive

```
user@host> show services monitoring rfc2544 test-id 2 extensive
```

```

Test id: 2, Test name: mytest
  Test mode: Reflect
  Test family: INET
  Test state: Terminated
  Status: Stopped from CLI

```

```
Test start time: 2021-02-11 03:38:25 PST
Test finish time: 2021-02-11 03:38:43 PST
Counters last cleared: -
```

INET family Configuration:

```
Destination IPv4 address: 192.0.2.4
Destination UDP port: 7890
Source IPv4 address: 192.0.2.3
```

Elapsed time	Reflected Packets	Reflected Bytes
18	0	0

show services monitoring rfc2544 summary

```
user@host> show services monitoring rfc2544 summary
[Number of active tests: active-test-count],
[Number of completed tests: completed-test-count],
[Number of terminated tests: terminated-test-count]
```

Release Information

Command introduced in Junos OS Evolved Release 21.1R1.

active-tests, *completed-tests*, *detail*, *summary*, and *terminated-tests* options introduced in Junos OS Evolved Release 22.4R1.

show services monitoring rpm history-results

IN THIS SECTION

- [Syntax | 1720](#)
- [Description | 1720](#)
- [Options | 1720](#)
- [Required Privilege Level | 1721](#)

- [Output Fields | 1721](#)
- [Sample Output | 1723](#)
- [Release Information | 1727](#)

Syntax

```
show services monitoring rpm history-results
owner name
<brief | detail>
<since time>
<source-address address>
<target address>
<test name>
```

Description

Display the results stored for the specified real-time performance monitoring (RPM) probes.

Options

owner <i>name</i>	(Required) Display information for probes with the specified probe owner.
brief detail	(Optional) Display the specified level of output. <ul style="list-style-type: none"> • Default: brief
since <i>time</i>	(Optional) Display information from the specified time. Specify time as <i>yyyy-mm-dd.hh:mm:ss</i> .
source-address <i>address</i>	(Optional) Display information only for probes with the specified source address.
target <i>address</i>	(Optional) Display information only for probes with the specified target address.
test <i>name</i>	(Optional) Display information only for the specified test. <p>Do not configure test if you configure any of the following options: source-address or target. These options do not work when you configure test.</p>

Required Privilege Level

view

Output Fields

[Table 170 on page 1721](#) lists the output fields for the `show services monitoring rpm history-results` command. Output fields are listed in the approximate order in which they appear.

Table 170: show services monitoring rpm history-results Output Fields

Field Name	Field Description	Level of Output
Owner	Owner name. When you configure the probe owner statement at the <code>[edit services monitoring rpm]</code> hierarchy level, this field displays the configured owner name.	All levels
Test	Name of a test representing a collection of probes. When you configure the test <i>name</i> statement at the <code>[edit services monitoring rpm owner]</code> hierarchy level, the field displays the configured test name.	All levels
Probe sent	Timestamp when the probe was sent.	owner, brief
Probe received	Timestamp when the probe result was determined.	owner, brief
Round trip time	Average ping round-trip time (RTT), in microseconds.	owner, brief
Probe-type	Configured probe type for the test.	detail

Table 170: show services monitoring rpm history-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Probe results	<p>Result of a particular probe performed by a remote host, including whether a probe response was received. The following information is contained in the results:</p> <ul style="list-style-type: none"> • Probe sent time—Timestamp for when the probe was sent. • Probe rcvd time—Timestamp for when the probe was sent, and whether offload timestamping is configured. • Rtt—Average ping round-trip time (RTT), in microseconds. • Rtt jitter—Difference, in microseconds, between the maximum and minimum RTT. • Egress jitter—For the probe type icmp-ping-timestamp, the egress delay measured for this probe, in microseconds. • Ingress jitter—For the probe type icmp-ping-timestamp, the ingress delay measured for this probe, in microseconds. 	detail
Results over current test	Displays the results for the current test by probe at the time each probe was completed.	detail
Probes sent	Number of probes sent with the current test.	detail
Probes received	Number of probe responses received within the current test.	detail
Loss percentage	Percentage of lost probes for the current test.	detail

Table 170: show services monitoring rpm history-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Measurement	Measurement for round trip time: <ul style="list-style-type: none"> • Samples—Number of samples. • Minimum—Minimum RTT measured over the course of the current test. • Maximum—Maximum RTT measured over the course of the current test. • Average—Average RTT measured over the course of the current test. • Stddev—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test. 	detail

Sample Output

show services monitoring rpm history-results owner

```
user@host> show services monitoring rpm history-results owner icmp-evo
```

Owner, Test	Probe sent	Probe received	Round trip time
icmp-evo, icmp-evo-2	11/18/20 12:08:28.621987	11/18/20 12:08:28.648443	26444 usec
icmp-evo, icmp-evo-4	11/18/20 12:08:28.642527	11/18/20 12:08:28.660700	18171 usec
icmp-evo, icmp-evo-1	11/18/20 12:08:28.642533	11/18/20 12:08:28.663284	20749 usec
icmp-evo, icmp-evo-3	11/18/20 12:08:28.653151	11/18/20 12:08:28.669098	15945 usec
icmp-evo, icmp-evo-2	11/18/20 12:08:33.616890	11/18/20 12:08:33.628069	11176 usec
icmp-evo, icmp-evo-1	11/18/20 12:08:33.637811	11/18/20 12:08:33.656368	18556 usec
icmp-evo, icmp-evo-4	11/18/20 12:08:33.637819	11/18/20 12:08:33.652862	15039 usec
icmp-evo, icmp-evo-3	11/18/20 12:08:33.648109	11/18/20 12:08:33.662094	13983 usec
icmp-evo, icmp-evo-2	11/18/20 12:09:33.624341	11/18/20 12:09:33.636556	12205 usec
icmp-evo, icmp-evo-3	11/18/20 12:09:33.655442	11/18/20 12:09:33.681148	25704 usec
icmp-evo, icmp-evo-1	11/18/20 12:09:33.655454	11/18/20 12:09:33.690099	34644 usec
icmp-evo, icmp-evo-4	11/18/20 12:09:33.655456	11/18/20 12:09:33.679532	24073 usec
icmp-evo, icmp-evo-2	11/18/20 12:09:38.624449	11/18/20 12:09:38.636945	12493 usec

icmp-evo, icmp-evo-4	11/18/20 12:09:38.655661	11/18/20 12:09:38.675016	19351 usec
icmp-evo, icmp-evo-1	11/18/20 12:09:38.655690	11/18/20 12:09:38.676942	21250 usec
icmp-evo, icmp-evo-3	11/18/20 12:09:38.656034	11/18/20 12:09:38.682802	26766 usec
icmp-evo, icmp-evo-2	11/18/20 12:10:38.634733	11/18/20 12:10:38.646682	11929 usec
icmp-evo, icmp-evo-3	11/18/20 12:10:38.675453	11/18/20 12:10:38.695945	20489 usec
icmp-evo, icmp-evo-1	11/18/20 12:10:38.675859	11/18/20 12:10:38.700085	24225 usec
icmp-evo, icmp-evo-4	11/18/20 12:10:38.675945	11/18/20 12:10:38.701973	26026 usec
icmp-evo, icmp-evo-2	11/18/20 12:10:43.641129	11/18/20 12:10:43.657255	16108 usec
icmp-evo, icmp-evo-4	11/18/20 12:10:43.682202	11/18/20 12:10:43.701403	19190 usec
icmp-evo, icmp-evo-1	11/18/20 12:10:43.682684	11/18/20 12:10:43.704907	22220 usec
icmp-evo, icmp-evo-3	11/18/20 12:10:43.682981	11/18/20 12:10:43.710825	27843 usec
icmp-evo, icmp-evo-2	11/18/20 12:11:43.658117	11/18/20 12:11:43.680440	22312 usec
icmp-evo, icmp-evo-4	11/18/20 12:11:43.699092	11/18/20 12:11:43.716991	17896 usec
icmp-evo, icmp-evo-3	11/18/20 12:11:43.709647	11/18/20 12:11:43.733250	23601 usec
icmp-evo, icmp-evo-1	11/18/20 12:11:43.709653	11/18/20 12:11:43.737341	27687 usec
icmp-evo, icmp-evo-2	11/18/20 12:11:48.656468	11/18/20 12:11:48.670180	13708 usec
icmp-evo, icmp-evo-4	11/18/20 12:11:48.697172	11/18/20 12:11:48.708817	11642 usec
icmp-evo, icmp-evo-1	11/18/20 12:11:48.707382	11/18/20 12:11:48.729500	22116 usec
icmp-evo, icmp-evo-3	11/18/20 12:11:48.707621	11/18/20 12:11:48.735767	28144 usec
icmp-evo, icmp-evo-2	11/18/20 12:12:48.669000	11/18/20 12:12:48.688289	19279 usec
icmp-evo, icmp-evo-4	11/18/20 12:12:48.710047	11/18/20 12:12:48.721077	11027 usec
icmp-evo, icmp-evo-1	11/18/20 12:12:48.731012	11/18/20 12:12:48.749973	18959 usec

... (output truncated)

show services monitoring rpm history-results owner detail

```
user@host> show services monitoring rpm history-results owner icmp-junos detail
```

Owner: icmp-junos, Test: icmp-junos-2, Probe type: icmp-ping

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:40.911234

Probe rcvd time: 11/18/20 12:12:40.918908, Client offload timestamping

Rtt: 7674 usec, Rtt jitter: 0 usec

Results over current test:

Probes sent: 1, Probes received: 1, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 1, Minimum: 7674, Maximum: 7674, Average: 7674, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-3, Probe type: icmp-ping-timestamp

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:41.821260

Probe rcvd time: 11/18/20 12:12:41.832135, Client offload timestamping

Rtt: 10875 usec, Rtt jitter: 0 usec

Egress jitter: 0 usec, Ingress jitter: 0 usec

Results over current test:

Probes sent: 1, Probes received: 1, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 1, Minimum: 10875, Maximum: 10875, Average: 10875, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-4, Probe type: icmp-ping-timestamp

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:41.821327

Probe rcvd time: 11/18/20 12:12:41.831576, Client offload timestamping

Rtt: 10249 usec, Rtt jitter: 0 usec

Egress jitter: 0 usec, Ingress jitter: 0 usec

Results over current test:

Probes sent: 1, Probes received: 1, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 1, Minimum: 10249, Maximum: 10249, Average: 10249, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-1, Probe type: icmp-ping

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:45.167546

Probe rcvd time: 11/18/20 12:12:45.174508, Client offload timestamping

Rtt: 6962 usec, Rtt jitter: 0 usec

Results over current test:

Probes sent: 1, Probes received: 1, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 1, Minimum: 6962, Maximum: 6962, Average: 6962, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-2, Probe type: icmp-ping

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:45.906709

Probe rcvd time: 11/18/20 12:12:45.912594, Client offload timestamping

Rtt: 5885 usec, Rtt jitter: 1789 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 5885, Maximum: 7674, Average: 6779, Stddev: 898

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 1789, Maximum: 1789, Average: 1789, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-4, Probe type: icmp-ping-timestamp

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:46.815590

Probe rcvd time: 11/18/20 12:12:46.823159, Client offload timestamping

Rtt: 7569 usec, Rtt jitter: 2680 usec

Egress jitter: 2263 usec, Ingress jitter: 417 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 7569, Maximum: 10249, Average: 8909, Stddev: 1340

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 2680, Maximum: 2680, Average: 2680, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-3, Probe type: icmp-ping-timestamp

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:46.815597

Probe rcvd time: 11/18/20 12:12:46.823684, Client offload timestamping

Rtt: 8087 usec, Rtt jitter: 2788 usec

Egress jitter: 2337 usec, Ingress jitter: 451 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 8087, Maximum: 10875, Average: 9481, Stddev: 1394

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 2788, Maximum: 2788, Average: 2788, Stddev: 0

Owner: icmp-junos, Test: icmp-junos-1, Probe type: icmp-ping

Probe results:

Probe response received

Probe sent time: 11/18/20 12:12:50.170153

Probe rcvd time: 11/18/20 12:12:50.177224, Client offload timestamping

Rtt: 7071 usec, Rtt jitter: 109 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 6962, Maximum: 7071, Average: 7016, Stddev: 99

Measurement: Round trip jitter (usec)

```
Samples: 2, Minimum: 109, Maximum: 109, Average: 109, Stddev: 0
... (output truncated)
```

Release Information

Command introduced in Junos OS Evolved Release 20.1R1.

show services monitoring rpm probe-results

IN THIS SECTION

- [Syntax | 1727](#)
- [Description | 1727](#)
- [Options | 1728](#)
- [Required Privilege Level | 1728](#)
- [Output Fields | 1728](#)
- [Sample Output | 1734](#)
- [Release Information | 1740](#)

Syntax

```
show services monitoring rpm probe-results
<owner name>
<source-address address>
<status (FAIL | PASS) >
<target address>
<test name>
```

Description

Display the results of the most recent real-time performance monitoring (RPM) probes.

Options

- owner *name*

(Optional) Display information only for probes with the specified probe owner name.
owner is required if you want to specify the test option.
- source-address
address

(Optional) Display information only for probes with the specified source address.
- status

(Optional) Display information only for probes with the specified type of test result.
Specify one of the following:

FAIL

Failed tests

PASS

Passed tests
- target *address*

(Optional) Display information only for probes with the specified target address.
- test *name*

(Optional) Display information only for the specified test. You must also specify the
owner option.

Required Privilege Level

view

Output Fields

Table 171 on page 1728 lists the output fields for the show services monitoring rpm probe-results command. Output fields are listed in the approximate order in which they appear.

Table 171: show services rpm probe-results Output Fields

Field Name	Field Description	Level of Output
Owner	Owner name. When you configure the probe owner statement at the [edit services monitoring rpm] hierarchy level, this field displays the configured owner name.	owner, source-address, target, test

Table 171: show services rpm probe-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Test	Name of a test representing a collection of probes. When you configure the test <i>name</i> statement at the [edit services monitoring rpm owner] hierarchy level, the field displays the configured test name.	All levels
Target address	Destination IPv4 address used for the probes. This field is displayed when the probes are sent to the configured IPv4 or IPv6 targets or RPM servers.	owner, source-address, target, test
Source address	Source address used for the probes.	owner, source-address, target, test
Probe type	Protocol configured on the receiving probe server: icmp-ping, icmp-ping-timestamp, udp-ping, udp-ping-timestamp.	owner, source-address, target, test
Test size	Number of probes within a test.	owner, source-address, target, test,

Table 171: show services rpm probe-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Probe results	<p>Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:</p> <ul style="list-style-type: none"> • Probe response received • Probe sent time—Timestamp when the probe results were sent. • Probe rcvd time—Timestamp when the probe results were received. • Client and server hardware timestamps or Client offload timestamping—If timestamps are configured, an entry appears at this point. • Rtt—Average ping round-trip time (RTT), in microseconds. • Round trip jitter—Round-trip jitter, in microseconds. • Egress jitter—Egress delay, in microseconds. • Ingress jitter—Ingress delay, in microseconds. 	owner, source-address, target, test

Table 171: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Results over current test	<p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent within the current test. • Probes received—Number of probe responses received within the current test. • Loss percentage—Percentage of lost probes for the current test. • Measurement—Measurement type: round-trip time or round-trip jitter <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT measured over the course of the current test. • Maximum—Maximum RTT measured over the course of the current test. • Average—Average RTT measured over the course of the current test. • Stddev—Standard deviation, in microseconds. 	owner, source-address, target, test

Table 171: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Results over last test	<p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent for the most recently completed test. Probes received—Number of probe responses received for the most recently completed test. Loss percentage—Percentage of lost probes for the most recently completed test. Measurement—Measurement type: round-trip time or round-trip jitter <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT measured for the most recently completed test. Maximum—Maximum RTT measured for the most recently completed test. Average—Average RTT for the most recently completed test. Stddev—Standard deviation, in microseconds. 	owner, source-address, target, test

Table 171: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Results over all tests	<p>Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent in all tests. Probes received—Number of probe responses received in all tests. Loss percentage—Percentage of lost probes in all tests. Measurement—Measurement type: round-trip time or round-trip jitter <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT measured over the course of the current test. Maximum—Maximum RTT measured over the course of the current test. Average—Average RTT measured over the course of the current test. Stddev—Standard deviation, in microseconds. 	owner, source-address, target, test
Error Stats	<p>Displays error statistics for each probe.</p> <ul style="list-style-type: none"> Invalid client recv timestamp—Number of client receive timestamp less than client send timestamp. Invalid server send timestamp—Number of server send timestamp less than server receive timestamp. Invalid server processing time—Number of server side spent time greater than RTT. <p>NOTE: Error Stats is displayed in the output only if non-zero statistics exists.</p>	owner, source-address, target, test

Sample Output

show services monitoring rpm probe-results

```
user@host> show services monitoring rpm probe-results
```

```
Owner: icmp-evo, Test: icmp-evo-1
```

```
Target address: 10.0.1.2, Probe type: icmp-ping, Test size: 2
```

```
Probe results:
```

```
Probe response received
```

```
Probe sent time: 11/18/20 12:39:59.170864
```

```
Probe rcvd time: 11/18/20 12:39:59.180558, Client and server offload timestamping
```

```
Rtt: 9691 usec, Rtt jitter: 7501 usec
```

```
Egress jitter: 5677 usec, Ingress jitter: 1824 usec
```

```
Results over current test:
```

```
Probes sent: 2, Probes received: 2, Loss percentage: 0
```

```
Measurement: Round trip time (usec)
```

```
Samples: 2, Minimum: 9691, Maximum: 17192, Average: 13441, Stddev: 3752
```

```
Measurement: Round trip jitter (usec)
```

```
Samples: 2, Minimum: 7501, Maximum: 7501, Average: 7501, Stddev: 0
```

```
Results over last test:
```

```
Probes sent: 2, Probes received: 2, Loss percentage: 0
```

```
Measurement: Round trip time (usec)
```

```
Samples: 2, Minimum: 12453, Maximum: 12812, Average: 12632, Stddev: 211
```

```
Measurement: Round trip jitter (usec)
```

```
Samples: 2, Minimum: 359, Maximum: 359, Average: 359, Stddev: 0
```

```
Results over all tests:
```

```
Probes sent: 438, Probes received: 434, Loss percentage: 0.91
```

```
Measurement: Round trip time (usec)
```

```
Samples: 434, Minimum: 8200, Maximum: 489180, Average: 16060, Stddev: 23615
```

```
Measurement: Round trip jitter (usec)
```

```
Samples: 434, Minimum: 23, Maximum: 470888, Average: 5896, Stddev: 32214
```

```
Owner: icmp-evo, Test: icmp-evo-2
```

```
Target address: 192.168.33.33, Source address: 192.168.11.11, Probe type: icmp-ping, Test size: 2
```

```
Probe results:
```

```
Probe response received
```

```
Probe sent time: 11/18/20 12:39:59.086984
```

```
Probe rcvd time: 11/18/20 12:39:59.102336, Client and server offload timestamping
```

```
Rtt: 15349 usec, Rtt jitter: 13957 usec
```

```
Egress jitter: 12921 usec, Ingress jitter: 1036 usec
```

```
Results over current test:
```

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 15349, Maximum: 29306, Average: 22327, Stddev: 6980

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 13957, Maximum: 13957, Average: 13957, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 17559, Maximum: 19615, Average: 18587, Stddev: 1028

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 2056, Maximum: 2056, Average: 2056, Stddev: 0

Results over all tests:

Probes sent: 438, Probes received: 434, Loss percentage: 0.91

Measurement: Round trip time (usec)

Samples: 434, Minimum: 9006, Maximum: 502588, Average: 17743, Stddev: 24170

Measurement: Round trip jitter (usec)

Samples: 434, Minimum: 11, Maximum: 488815, Average: 7183, Stddev: 33038

Owner: icmp-evo, Test: icmp-evo-3

Target address: 10.0.1.2, Probe type: icmp-ping-timestamp, Test size: 2

Probe results:

Probe response received

Probe sent time: 11/18/20 12:39:59.181526

Probe rcvd time: 11/18/20 12:39:59.194873, Client and server offload timestamping

Rtt: 13344 usec, Rtt jitter: 7091 usec

Egress jitter: 8475 usec, Ingress jitter: 1384 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 13344, Maximum: 20435, Average: 16889, Stddev: 3547

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 7091, Maximum: 7091, Average: 7091, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 12394, Maximum: 17642, Average: 15018, Stddev: 2624

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 5248, Maximum: 5248, Average: 5248, Stddev: 0

Results over all tests:

Probes sent: 438, Probes received: 434, Loss percentage: 0.91

Measurement: Round trip time (usec)

Samples: 434, Minimum: 8268, Maximum: 445552, Average: 16728, Stddev: 21222

Measurement: Round trip jitter (usec)

```

Samples: 434, Minimum: 1, Maximum: 431893, Average: 5545, Stddev: 29251

Owner: icmp-evo, Test: icmp-evo-4
Target address: 10.0.33.33, Source address: 10.0.11.11, Probe type: icmp-ping-timestamp, Test
size: 2
Probe results:
  Probe response received
  Probe sent time: 11/18/20 12:39:59.097193
  Probe rcvd time: 11/18/20 12:39:59.105891, Client and server offload timestamping
  Rtt: 8696 usec, Rtt jitter: 13120 usec
  Egress jitter: 9579 usec, Ingress jitter: 3541 usec
Results over current test:
  Probes sent: 2, Probes received: 2, Loss percentage: 0
  Measurement: Round trip time (usec)
    Samples: 2, Minimum: 8696, Maximum: 21816, Average: 15256, Stddev: 6560
  Measurement: Round trip jitter (usec)
    Samples: 2, Minimum: 13120, Maximum: 13120, Average: 13120, Stddev: 0
Results over last test:
  Probes sent: 2, Probes received: 2, Loss percentage: 0
  Measurement: Round trip time (usec)
    Samples: 2, Minimum: 10625, Maximum: 16523, Average: 13574, Stddev: 2949
  Measurement: Round trip jitter (usec)
    Samples: 2, Minimum: 5898, Maximum: 5898, Average: 5898, Stddev: 0
Results over all tests:
  Probes sent: 438, Probes received: 434, Loss percentage: 0.91
  Measurement: Round trip time (usec)
    Samples: 434, Minimum: 8494, Maximum: 490791, Average: 16645, Stddev: 23635
  Measurement: Round trip jitter (usec)
    Samples: 434, Minimum: 8, Maximum: 477348, Average: 6244, Stddev: 32538
... (output truncated)

```

show services monitoring rpm probe-results owner

```
user@host> show services monitoring rpm probe-results owner icmp-junos
```

```

Owner: icmp-junos, Test: icmp-junos-1
Target address: 10.3.1.2, Probe type: icmp-ping, Test size: 2
Probe results:
  Probe response received
  Probe sent time: 11/18/20 12:42:05.426631
  Probe rcvd time: 11/18/20 12:42:05.438092, Client offload timestamping

```

Rtt: 11461 usec, Rtt jitter: 4982 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 6479, Maximum: 11461, Average: 8970, Stddev: 2491

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 4982, Maximum: 4982, Average: 4982, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 5845, Maximum: 7279, Average: 6562, Stddev: 717

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 1434, Maximum: 1434, Average: 1434, Stddev: 0

Results over all tests:

Probes sent: 442, Probes received: 438, Loss percentage: 0.9

Measurement: Round trip time (usec)

Samples: 438, Minimum: 4783, Maximum: 41643, Average: 8081, Stddev: 3003

Measurement: Round trip jitter (usec)

Samples: 438, Minimum: 3, Maximum: 35766, Average: 2564, Stddev: 3372

Owner: icmp-junos, Test: icmp-junos-2

Target address: 192.168.22.22, Source address: 192.168.11.11, Probe type: icmp-ping, Test size: 2

Probe results:

Probe response received

Probe sent time: 11/18/20 12:42:01.146976

Probe rcvd time: 11/18/20 12:42:01.160419, Client offload timestamping

Rtt: 13443 usec, Rtt jitter: 6738 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 6705, Maximum: 13443, Average: 10074, Stddev: 3369

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 6738, Maximum: 6738, Average: 6738, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 5884, Maximum: 11405, Average: 8644, Stddev: 2762

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 5521, Maximum: 5521, Average: 5521, Stddev: 0

Results over all tests:

Probes sent: 442, Probes received: 439, Loss percentage: 0.67

Measurement: Round trip time (usec)

Samples: 439, Minimum: 5479, Maximum: 42406, Average: 8884, Stddev: 3884

Measurement: Round trip jitter (usec)

Samples: 439, Minimum: 7, Maximum: 36177, Average: 3384, Stddev: 4437

Owner: icmp-junos, Test: icmp-junos-3

Target address: 10.4.1.2, Probe type: icmp-ping-timestamp, Test size: 2

Probe results:

Probe response received

Probe sent time: 11/18/20 12:42:02.128854

Probe rcvd time: 11/18/20 12:42:02.142671, Client offload timestamping

Rtt: 13817 usec, Rtt jitter: 6024 usec

Egress jitter: 2366 usec, Ingress jitter: 3658 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 7793, Maximum: 13817, Average: 10805, Stddev: 3012

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 6024, Maximum: 6024, Average: 6024, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 8043, Maximum: 10271, Average: 9157, Stddev: 1114

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 2228, Maximum: 2228, Average: 2228, Stddev: 0

Results over all tests:

Probes sent: 442, Probes received: 440, Loss percentage: 0.45

Measurement: Round trip time (usec)

Samples: 440, Minimum: 5486, Maximum: 102674, Average: 12544, Stddev: 11088

Measurement: Round trip jitter (usec)

Samples: 440, Minimum: 1, Maximum: 92729, Average: 6896, Stddev: 13940

Owner: icmp-junos, Test: icmp-junos-4

Target address: 10.0.22.22, Source address: 10.0.11.11, Probe type: icmp-ping-timestamp, Test size: 2

Probe results:

Probe response received

Probe sent time: 11/18/20 12:42:02.098433

Probe rcvd time: 11/18/20 12:42:02.140279, Client offload timestamping

Rtt: 41846 usec, Rtt jitter: 33904 usec

Egress jitter: 19853 usec, Ingress jitter: 14051 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 7942, Maximum: 41846, Average: 24894, Stddev: 16952

```

Measurement: Round trip jitter (usec)
  Samples: 2, Minimum: 33904, Maximum: 33904, Average: 33904, Stddev: 0
Results over last test:
Probes sent: 2, Probes received: 2, Loss percentage: 0
Measurement: Round trip time (usec)
  Samples: 2, Minimum: 8066, Maximum: 13446, Average: 10756, Stddev: 2690
Measurement: Round trip jitter (usec)
  Samples: 2, Minimum: 5380, Maximum: 5380, Average: 5380, Stddev: 0
Results over all tests:
Probes sent: 442, Probes received: 439, Loss percentage: 0.67
Measurement: Round trip time (usec)
  Samples: 439, Minimum: 5771, Maximum: 100477, Average: 12726, Stddev: 11125
Measurement: Round trip jitter (usec)
  Samples: 439, Minimum: 6, Maximum: 92178, Average: 7331, Stddev: 14240

```

show services monitoring rpm probe-results owner test

```
user@host> show services monitoring rpm probe-results owner icmp-junos test icmp-junos-1
```

```

Owner: icmp-junos, Test: icmp-junos-1
Target address: 10.3.1.2, Probe type: icmp-ping, Test size: 2
Probe results:
  Probe response received
  Probe sent time: 11/18/20 12:48:35.484952
  Probe rcvd time: 11/18/20 12:48:35.493256, Client offload timestamping
  Rtt: 8304 usec, Rtt jitter: 2524 usec
Results over current test:
Probes sent: 2, Probes received: 2, Loss percentage: 0
Measurement: Round trip time (usec)
  Samples: 2, Minimum: 5780, Maximum: 8304, Average: 7042, Stddev: 1262
Measurement: Round trip jitter (usec)
  Samples: 2, Minimum: 2524, Maximum: 2524, Average: 2524, Stddev: 0
Results over last test:
Probes sent: 2, Probes received: 2, Loss percentage: 0
Measurement: Round trip time (usec)
  Samples: 2, Minimum: 7124, Maximum: 7768, Average: 7446, Stddev: 322
Measurement: Round trip jitter (usec)
  Samples: 2, Minimum: 644, Maximum: 644, Average: 644, Stddev: 0
Results over all tests:
Probes sent: 454, Probes received: 450, Loss percentage: 0.88
Measurement: Round trip time (usec)

```

```
Samples: 450, Minimum: 4783, Maximum: 41643, Average: 8072, Stddev: 2976
Measurement: Round trip jitter (usec)
Samples: 450, Minimum: 2, Maximum: 35766, Average: 2535, Stddev: 3340
```

Release Information

Command introduced in Junos OS Evolved Release 20.1R1.

show services monitoring twamp client control-info

IN THIS SECTION

- [Syntax | 1740](#)
- [Description | 1740](#)
- [Options | 1741](#)
- [Required Privilege Level | 1741](#)
- [Output Fields | 1741](#)
- [Sample Output | 1742](#)
- [Release Information | 1742](#)

Syntax

```
show services monitoring twamp client control-info
<control-connection control-connection-name>
```

Description

Display information about the control connections established between the Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify the control-connection option when you issue the command.

Options

`control-connection` *control-connection-name*

(Optional) Display information about the specified control connection.

Required Privilege Level

view

Output Fields

Table 172 on page 1741 lists the output fields for the `show services monitoring twamp client control-info` command. Output fields are listed in the approximate order in which they appear.

Table 172: show services monitoring twamp client control-info Output Fields

Field Name	Field Description	Levels of Output
Control name	Control-connection name that uniquely identifies the connection between the TWAMP server and a particular client	All
Client address:port	Client IP address and port number	All
Server address:port	Server IP address and port number	All
Active tests	Number of active tests	All
Control status	The status of the connection; can be sleeping, testing, or disconnected	All
Auth mode	Authentication mode	All

Sample Output

show services monitoring twamp client control-info

```
user@host> show services monitoring twamp client control-info
```

Control name	Client address:port	Server address:port	Active tests	Control status	Auth mode
mngd-deflt-1	10.6.1.1:57629	10.6.1.2:60000	2	testing	none
mngd-deflt-2	10.5.1.1:52777	10.5.1.2:862	0	sleeping	none
mngd-inst-1	10.0.1.1:43033	10.0.3.3:60001	2	testing	none
mngd-inst-2	10.0.1.1:34167	10.0.2.2:862	2	testing	none

show services monitoring twamp client control-info control-connection

```
user@host> show services monitoring twamp client control-info control-connection mngd-inst-1
```

Control name	Client address:port	Server address:port	Active tests	Control status	Auth mode
mngd-inst-1	10.0.1.1:43033	10.0.3.3:60001	2	testing	none

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

[control-connection \(Junos OS Evolved\) | 1019](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

show services monitoring twamp client history-results

IN THIS SECTION

- [Syntax | 1743](#)
- [Description | 1743](#)
- [Options | 1744](#)
- [Required Privilege Level | 1744](#)
- [Output Fields | 1744](#)
- [Sample Output | 1746](#)
- [Release Information | 1749](#)

Syntax

```
show services rpm twamp client history-results
control-connection control-connection-name
<detail>
<since time>
<source-address address>
<target address>
<test-session test-session-name>
```

Description

Display standard information about the results of the last 50 probes for a Two-Way Active Measurement Protocol (TWAMP) control connection. You can also view the history results of the probes or test packets sent from a TWAMP client to a TWAMP server by source address, by target address, or for a specific test session associated with the control connection. To change the number of probe results displayed, configure the history-size statement at the [edit monitoring twamp client control-connection *control-connection-name* test-session *test-session-name* hierarchy level.

Options

<control-connection <i>control-connection-name</i>	(Required) Display information for the specified control connection between a TWAMP client and a TWAMP server.
detail	(Optional) Display detailed information about the control connection.
since <i>time</i>	(Optional) Display information from the specified time. Specify time as <i>yyyy-mm-dd.hh:mm:ss</i> .
source-address <i>address</i>	(Optional) Display information only for those probes with the specified IPv4 source address.
target <i>address</i>	(Optional) Display information only for those probes with the specified IPv4 target address.
test-session <i>test-session-name</i>	(Optional) Display information for the specified test session associated with the control connection between a TWAMP client and a TWAMP server.

Required Privilege Level

view

Output Fields

[Table 173 on page 1744](#) lists the output fields for the `show services monitoring twamp client history-results` command. Output fields are listed in the approximate order in which they appear.

Table 173: show services monitoring twamp client history-results Output Fields

Field Name	Field Description	Level of Output
Owner	Probe owner or the TWAMP client.	All levels
Test	Name of a test for a TWAMP probe instance.	All levels
Probe received	Timestamp when the probe result was determined.	All levels
Round trip time	Average ping round-trip time (RTT), in microseconds.	All levels

Table 173: show services monitoring twamp client history-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Probe results	<p>Result of a particular probe performed by a remote host. The following information is contained in the results:</p> <ul style="list-style-type: none"> • Probe response received—Timestamp when the probe result was determined. • Probe sent time—Timestamp when the probe was sent. • Probe rcvd time—Timestamp when the probe was received. • Rtt—Average ping round-trip time (RTT), in microseconds. • Rtt jitter—Average ping round-trip time (RTT) jitter, in microseconds. • Egress jitter—Average ping egress jitter, in microseconds. • Ingress jitter—Average ping ingress jitter, in microseconds. 	detail
Results over current test	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	detail
Probes sent	Number of probes sent with the current test.	detail
Probes received	Number of probe responses received within the current test.	detail
Loss percentage	Percentage of lost probes for the current test.	detail
Measurement	<p>Measurements for round-trip time (RTT).</p> <ul style="list-style-type: none"> • Minimum—Minimum RTT measured over the course of the current test. • Maximum—Maximum RTT measured over the course of the current test. • Average—Average RTT measured over the course of the current test. 	detail

Sample Output

show services monitoring twamp client history-results control-connection

```

user@host> show services monitoring twamp client history-results control-connection mngd-deflt
Owner, Test                Probe sent                Probe received            Round trip
time
mngd-deflt, test-2         08/27/20 18:32:40.532435  08/27/20 18:32:40.553237  20216 usec
mngd-deflt, test-1         08/27/20 18:32:40.532536  08/27/20 18:32:40.555498  22921 usec
mngd-deflt, test-1         08/27/20 18:33:40.623973  08/27/20 18:33:40.647633  23597 usec
mngd-deflt, test-2         08/27/20 18:33:40.624073  08/27/20 18:33:40.648176  24070 usec
mngd-deflt, test-2         08/27/20 18:34:40.630696  08/27/20 18:34:40.641727  10992 usec
mngd-deflt, test-1         08/27/20 18:34:40.630741  08/27/20 18:34:40.644139  13350 usec
mngd-deflt, test-1         08/27/20 18:35:40.696545  08/27/20 18:35:40.711118  14501 usec
mngd-deflt, test-2         08/27/20 18:35:40.696878  08/27/20 18:35:40.711692  14778 usec
mngd-deflt, test-2         08/27/20 18:36:40.694764  08/27/20 18:36:40.711443  16643 usec
mngd-deflt, test-1         08/27/20 18:36:40.694885  08/27/20 18:36:40.713718  18792 usec
mngd-deflt, test-1         08/27/20 18:37:40.765843  08/27/20 18:37:40.778095  12209 usec
mngd-deflt, test-2         08/27/20 18:37:40.766109  08/27/20 18:37:40.778821  12678 usec
mngd-deflt, test-2         08/27/20 18:38:40.768726  08/27/20 18:38:40.781763  13001 usec
mngd-deflt, test-1         08/27/20 18:38:40.768772  08/27/20 18:38:40.782283  13478 usec
mngd-deflt, test-1         08/27/20 18:39:40.838053  08/27/20 18:39:40.862990  24891 usec
mngd-deflt, test-2         08/27/20 18:39:40.838387  08/27/20 18:39:40.863068  24618 usec
mngd-deflt, test-2         08/27/20 18:40:40.841528  08/27/20 18:40:40.859848  18216 usec
mngd-deflt, test-1         08/27/20 18:40:40.841584  08/27/20 18:40:40.861503  19865 usec
mngd-deflt, test-1         08/27/20 18:41:40.929306  08/27/20 18:41:40.947766  18424 usec
mngd-deflt, test-2         08/27/20 18:41:40.929475  08/27/20 18:41:40.948378  18871 usec
mngd-deflt, test-2         08/27/20 18:42:40.923137  08/27/20 18:42:40.941889  18673 usec
mngd-deflt, test-1         08/27/20 18:42:40.923192  08/27/20 18:42:40.943715  20475 usec
mngd-deflt, test-1         08/27/20 18:43:41.003976  08/27/20 18:43:41.021443  17453 usec
mngd-deflt, test-2         08/27/20 18:43:41.004567  08/27/20 18:43:41.022093  17492 usec
mngd-deflt, test-2         08/27/20 18:44:41.010953  08/27/20 18:44:41.032575  21564 usec
mngd-deflt, test-1         08/27/20 18:44:41.011250  08/27/20 18:44:41.034401  23110 usec
mngd-deflt, test-1         08/27/20 18:45:41.106478  08/27/20 18:45:41.124683  18173 usec
mngd-deflt, test-2         08/27/20 18:45:41.106667  08/27/20 18:45:41.125132  18433 usec
mngd-deflt, test-2         08/27/20 18:46:41.109541  08/27/20 18:46:41.132243  21855 usec
mngd-deflt, test-1         08/27/20 18:46:41.110129  08/27/20 18:46:41.137104  26935 usec
mngd-deflt, test-1         08/27/20 18:47:41.200747  08/27/20 18:47:41.217530  16731 usec
mngd-deflt, test-2         08/27/20 18:47:41.201048  08/27/20 18:47:41.218069  16982 usec
mngd-deflt, test-2         08/27/20 18:48:41.200511  08/27/20 18:48:41.226252  25662 usec
mngd-deflt, test-1         08/27/20 18:48:41.200678  08/27/20 18:48:41.227703  27010 usec
mngd-deflt, test-1         08/27/20 18:49:41.273230  08/27/20 18:49:41.289343  16094 usec

```

mngd-deflt, test-2	08/27/20 18:49:41.273428	08/27/20 18:49:41.289450	16003 usec
mngd-deflt, test-2	08/27/20 18:50:41.283700	08/27/20 18:50:41.296895	13179 usec
mngd-deflt, test-1	08/27/20 18:50:41.283950	08/27/20 18:50:41.297335	13373 usec
mngd-deflt, test-1	08/27/20 18:51:41.342941	08/27/20 18:51:41.368757	25778 usec
mngd-deflt, test-2	08/27/20 18:51:41.343179	08/27/20 18:51:41.368799	25585 usec
mngd-deflt, test-2	08/27/20 18:52:41.360424	08/27/20 18:52:41.376456	15989 usec
mngd-deflt, test-1	08/27/20 18:52:41.360560	08/27/20 18:52:41.377131	16539 usec
mngd-deflt, test-1	08/27/20 18:53:41.426307	08/27/20 18:53:41.440569	14224 usec
mngd-deflt, test-2	08/27/20 18:53:41.426425	08/27/20 18:53:41.440605	14137 usec
mngd-deflt, test-2	08/27/20 18:54:41.431346	08/27/20 18:54:41.443324	11958 usec
mngd-deflt, test-1	08/27/20 18:54:41.432825	08/27/20 18:54:41.443912	11074 usec
mngd-deflt, test-1	08/27/20 18:55:41.498060	08/27/20 18:55:41.511760	13664 usec
mngd-deflt, test-2	08/27/20 18:55:41.498084	08/27/20 18:55:41.512597	14469 usec
mngd-deflt, test-2	08/27/20 18:56:41.497135	08/27/20 18:56:41.513473	15792 usec
mngd-deflt, test-1	08/27/20 18:56:41.497152	08/27/20 18:56:41.513544	16380 usec
mngd-deflt, test-1	08/27/20 18:57:41.579541	08/27/20 18:57:41.592352	12750 usec
mngd-deflt, test-2	08/27/20 18:57:41.579718	08/27/20 18:57:41.592905	13143 usec
mngd-deflt, test-2	08/27/20 18:58:41.575705	08/27/20 18:58:41.587480	11738 usec
mngd-deflt, test-1	08/27/20 18:58:41.575725	08/27/20 18:58:41.591966	16195 usec
mngd-deflt, test-1	08/27/20 18:59:41.644790	08/27/20 18:59:41.661570	16763 usec
mngd-deflt, test-2	08/27/20 18:59:41.644812	08/27/20 18:59:41.662216	17390 usec
mngd-deflt, test-2	08/27/20 19:00:41.646420	08/27/20 19:00:41.658933	12495 usec
mngd-deflt, test-1	08/27/20 19:00:41.646725	08/27/20 19:00:41.659386	12646 usec
mngd-deflt, test-1	08/27/20 19:01:41.710108	08/27/20 19:01:41.725249	15097 usec
mngd-deflt, test-2	08/27/20 19:01:41.710156	08/27/20 19:01:41.726253	16059 usec
mngd-deflt, test-2	08/27/20 19:02:41.710661	08/27/20 19:02:41.728264	17552 usec
mngd-deflt, test-1	08/27/20 19:02:41.711725	08/27/20 19:02:41.728481	16725 usec
mngd-deflt, test-1	08/27/20 19:03:41.775533	08/27/20 19:03:41.792378	16816 usec
mngd-deflt, test-2	08/27/20 19:03:41.775601	08/27/20 19:03:41.792795	17133 usec
mngd-deflt, test-2	08/27/20 19:04:41.781272	08/27/20 19:04:41.799808	18327 usec
mngd-deflt, test-1	08/27/20 19:04:41.781391	08/27/20 19:04:41.801613	20187 usec
mngd-deflt, test-1	08/27/20 19:05:41.862516	08/27/20 19:05:41.883037	20497 usec
mngd-deflt, test-2	08/27/20 19:05:41.862527	08/27/20 19:05:41.883585	21018 usec
mngd-deflt, test-2	08/27/20 19:06:41.869384	08/27/20 19:06:41.883258	13837 usec
mngd-deflt, test-1	08/27/20 19:06:41.869466	08/27/20 19:06:41.884977	15475 usec
mngd-deflt, test-1	08/27/20 19:07:41.945802	08/27/20 19:07:41.964598	18758 usec
mngd-deflt, test-2	08/27/20 19:07:41.945817	08/27/20 19:07:41.965232	19379 usec
mngd-deflt, test-2	08/27/20 19:08:41.953929	08/27/20 19:08:41.972951	18950 usec
mngd-deflt, test-1	08/27/20 19:08:41.954886	08/27/20 19:08:41.974187	19258 usec
mngd-deflt, test-1	08/27/20 19:09:42.037577	08/27/20 19:09:42.056201	18570 usec
mngd-deflt, test-2	08/27/20 19:09:42.037629	08/27/20 19:09:42.056842	19171 usec
mngd-deflt, test-2	08/27/20 19:10:42.039524	08/27/20 19:10:42.056397	16793 usec
mngd-deflt, test-1	08/27/20 19:10:42.039601	08/27/20 19:10:42.059688	20038 usec

mngd-deflt, test-1	08/27/20 19:11:42.100698	08/27/20 19:11:42.118606	17881 usec
mngd-deflt, test-2	08/27/20 19:11:42.100942	08/27/20 19:11:42.119158	18188 usec
mngd-deflt, test-2	08/27/20 19:12:42.102936	08/27/20 19:12:42.122253	18728 usec
mngd-deflt, test-1	08/27/20 19:12:42.103062	08/27/20 19:12:42.125590	22511 usec
mngd-deflt, test-1	08/27/20 19:13:42.170855	08/27/20 19:13:42.192140	21239 usec
mngd-deflt, test-2	08/27/20 19:13:42.171132	08/27/20 19:13:42.192333	21158 usec
mngd-deflt, test-2	08/27/20 19:14:42.162464	08/27/20 19:14:42.182361	19794 usec
mngd-deflt, test-1	08/27/20 19:14:42.162601	08/27/20 19:14:42.184849	22206 usec
mngd-deflt, test-1	08/27/20 19:15:42.239051	08/27/20 19:15:42.257361	18283 usec
mngd-deflt, test-2	08/27/20 19:15:42.239516	08/27/20 19:15:42.257933	18388 usec
mngd-deflt, test-2	08/27/20 19:16:42.242182	08/27/20 19:16:42.263899	21610 usec
mngd-deflt, test-1	08/27/20 19:16:42.242245	08/27/20 19:16:42.265736	23434 usec
mngd-deflt, test-1	08/27/20 19:17:42.346073	08/27/20 19:17:42.363597	17477 usec
mngd-deflt, test-2	08/27/20 19:17:42.346432	08/27/20 19:17:42.364867	18345 usec
mngd-deflt, test-2	08/27/20 19:18:42.349453	08/27/20 19:18:42.378702	29204 usec
mngd-deflt, test-1	08/27/20 19:18:42.349502	08/27/20 19:18:42.378729	29171 usec
mngd-deflt, test-1	08/27/20 19:19:42.430651	08/27/20 19:19:42.444923	13792 usec
mngd-deflt, test-2	08/27/20 19:19:42.430842	08/27/20 19:19:42.446115	15236 usec
mngd-deflt, test-2	08/27/20 19:20:42.428240	08/27/20 19:20:42.443861	15571 usec
mngd-deflt, test-1	08/27/20 19:20:42.428294	08/27/20 19:20:42.444358	16020 usec
mngd-deflt, test-1	08/27/20 19:21:42.494229	08/27/20 19:21:42.509517	15217 usec
mngd-deflt, test-2	08/27/20 19:21:42.494436	08/27/20 19:21:42.510934	16478 usec

show services monitoring twamp client history-results control-connection detail

```
user@host> show services monitoring twamp client history-results control-connection mngd-deflt
detail
```

```
Owner: mngd-deflt, Test: test-1, Probe type: twamp
```

```
Probe results:
```

```
Probe response received
```

```
Probe sent time: 08/27/20 18:35:40.696545
```

```
Probe rcvd time: 08/27/20 18:35:40.711118, Client and server offload timestamping
```

```
Rtt: 14501 usec, Rtt jitter: 0 usec
```

```
Egress jitter: 0 usec, Ingress jitter: 0 usec
```

```
Results over current test:
```

```
Probes sent: 1, Probes received: 1, Loss percentage: 0
```

```
Measurement: Round trip time (usec)
```

```
Samples: 1, Minimum: 14501, Maximum: 14501, Average: 14501, Stddev: 0
```

```
Owner: mngd-deflt, Test: test-2, Probe type: twamp
```

```
Probe results:
```

```

Probe response received
Probe sent time: 08/27/20 18:35:40.696878
Probe rcvd time: 08/27/20 18:35:40.711692, Client and server offload timestamping
Rtt: 14778 usec, Rtt jitter: 0 usec
Egress jitter: 0 usec, Ingress jitter: 0 usec
Results over current test:
  Probes sent: 1, Probes received: 1, Loss percentage: 0
  Measurement: Round trip time (usec)
    Samples: 1, Minimum: 14778, Maximum: 14778, Average: 14778, Stddev: 0
  ...

```

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

[control-connection \(Junos OS Evolved\) | 1019](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

show services monitoring twamp client probe-results

IN THIS SECTION

- [Syntax | 1750](#)
- [Description | 1750](#)
- [Options | 1750](#)
- [Required Privilege Level | 1751](#)
- [Output Fields | 1751](#)
- [Sample Output | 1755](#)
- [Release Information | 1759](#)

Syntax

```
show services monitoring twamp client probe-results
<control-connection control-connection-name>
<source-address>
<status (FAIL | PASS)>
<target>
<test-session test-session-name>
```

Description

Display the results of the most recent Two-Way Active Measurement Protocol (TWAMP) probes. By default, all probe results are displayed, unless you specify the control-connection option when you issue the command. The control-connection option is required when you want to issue the command with one of the other options. For example, to see the probe results for a particular test session, you must also include the control-connection option to specify which connection contains the test session: **show services monitoring twamp client probe-results control-connection *control-connection-name* test-session *test-session-name***.

Options

control-connection <i>control-connection-name</i>	(Optional, but required if specifying one of other options as well) Display probe results for the specified control connection between a TWAMP client and a TWAMP server.
source-address <i>address</i>	(Optional) Display results only for those probes with the specified IPv4 source address associated with a particular control connection.
status (FAIL PASS)	<p>(Optional) Display results only for those probes with the specified status associated with a particular control connection.</p> <ul style="list-style-type: none"> • Values: Specify one of the following: <ul style="list-style-type: none"> • FAIL—Display results for failed probes. • PASS—Display results for passed probes.
target <i>address</i>	(Optional) Display results only for those probes with the specified IPv4 target address associated with a particular control connection.
test-session <i>test-session-name</i>	(Optional) Display probe results for the specified test session associated with the control connection between a TWAMP client and a TWAMP server.

Required Privilege Level

view

Output Fields

[Table 174 on page 1751](#) lists the output fields for the `show services monitoring twamp client probe-results` command. Output fields are listed in the approximate order in which they appear.

Table 174: show services monitoring twamp client probe-results Output Fields

Field Name	Field Description	Level of Output
Owner	Owner name. Entries are listed in alphabetical order by owner name and then by test name. When you configure the control-connection statement at the <code>[edit services monitoring twamp client]</code> hierarchy level, this field displays the configured control-connection name.	All levels
Test	Name of a test representing a collection of probes. When you configure the test-session test-name statement at the <code>[edit services monitoring twamp client control-connection]</code> hierarchy level, the field displays the configured test-session name.	All levels
Target address	Destination address used for the probes. This field is displayed when the probes are sent to the configured targets.	All levels
Source address	Source address used for the probes.	All levels
Probe type	Protocol configured on the receiving probe server: twamp	All levels
Test size	Number of probes within a test.	All levels

Table 174: show services monitoring twamp client probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Probe results	<p>Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:</p> <ul style="list-style-type: none"> • Probe response received <ul style="list-style-type: none"> • Probe sent time—Timestamp when the probe's results were sent. • Probe rcvd time—Timestamp when the probe's results were received. • Client and server offload timestamping—If PFE timestamping is configured, an entry appears at this point. • Rtt—Average ping round-trip time (RTT), in microseconds. • Rtt jitter—Round-trip jitter, in microseconds. • Egress jitter—Egress jitter, in microseconds. • Ingress jitter—Ingress jitter, in microseconds. 	All levels

Table 174: show services monitoring twamp client probe-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Results over current test	<p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent within the current test. • Probes received—Number of probe responses received within the current test. • Loss percentage—Percentage of lost probes for the current test. • Measurement—Measurement type. Possible values are round-trip time (RTT) and round-trip jitter <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT measured over the course of the current test. • Maximum—Maximum RTT measured over the course of the current test. • Average—Average RTT measured over the course of the current test. • Stddev—Standard deviation, in microseconds. 	All levels

Table 174: show services monitoring twamp client probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Results over last test	<p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed.</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent for the most recently completed test. Probes received—Number of probe responses received for the most recently completed test. Loss percentage—Percentage of lost probes for the most recently completed test. Measurement—Measurement type. Possible values are round-trip time (RTT) and round-trip jitter. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT measured for the most recently completed test. Maximum—Maximum RTT measured for the most recently completed test. Average—Average RTT measured for the most recently completed test. Stddev—Standard deviation, in microseconds. 	All levels

Table 174: show services monitoring twamp client probe-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Results over all tests	<p>Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent in all tests. • Probes received—Number of probe responses received in all tests. • Loss percentage—Percentage of lost probes in all tests. • Measurement—Measurement type. Possible values are round-trip time (RTT) and positive round-trip jitter. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT measured over the course of the current test. • Maximum—Maximum RTT measured over the course of the current test. • Average—Average RTT measured over the course of the current test. • Stddev—Standard deviation, in microseconds. 	All levels

Sample Output

show services monitoring twamp client probe-results

```
user@host> show services monitoring twamp client probe-results
```

```
Owner: light-deflt, Test: test-5
```

```
Target address: 10.0.1.2, Source address: 10.0.1.1, Probe type: twamp, Test size: 2
```

```
Probe results:
```

```
Probe response received
```

```
Probe sent time: 08/27/20 19:50:25.549385
```

```
Probe rcvd time: 08/27/20 19:50:25.564693, Client and server offload timestamping
```

```
Rtt: 15224 usec, Rtt jitter: 673 usec
```

```
Egress jitter: 850 usec, Ingress jitter: 177 usec
```

```
Results over current test:
```

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 15224, Maximum: 15897, Average: 15560, Stddev: 358

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 673, Maximum: 673, Average: 673, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 14762, Maximum: 16395, Average: 15578, Stddev: 825

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 1633, Maximum: 1633, Average: 1633, Stddev: 0

Results over all tests:

Probes sent: 714, Probes received: 714, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 714, Minimum: 8740, Maximum: 307066, Average: 20625, Stddev: 12145

Measurement: Round trip jitter (usec)

Samples: 714, Minimum: 6, Maximum: 292383, Average: 5420, Stddev: 15797

Owner: light-deflt, Test: test-6

Target address: 192.168.33.33, Source address: 192.168.11.11, Probe type: twamp, Test size: 2

Probe results:

Probe response received

Probe sent time: 08/27/20 19:50:25.549497

Probe rcvd time: 08/27/20 19:50:25.564740, Client and server offload timestamping

Rtt: 15180 usec, Rtt jitter: 125 usec

Egress jitter: 114 usec, Ingress jitter: 11 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 15180, Maximum: 15305, Average: 15242, Stddev: 138

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 125, Maximum: 125, Average: 125, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 14823, Maximum: 16362, Average: 15592, Stddev: 779

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 1539, Maximum: 1539, Average: 1539, Stddev: 0

Results over all tests:

Probes sent: 712, Probes received: 712, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 712, Minimum: 9961, Maximum: 57700, Average: 20115, Stddev: 5615

Measurement: Round trip jitter (usec)

Samples: 712, Minimum: 7, Maximum: 44088, Average: 4691, Stddev: 4920

Owner: mngd-deflt, Test: test-1

Target address: 10.0.1.2, Source address: 10.0.1.1, Probe type: twamp, Test size: 2

Probe results:

Probe response received

Probe sent time: 08/27/20 19:50:43.572219

Probe rcvd time: 08/27/20 19:50:43.597556, Client and server offload timestamping

Rtt: 25299 usec, Rtt jitter: 12621 usec

Egress jitter: 1312 usec, Ingress jitter: 11309 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 12678, Maximum: 25299, Average: 18988, Stddev: 6312

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 12621, Maximum: 12621, Average: 12621, Stddev: 0

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 13690, Maximum: 13953, Average: 13821, Stddev: 176

Measurement: Round trip jitter (usec)

Samples: 2, Minimum: 263, Maximum: 263, Average: 263, Stddev: 0

Results over all tests:

Probes sent: 714, Probes received: 714, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 714, Minimum: 8496, Maximum: 69366, Average: 18243, Stddev: 5540

Measurement: Round trip jitter (usec)

Samples: 714, Minimum: 10, Maximum: 57276, Average: 5810, Stddev: 5782

Owner: mngd-deflt, Test: test-2

Target address: 192.168.33.33, Source address: 192.168.11.11, Probe type: twamp, Test size: 2

Probe results:

Probe response received

Probe sent time: 08/27/20 19:50:43.572124

Probe rcvd time: 08/27/20 19:50:43.597511, Client and server offload timestamping

Rtt: 25338 usec, Rtt jitter: 11895 usec

Egress jitter: 75 usec, Ingress jitter: 11820 usec

Results over current test:

Probes sent: 2, Probes received: 2, Loss percentage: 0

Measurement: Round trip time (usec)

Samples: 2, Minimum: 13443, Maximum: 25338, Average: 19390, Stddev: 5949

Results over last test:

Probes sent: 2, Probes received: 2, Loss percentage: 0


```

Measurement: Round trip time (usec)
  Samples: 2, Minimum: 12612, Maximum: 15279, Average: 13945, Stddev: 1338
Measurement: Round trip jitter (usec)
  Samples: 2, Minimum: 2667, Maximum: 2667, Average: 2667, Stddev: 0
Results over all tests:
  Probes sent: 714, Probes received: 714, Loss percentage: 0
Measurement: Round trip time (usec)
  Samples: 714, Minimum: 8403, Maximum: 69232, Average: 18341, Stddev: 5290
Measurement: Round trip jitter (usec)
  Samples: 714, Minimum: 0, Maximum: 56810, Average: 4766, Stddev: 5358
...

```

show services monitoring twamp client probe-results control-connection test-session

```

user@host> show services monitoring twamp client probe-results control-connection mngd-deflt
test-session test-2

Owner: mngd-deflt, Test: test-2
Target address: 192.168.33.33, Source address: 192.168.11.11, Probe type: twamp, Test size: 2
Probe results:
  Probe response received
  Probe sent time: 08/27/20 19:50:43.572124
  Probe rcvd time: 08/27/20 19:50:43.597511, Client and server offload timestamping
  Rtt: 25338 usec, Rtt jitter: 11895 usec
  Egress jitter: 75 usec, Ingress jitter: 11820 usec
Results over current test:
  Probes sent: 2, Probes received: 2, Loss percentage: 0
Measurement: Round trip time (usec)
  Samples: 2, Minimum: 13443, Maximum: 25338, Average: 19390, Stddev: 5949
Results over last test:
  Probes sent: 2, Probes received: 2, Loss percentage: 0
Measurement: Round trip time (usec)
  Samples: 2, Minimum: 12612, Maximum: 15279, Average: 13945, Stddev: 1338
Measurement: Round trip jitter (usec)
  Samples: 2, Minimum: 2667, Maximum: 2667, Average: 2667, Stddev: 0
Results over all tests:
  Probes sent: 714, Probes received: 714, Loss percentage: 0
Measurement: Round trip time (usec)
  Samples: 714, Minimum: 8403, Maximum: 69232, Average: 18341, Stddev: 5290
Measurement: Round trip jitter (usec)

```

Samples: 714, Minimum: 0, Maximum: 56810, Average: 4766, Stddev: 5358

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

[control-connection \(Junos OS Evolved\) | 1019](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

show services monitoring twamp client test-info

IN THIS SECTION

- [Syntax | 1759](#)
- [Description | 1760](#)
- [Options | 1760](#)
- [Required Privilege Level | 1760](#)
- [Output Fields | 1760](#)
- [Sample Output | 1761](#)
- [Release Information | 1762](#)

Syntax

```
show services monitoring twamp client test-info  
<control-connection control-connection-name>  
<test-session test-session-name>
```

Description

Display information about the test sessions established between the Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, information for all established control connections and test sessions are displayed, unless you specify the control-connection option when you issue the command. The control-connection option is required when you want to issue the command with one of the other options. For example, to display the information for a particular test session, you must also include the control-connection option to specify which connection contains the test session: **show services monitoring twamp client test-info control-connection *control-connection-name* test-session *test-session-name***.

Options

- control-connection *control-connection-name*** (Optional, but required if specifying the test-session option) Display information about the test sessions on the specified control connection, which is established for control packets exchanged between a TWAMP client and a TWAMP server.
- test-session *test-session-name*** (Optional) Display information about the specified test session, which is established for data packets transmitted between a TWAMP client and a TWAMP server, and the control connection associated with the test session.

Required Privilege Level

view

Output Fields

Table 175 on page 1760 lists the output fields for the show services monitoring twamp client test-info command. Output fields are listed in the approximate order in which they appear.

Table 175: show services monitoring twamp client test-info Output Fields

Field Name	Field Description	Level of Output
Control name	Name that uniquely identifies the control connection between the TWAMP server and a particular client	All
Test name	Name that uniquely identifies the control connection between the TWAMP server and a particular client	

Table 175: show services monitoring twamp client test-info Output Fields (Continued)

Field Name	Field Description	Level of Output
Client address:port	Client IP address and port number	All
Server address:port	Server IP address and port number	All
Test status	Status of the test session; can be active or idle.	All

Sample Output

show services monitoring twamp client test-info

```
user@host> show services monitoring twamp client test-info
```

Control name	Test name	Client address:port	Server address:port	Test status
light-deflt	test-5	10.0.1.1:51274	10.0.1.2:61000	active
light-deflt	test-6	192.168.11.11:35197	192.168.33.33:61000	active
light-inst	test-7	10.2.1.1:45441	10.0.1.2:61001	active
light-inst	test-8	10.0.11.11:43598	10.0.33.33:61001	active
mngd-deflt	test-1	10.0.1.1:50508	10.0.1.2:10015	idle
mngd-deflt	test-2	192.168.11.11:44246	192.168.33.33:10017	idle
mngd-inst	test-3	10.2.1.1:42503	10.2.1.2:10016	idle
mngd-inst	test-4	10.0.11.11:39008	10.0.33.33:10018	idle

show services monitoring twamp client test-info control-connection test-session

```
user@host> show services monitoring twamp client test-info control-connection light-deflt test-session test-6
```

Control name	Test name	Client address:port	Server address:port	Test status
light-deflt	test-6	192.168.11.11:35197	192.168.33.33:61000	active

Release Information

Command introduced in Junos Evolved 20.3R1 .

RELATED DOCUMENTATION

[control-connection \(Junos OS Evolved\) | 1019](#)

[Understand Two-Way Active Measurement Protocol | 686](#)

show services monitoring twamp server control-info

IN THIS SECTION

- [Syntax | 1762](#)
- [Description | 1762](#)
- [Options | 1763](#)
- [Required Privilege Level | 1763](#)
- [Output Fields | 1763](#)
- [Sample Output | 1764](#)
- [Release Information | 1764](#)

Syntax

```
show services monitoring twamp server control-info  
<control-connection control-identifier>
```

Description

Display information about the control connections established between the Two-Way Active Measurement Protocol (TWAMP) server and control-clients for managed servers. By default, all established control connections are displayed, unless you specify the control-connection option when

you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

control-connection
control-identifier (Optional) Specify the numeric identifier for a control connection to display information about that control connection.

Required Privilege Level

view

Output Fields

[Table 176 on page 1763](#) lists the output fields for the `show services monitoring twamp server control-info` command. Output fields are listed in the approximate order in which they appear.

Table 176: show services monitoring twamp server control-info Output Fields

Field Name	Field Description	Level of Output
Control identifier	Numeric identifier that uniquely identifies the session between the TWAMP server and a particular client	All
Client address:port	Client IP address and port number	All
Server address:port	Server IP address and port number	All
Active tests	Number of active tests	All
Control status	Status of the control connection; can be sleeping or testing	All
Auth mode	Authentication mode	All

Sample Output

show services monitoring twamp server control-info

```
user@host> show services monitoring twamp server control-info
```

Control identifier	Client address:port	Server address:port	Active tests	Control status	Auth mode
143	10.0.3.3:36289	10.0.1.1:862	2	testing	none
146	10.1.1.2:40365	10.1.1.1:862	2	testing	none
149	10.5.1.2:49613	10.5.1.1:862	0	sleeping	none
152	10.2.1.2:56958	10.0.1.1:862	2	testing	none

show services monitoring twamp server control-info control-connection

```
user@host> show services monitoring twamp server control-info control-connection 143
```

Control identifier	Client address:port	Server address:port	Active tests	Control status	Auth mode
143	10.0.3.3:36289	10.0.1.1:862	2	testing	none

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

| [Understand Two-Way Active Measurement Protocol](#) | 686

show services monitoring twamp server test-info

IN THIS SECTION

● [Syntax](#) | 1765

- [Description | 1765](#)
- [Options | 1765](#)
- [Required Privilege Level | 1766](#)
- [Output Fields | 1766](#)
- [Sample Output | 1766](#)
- [Release Information | 1767](#)

Syntax

```
show services monitoring twamp server test-info
<control-connection control-identifier>
<test-session test-identifier>
```

Description

Display information about the test sessions established between the Two-Way Active Measurement Protocol (TWAMP) server and control-clients. By default, all established test sessions are displayed, unless you specify the control-connection option when you issue the command. The control-connection option is required when you want to issue the command with the test-session option. For example, to display the information for a particular test session, you must also include the control-connection option to specify which connection contains the test session: **show services monitoring twamp server test-info control-connection *control-identifier* test-session *test-identifier***.

Options

control-connection <i>control-identifier</i>	(Optional, but required if specifying the test-session option) Specify a numeric identifier to display information about the specified control connection, which is established for control packets exchanged between a TWAMP client and a TWAMP server.
test-session <i>test-identifier</i>	(Optional) Specify a numeric identifier to display information about a specific test session.

Required Privilege Level

view

Output Fields

Table 177 on page 1766 lists the output fields for the show services monitoring twamp server test-info command. Output fields are listed in the approximate order in which they appear.

Table 177: show services monitoring twamp server test-info Output Fields

Field Name	Field Description	Level of Output
Control identifier	Identification number that uniquely identifies the control session between the TWAMP server and a particular client	All
Test identifier	Identification number that uniquely identifies the test session between the TWAMP server and a particular client	All
Client address:port	Client IP address and port number	All
Server address:port	Server IP address and port number	All
Test status	Status of the test session; can be active or idle	All
Auth mode	Authentication mode	All

Sample Output

show services monitoring twamp server test-info

```

user@host> show services monitoring twamp server test-info

```

Control identifier	Test identifier	Client address:port	Server address:port	Test status
149	7045	10.5.1.2:13474	192.168.1.1:13474	active
149	7046	10.5.1.2:13475	10.5.1.1:13475	active

152	7043	10.2.1.2:13472	10.0.1.1:13472	active
152	7044	10.2.1.2:13473	10.2.1.1:13473	active

show services monitoring twamp server test-info control-connection

```
user@host> show services monitoring twamp server test-info control-connection 149
```

Control identifier	Test identifier	Client address:port	Server address:port	Test status
149	7045	10.5.1.2:13474	192.168.1.1:13474	active
149	7046	10.5.1.2:13475	10.5.1.1:13475	active

show services monitoring twamp server test-info control-connection test-session

```
user@host> show services monitoring twamp server test-info control-connection 149 test-session 7045
```

Control identifier	Test identifier	Client address:port	Server address:port	Test status
149	7045	10.5.1.2:13474	192.168.1.1:13474	active

Release Information

Command introduced in Junos OS Evolved 20.3R1.

RELATED DOCUMENTATION

[Understand Two-Way Active Measurement Protocol](#) | 686

show services rpm active-servers

IN THIS SECTION

- [Syntax | 1768](#)
- [Description | 1768](#)
- [Options | 1768](#)
- [Required Privilege Level | 1768](#)
- [Output Fields | 1768](#)
- [Sample Output | 1769](#)
- [Release Information | 1769](#)

Syntax

```
show services rpm active-servers
```

Description

Display the protocols and corresponding ports for which a router or switch is configured as a real-time performance monitoring (RPM) server.

Options

This command has no options.

Required Privilege Level

view

Output Fields

[Table 178 on page 1769](#) lists the output fields for the `show services rpm active-servers` command. Output fields are listed in the approximate order in which they appear.

Table 178: show services rpm active-servers Output Fields

Field Name	Field Description
Protocol	Protocol configured on the receiving probe server. The protocol can be the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP).
Port	Port configured on the receiving probe server.
Destination interface name	Output interface name for the probes.

Sample Output

show services rpm active-servers

```
user@host> show services rpm active-servers
  Protocol: TCP, Port: 50000, Destination interface name: lt-0/0/0.0
  Protocol: UDP, Port: 50001, Destination interface name: lt-0/0/0.0
```

Release Information

Command introduced before Junos OS Release 7.4.

show services rpm history-results

IN THIS SECTION

- [Syntax | 1770](#)
- [Description | 1770](#)
- [Options | 1770](#)
- [Required Privilege Level | 1771](#)
- [Output Fields | 1771](#)

- [Sample Output | 1773](#)
- [Release Information | 1775](#)

Syntax

```
show services rpm history-results
<brief | detail>
<dst-interface interface-name>
<owner owner>
<limit number>
<since time>
<source-address address>
<target-address address>
<test name>
```

Description

Display the results stored for the specified real-time performance monitoring (RPM) probes.

Options

- | | |
|--|--|
| none | (Optional) Display the results of the last 50 probes for all RPM instances. |
| brief detail | (Optional) Display the specified level of output. |
| dst-interface
<i>interface-name</i> | (Optional) Display information only for RPM probes that are generated on this MS-MPC or MS-MIC services interface. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the <code>owner</code> option. |
| limit <i>number</i> | <p>(Optional) Limit the number of results that are displayed. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the <code>owner</code> option.</p> <ul style="list-style-type: none"> • Range: 1 through 4,294,967,295 • Default: 100 |

- owner *owner*** (Optional) Display information only for probes with the specified probe owner. You must configure `owner` if you configure any of the following options: `dst-interface`, `limit`, `source-address`, or `target-address`.
- since *time*** (Optional) Display information from the specified time. Specify time as *yyyy-mm-dd hh:mm:ss*.
- source-address *address*** (Optional) Display information only for probes with the specified source address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the `owner` option.
- target-address *address*** (Optional) Display information only for probes with the specified target address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. You must also configure the `owner` option.
- test *name*** (Optional starting in Junos OS Release 18.1R1) Display information only for the specified test.
- Do not configure `test` if you configure any of the following options: `dst-interface`, `limit`, `source-address`, or `target-address`. These options do not work when you configure `test`.

Required Privilege Level

view

Output Fields

[Table 179 on page 1771](#) lists the output fields for the `show services rpm history-results` command. Output fields are listed in the approximate order in which they appear.

Table 179: show services rpm history-results Output Fields

Field Name	Field Description	Level of Output
Owner	Probe owner.	All levels
Test	Name of a test for a probe instance.	All levels
Probe received	Timestamp when the probe result was determined.	All levels

Table 179: show services rpm history-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Round trip time	Average ping round-trip time (RTT), in microseconds.	All levels
Probe results	<p>Result of a particular probe performed by a remote host. The following information is contained in the results:</p> <ul style="list-style-type: none"> • Response received—Timestamp when the probe result was determined. • Rtt—Average ping round-trip time (RTT), in microseconds. 	detail
Results over current test	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	detail
Probes sent	Number of probes sent with the current test.	detail
Probes received	Number of probe responses received within the current test.	detail
Loss percentage	Percentage of lost probes for the current test.	detail

Table 179: show services rpm history-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Measurement	<p>Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-pin-timestamp, the egress and ingress delay:</p> <ul style="list-style-type: none"> • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. • Jitter—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test. • Stddev—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test. 	detail

Sample Output

show services rpm history-results owner test

```

user@host> show services rpm history-results owner p1 test t1
  Owner, Test      Probe received      Round trip time
p1, t1            Wed Aug 12 01:02:35 2009      315 usec
p1, t1            Wed Aug 12 01:02:36 2009      266 usec
p1, t1            Wed Aug 12 01:02:37 2009      314 usec
p1, t1            Wed Aug 12 01:02:38 2009      388 usec
p1, t1            Wed Aug 12 01:02:39 2009      316 usec
p1, t1            Wed Aug 12 01:02:40 2009      271 usec
p1, t1            Wed Aug 12 01:02:41 2009      314 usec
p1, t1            Wed Aug 12 01:02:42 2009      1180 usec

```


show services rpm history-results owner test detail

```
user@host> show services rpm history-results owner p1 test t1 detail
```

```
Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
```

```
Probe results:
```

```
Response received, Wed Aug 12 01:02:35 2009,
```

```
Client and server hardware timestamps
```

```
Rtt: 315 usec
```

```
Results over current test:
```

```
Probes sent: 1, Probes received: 1, Loss percentage: 0
```

```
Measurement: Round trip time
```

```
Samples: 1, Minimum: 315 usec, Maximum: 315 usec, Average: 315 usec,
```

```
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 315 usec
```

```
Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
```

```
Probe results:
```

```
Response received, Wed Aug 12 01:02:36 2009,
```

```
Client and server hardware timestamps
```

```
Rtt: 266 usec, Round trip jitter: -50 usec,
```

```
Round trip interarrival jitter: 3 usec
```

```
Results over current test:
```

```
Probes sent: 2, Probes received: 2, Loss percentage: 0
```

```
Measurement: Round trip time
```

```
Samples: 2, Minimum: 266 usec, Maximum: 315 usec, Average: 291 usec,
```

```
Peak to peak: 49 usec, Stddev: 24 usec, Sum: 581 usec
```

```
Measurement: Negative round trip jitter
```

```
Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
```

```
Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec
```

```
Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
```

```
Probe results:
```

```
Response received, Wed Aug 12 01:02:37 2009,
```

```
Client and server hardware timestamps
```

```
Rtt: 314 usec, Round trip jitter: 49 usec,
```

```
Round trip interarrival jitter: 6 usec
```

```
Results over current test:
```

```
Probes sent: 3, Probes received: 3, Loss percentage: 0
```

```
Measurement: Round trip time
```

```
Samples: 3, Minimum: 266 usec, Maximum: 315 usec, Average: 298 usec,
```

```
Peak to peak: 49 usec, Stddev: 23 usec, Sum: 895 usec
```

```
Measurement: Positive round trip jitter
```

```
Samples: 1, Minimum: 49 usec, Maximum: 49 usec, Average: 49 usec,
```

```

    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 49 usec
Measurement: Negative round trip jitter
    Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

Owner: p1, Test: t1, Probe type: icmp-ping-timestamp
Probe results:
    Response received, Wed Aug 12 01:02:38 2009,
    Client and server hardware timestamps
    Rtt: 388 usec, Round trip jitter: 74 usec,
    Round trip interarrival jitter: 10 usec
Results over current test:
    Probes sent: 4, Probes received: 4, Loss percentage: 0
Measurement: Round trip time
    Samples: 4, Minimum: 266 usec, Maximum: 388 usec, Average: 321 usec,
    Peak to peak: 122 usec, Stddev: 44 usec, Sum: 1283 usec
Measurement: Positive round trip jitter
    Samples: 2, Minimum: 49 usec, Maximum: 74 usec, Average: 62 usec,
    Peak to peak: 25 usec, Stddev: 12 usec, Sum: 123 usec
Measurement: Negative round trip jitter
    Samples: 1, Minimum: 50 usec, Maximum: 50 usec, Average: 50 usec,
    Peak to peak: 0 usec, Stddev: 0 usec, Sum: 50 usec

```

Release Information

Command introduced before Junos OS Release 7.4.

`dst-interface`, `limit`, `source-address`, and `target-address` options introduced in Junos OS Release 18.1R1 on MX Series.

`owner` and `test` options became optional in Junos OS Release 18.1R1 on MX Series.

show services rpm probe-results

IN THIS SECTION

- [Syntax | 1776](#)
- [Description | 1776](#)

- [Options | 1776](#)
- [Required Privilege Level | 1777](#)
- [Output Fields | 1777](#)
- [Sample Output | 1786](#)
- [Release Information | 1791](#)

Syntax

```
show services rpm probe-results
<dst-interface interface-name>
<limit number>
<owner owner>
<source-address address>
<status (fail | pass) >
<target-address address>
<terse>
<test name>
```

Description

Display the results of the most recent real-time performance monitoring (RPM) probes.

Options

All the following options require that you also configure the `owner` option.

- | | |
|---|--|
| dst-interface
<i>interface-name</i> | (Optional) Display information only for RPM probes that are configured on this MS-MPC or MS-MIC services interface. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. |
| limit <i>number</i> | (Optional) Limit the number of results that are displayed. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. <ul style="list-style-type: none"> ● Range: 1 through 4,294,967,295 ● Default: 100 |

none	Display information for all of the most recent RPM probes.				
owner <i>owner</i>	(Optional) Display information only for probes with the specified probe owner. You must configure <code>owner</code> if you configure any other options.				
source-address <i>address</i>	(Optional) Display information only for probes with the specified source address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.				
status	(Optional) Display information only for probes with the specified type of test result. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface. Specify one of the following: <table><tr><td>fail</td><td>Failed tests</td></tr><tr><td>pass</td><td>Passed tests</td></tr></table>	fail	Failed tests	pass	Passed tests
fail	Failed tests				
pass	Passed tests				
target-address <i>address</i>	(Optional) Display information only for probes with the specified target address. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.				
terse	(Optional) Display summary information. This option works only for RPM probes generated on an MS-MPC or MS-MIC services interface.				
test <i>name</i>	(Optional) Display information only for the specified test. Do not configure <code>test</code> if you configure any of the following options: <code>dst-interface</code> , <code>source-address</code> , or <code>target-address</code> . These options do not work when you configure <code>test</code> .				

Required Privilege Level

view

Output Fields

Table 180 on page 1778 lists the output fields for the `show services rpm probe-results` command. Output fields are listed in the approximate order in which they appear.

Table 180: show services rpm probe-results Output Fields

Field Name	Field Description	Level of Output
Owner	Owner name. When you configure the probe owner statement at the [edit services rpm] hierarchy level, this field displays the configured owner name. When you configure BGP neighbor discovery through RPM, the output for this field is Rpm-Bgp-Owner.	none dst-interface limit owner source-address target-address test
Test	Name of a test representing a collection of probes. When you configure the test test-name statement at the [edit services rpm probe owner] hierarchy level, the field displays the configured test name. When you configure BGP neighbor discovery through RPM, the output for this field is Rpm-BGP-Test- <i>n</i> , where <i>n</i> is a cumulative number.	All levels
Target address	Destination IPv4 address used for the probes. This field is displayed when the probes are sent to the configured IPv4 or IPv6 targets or RPM servers.	none dst-interface limit owner source-address target-address terse test
Target inet6-address	Destination IPv6 address used for the probes. This field is displayed when the probes are sent to the configured IPv6 targets or RPM servers.	none dst-interface limit owner source-address target-address terse test
Source address	Source address used for the probes.	none dst-interface limit owner source-address target-address test
Probe type	Protocol configured on the receiving probe server: http-get, http-metadata-get, icmp-ping, icmp-ping-timestamp, tcp-ping, udp-ping, or udp-ping-timestamp.	none dst-interface limit owner source-address target-address test
Test size	Number of probes within a test.	none dst-interface limit owner source-address target-address test

Table 180: show services rpm probe-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Routing Instance Name	<p>(BGP neighbor discovery) Name of the configured (if any) routing instance, logical system name, or both, in which the probe is configured:</p> <ul style="list-style-type: none"> • When a routing instance is defined within a logical system, the logical system name is followed by the routing instance name. A slash (/) is used to separate the two entities. For example, if the routing instance called R1 is configured within the logical system called LS, the name in the output field is LS/R1. • When a routing instance is configured but the default logical system is used, the name in the output field is the name of the routing instance. • When a logical system is configured but the default routing instance is used, the name in the output field is the name of the logical system followed by default. A slash (/) is used to separate the two entities. For example, LS/default. 	<p>none dst-interface limit owner source-address target- address test</p>

Table 180: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Probe results	<p>Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:</p> <ul style="list-style-type: none"> • Response received • Probe sent time—Timestamp when the probe's results was sent. • Probe rcvd/timeout time—Timestamp when the probe's results was received. • Client and server hardware timestamps—If timestamps are configured, an entry appears at this point. • Rtt—Average ping round-trip time (RTT), in microseconds. • Egress jitter—Egress jitter, in microseconds. • Ingress jitter—Ingress jitter, in microseconds. • Round trip jitter—Round-trip jitter, in microseconds. • Egress interarrival jitter—Egress interarrival jitter, in microseconds. • Ingress interarrival jitter—Ingress interarrival jitter, in microseconds. • Round trip interarrival jitter—Round-trip interarrival jitter, in microseconds. 	<p>none dst-interface limit owner source-address target- address test</p>

Table 180: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Results over current test	<p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent within the current test. • Probes received—Number of probe responses received within the current test. • Loss percentage—Percentage of lost probes for the current test. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. 	none dst-interface limit owner source-address target-address test

Table 180: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none">• Peak to peak—Peak-to-peak difference, in microseconds.• Stddev—Standard deviation, in microseconds.• Sum—Statistical sum.	

Table 180: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Results over last test	<p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent for the most recently completed test. Probes received—Number of probe responses received for the most recently completed test. Loss percentage—Percentage of lost probes for the most recently completed test. Test completed—Time the most recent test was completed. Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test. Maximum—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test. Average—Average RTT, ingress delay, or egress delay measured for the most recently completed test. 	none dst-interface limit owner source-address target- address test

Table 180: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none">• Peak to peak—Peak-to-peak difference, in microseconds.• Stddev—Standard deviation, in microseconds.• Sum—Statistical sum.	

Table 180: show services rpm probe-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Results over all tests	<p>Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent in all tests. Probes received—Number of probe responses received in all tests. Loss percentage—Percentage of lost probes in all tests. Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types icmp-ping-timestamp and udp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. Peak to peak—Peak-to-peak difference, in microseconds. Stddev—Standard deviation, in microseconds. Sum—Statistical sum. 	none dst-interface limit owner source-address target- address test

Table 180: show services rpm probe-results Output Fields (Continued)

Field Name	Field Description	Level of Output
Error Stats	<p>Displays error statistics for each probe.</p> <ul style="list-style-type: none"> Invalid client rcv timestamp—Number of client receive timestamp less than client send timestamp. Invalid server send timestamp—Number of server send timestamp less than server receive timestamp. Invalid server processing time—Number of server side spent time greater than RTT. <p>NOTE: Error Stats is displayed in the output only if non-zero statistics exists.</p>	none dst-interface limit owner source-address target-address test
Last Probe Status	Status of the last probe that was sent for the current test (fail or pass).	status
Status	Status of the last completed test (up or down).	status terse
Source-IF	The MS-MPC or MS-MIC services interface that generates the RPM probes.	terse

Sample Output

show services rpm probe-results (IPv4 Targets)

```

user@host> show services rpm probe-results
  Owner: ADSN-J4300.ADSN-J2300.D2, Test: 75300002
  Target address: 172.16.54.172, Source address: 10.206.0.1,
  Probe type: udp-ping-timestamp, Test size: 10 probes
  Probe results:
    Response received
    Probe sent time: Tue Feb  6 14:53:15 2007,
    Probe rcvd/timeout time: Tue Feb 6 14:53:15 2007
    Client and server hardware timestamps
    Rtt: 575 usec, Egress jitter: 5 usec, Ingress jitter: 8 usec,

```

Round trip jitter: 12 usec, Egress interarrival jitter: 8 usec,
 Ingress interarrival jitter: 7 usec, Round trip interarrival jitter: 7 usec,
 Round trip interarrival jitter: 669 usec

Results over current test:

Probes sent: 10, Probes received: 10, Loss percentage: 0

Measurement: Round trip time

Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
 Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec

Measurement: Positive round trip jitter

Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
 Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec

Measurement: Negative round trip jitter

Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
 Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec

Measurement: Egress time

Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
 Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec

Measurement: Positive Egress jitter

Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
 Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec

Measurement: Negative Egress jitter

Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
 Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec

Measurement: Ingress time

Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
 Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec

Measurement: Positive Ingress jitter

Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
 Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec

Measurement: Negative Ingress jitter

Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
 Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec

Results over last test:

Probes sent: 10, Probes received: 10, Loss percentage: 0

Test completed on Tue Feb 6 14:53:16 2007

Measurement: Round trip time

Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
 Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec

Measurement: Positive round trip jitter

Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
 Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec

Measurement: Negative round trip jitter

Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,

```

    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Egress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over all tests:
Probes sent: 560, Probes received: 560, Loss percentage: 0
Measurement: Round trip time
    Samples: 560, Minimum: 805 usec, Maximum: 3114 usec, Average: 1756 usec,
    Peak to peak: 2309 usec, Stddev: 519 usec, Sum: xxxx usec
Measurement: Positive round trip jitter
    Samples: 257, Minimum: 0 usec, Maximum: 2054 usec, Average: 597 usec,
    Peak to peak: 2054 usec, Stddev: 427 usec, Sum: xxxx usec
Measurement: Negative round trip jitter
    Samples: 302, Minimum: 1 usec, Maximum: 1812 usec, Average: 511 usec,
    Peak to peak: 1811 usec, Stddev: 408 usec, Sum: xxxx usec
Measurement: Egress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Egress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Measurement: Ingress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
Measurement: Positive Ingress jitter

```

```

    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
Measurement: Negative Ingress jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Error Stats:
    Invalid client recv timestamp: 3, Invalid server send timestamp: 0
    Invalid server processing time: 0

```

show services rpm probe-results (IPv6 Targets)

```

user@host> show services rpm probe-results
Owner: p, Test: t1
Target inet6-address: 2001:db8:0:1:2a0:a502:0:1da,
Target Port : 34567 Test size: 1000000 probes
Probe results:
    Response received
    Probe sent time: Mon Dec 16 10:48:07 2013
    Probe rcvd/timeout time: Mon Dec 16 10:48:07 2013
    Client and server hardware timestamps
    Rtt: 236 usec, Round trip jitter: -10 usec, Round trip interarrival jitter: 484 usec
Results over current test:
    Probes sent: 10, Probes received: 10, Loss percentage: 0
    Measurement: Round trip time
        Samples: 10, Minimum: 231 usec, Maximum: 298 usec, Average: 268 usec, Peak to peak: 67
        usec, Stddev: 24 usec, Sum: 2682 usec
    Measurement: Positive round trip jitter
        Samples: 3, Minimum: 15 usec, Maximum: 1841 usec, Average: 750 usec, Peak to peak: 1826
        usec, Stddev: 787 usec, Sum: 2251 usec
    Measurement: Negative round trip jitter
        Samples: 7, Minimum: 10 usec, Maximum: 1244 usec, Average: 709 usec, Peak to peak: 1234
        usec, Stddev: 466 usec, Sum: 4961 usec
Results over last test:
    Probes sent: 10, Probes received: 10, Loss percentage: 0
    Test completed on Mon Dec 16 10:48:07 2013
    Measurement: Round trip time
        Samples: 10, Minimum: 231 usec, Maximum: 298 usec, Average: 268 usec, Peak to peak: 67
        usec, Stddev: 24 usec, Sum: 2682 usec
    Measurement: Positive round trip jitter
        Samples: 3, Minimum: 15 usec, Maximum: 1841 usec, Average: 750 usec, Peak to peak: 1826
        usec, Stddev: 787 usec, Sum: 2251 usec

```



```

Measurement: Negative round trip jitter
  Samples: 7, Minimum: 10 usec, Maximum: 1244 usec, Average: 709 usec, Peak to peak: 1234
usec, Stddev: 466 usec, Sum: 4961 usec
Results over all tests(From start of current control session):
  Probes sent: 490, Probes received: 488, Loss percentage: 0
Measurement: Round trip time
  Samples: 488, Minimum: 231 usec, Maximum: 306 usec, Average: 270 usec, Peak to peak: 75
usec, Stddev: 16 usec, Sum: 131586 usec
Measurement: Positive round trip jitter
  Samples: 254, Minimum: 0 usec, Maximum: 10151 usec, Average: 157 usec, Peak to peak:
10151 usec, Stddev: 873 usec, Sum: 39817 usec
Measurement: Negative round trip jitter
  Samples: 233, Minimum: 1 usec, Maximum: 10170 usec, Average: 171 usec, Peak to peak:
10169 usec, Stddev: 888 usec, Sum: 39889 usec

```

show services rpm probe-results owner terse

```

user@host> show services rpm probe-results owner owner1 terse

```

Test Name	Source-IF	Target Address	Status	Last Change
t_1	ms-2/2/0.10	192.0.2.1	UP	0D0H1M29S
t_2	ms-2/2/0.10	192.0.2.1	UP	0D0H1M29S
t_3	ms-2/2/0.10	192.0.2.1	UP	0D0H1M29S

show services rpm probe-results owner status fail

```

user@host> show services rpm probe-results owner owner1 status fail

```

Test Name	Last Probe Status	Status
t_1	FAIL	DOWN
t_2	FAIL	DOWN
t_3	FAIL	DOWN

show services rpm probe-results (BGP Neighbor Discovery)

```

user@host> show services rpm probe-results
Owner: Rpm-Bgp-Owner, Test: Rpm-Bgp-Test-1
Target address: 10.209.152.37, Probe type: icmp-ping, Test size: 5 probes
Routing Instance Name: LS1/RI1
Probe results:

```

```

Response received
Probe sent time: Fri Oct 28 05:20:23 2005
Probe rcvd/timeout time: Fri Oct 28 05:20:23 2005
Rtt: 662 usec
Results over current test:
Probes sent: 5, Probes received: 5, Loss percentage: 0
Measurement: Round trip time
  Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
  Jitter: 133 usec, Stddev: 53 usec
Results over all tests:
Probes sent: 5, Probes received: 5, Loss percentage: 0
Measurement: Round trip time
  Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
  Jitter: 133 usec, Stddev: 53 usec

```

Release Information

Command introduced before Junos OS Release 7.4.

dst-interface, limit, source-address, status, target-address, and terse options introduced in Junos OS Release 18.1R1 on MX Series.

show services rpm rfc2544-benchmarking

IN THIS SECTION

- [Syntax | 1792](#)
- [Description | 1792](#)
- [Options | 1792](#)
- [Required Privilege Level | 1793](#)
- [Output Fields | 1793](#)
- [Sample Output | 1794](#)
- [Release Information | 1800](#)

Syntax

```
show services rpm rfc2544-benchmarking
<aborted-tests (test-id [test-id] | brief | detail)>
<active-tests (test-id [test-id] | brief | detail)>
<completed-tests (test-id [test-id] | brief | detail)>
<summary>
```

Description

Display information about the results of each category or state of the RFC 2544-based benchmarking test, such as terminated tests, active tests, and completed tests, for each real-time performance monitoring (RPM) instance. You can view the results of each test state for all of the configured test IDs or for a specific test ID. Also, you can display statistics about the total number of tests of each state for a high-level, quick analysis. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

You can view the test results of multiple test IDs at the same time by entering the IDs in a single command. If you enter multiple test ID values, you must separate each number with a space.

Options

none	Display test results for all categories.
aborted-tests	(Optional) Display the list of tests that were terminated or stopped. This list includes tests that failed due to various error conditions and tests that you terminated by entering the test service rpm rfc2544-benchmarking test <i>test-name</i> stop command. The Status field in the output specifies the reason for the termination of the test.
active-tests	(Optional) Display the results of the set of tests that are currently running.
brief detail extensive	Display the specified level of output.
completed-tests	(Optional) Display the results of the set of tests that were successfully completed. A completed test is one that passes through all the test steps or states specified in RFC 2544. A test that is marked as completed after it went through all the states from the beginning to the end can still be reported as a failed test. For example, a failed test can be a test that sends the desired number of packets, but does not receive the frames back from the other end.
summary	(Optional) Display summary output.

`test-id test-id` Display test results for the specified unique identifier of the test.

Required Privilege Level

view

Output Fields

[Table 181 on page 1793](#) lists the output fields for the `show services rpm rfc2544-benchmarking` command. Output fields are listed in the approximate order in which they appear.

Table 181: show services rpm rfc2544-benchmarking Output Fields

Field Name	Field Description
Test information	Details of the performed RFC 2544 benchmarking test.
Test id	Unique identifier configured for the test.
Test name	Name configured for the test.
Test type	The type of statistical detail that is collected for the test, based on the configured test type. Throughput-related, latency, frame-loss, or back-to-back frames-related information is displayed for ACX Series routers. Reflected packets-related information is displayed for MX104 Series routers.
Test mode	Mode configured for the test on the router. Test modes are: <ul style="list-style-type: none"> Initiate-and-Terminate: Test frames are initiated from one end and terminated at the same end. This mode requires a reflector to be configured at the peer end to enable the test frames to be returned to the source. This mode is supported only on ACX Series routers. Reflect: Test frames that originate from one end are reflected at the other end on the selected service, such as IPv4 or Ethernet.
Test packet size	Size of the test packets in bytes. This field is valid only when the test mode is Initiate-and-Terminate.

Table 181: show services rpm rfc2544-benchmarking Output Fields (Continued)

Field Name	Field Description
Test state	State of the test that is in progress or active when the output is displayed.
Status	Indicates whether the test is currently in progress or has been terminated. This field is displayed for tests that are in progress or were terminated by entering the test <code>services rpm rfc2544-benchmarking test <test-name / test-id> stop</code> command.
Test start time	Time at which the test started in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS).
Test finish time	Time at which the test completed.
Counters last cleared	Date, time, and how long ago the statistics for the test were cleared. The format is <i>year-month-day hour:minute:second:timezone (hour:minute:second ago)</i> . For example, 2010-05-17 07:51:28 PDT (00:04:33 ago). If you did not clear the statistics previously at any point, Never is displayed.
Number of active tests	Total number of tests that are currently running.
Number of completed tests	Total number of tests that were successfully completed
Number of aborted tests	Total number of tests that were terminated or halted.

Sample Output

show services rpm rfc2544-benchmarking summary

```
user@host> show services rpm rfc2544-benchmarking summary
Tests summary :
```

```
Number of active tests: 0, Number of completed tests: 4,
Number of aborted tests: 52
```

This output indicates that no test iteration is currently in progress (at the time of issue of the command), 4 tests were completed successfully, and 52 tests were halted.

show services rpm rfc2544-benchmarking aborted-tests (ACX Series Router)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests
Test information :
  Test id: 1, Test name: test1, Test type: Throughput
  Test mode: Initiate-and-Terminate
  Test packet size: 64 1280
  Test state: RFC2544_TEST_STATE_STOPPED
  Status: User-aborted-via-cli
  Test start time: 2005-08-05 03:19:58 UTC
  Test finish time: 2005-08-05 03:20:00 UTC
  Counters last cleared: Never

  Test id: 2, Test name: test1, Test type: Throughput
  Test mode: Initiate-and-Terminate
  Test packet size: 64 1280
  Test state: RFC2544_TEST_STATE_STOPPED
  Status: User-aborted-via-cli
  Test start time: 2005-08-05 03:20:00 UTC
  Test finish time: 2005-08-05 03:20:02 UTC
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (ACX Series Router)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
Test information :
  Test id: 18, Test name: test1, Test type: Throughput
  Test mode: Initiate-and-Terminate
  Test packet size: 64 1280
  Test state: RFC2544_TEST_STATE_COMPLETED
  Test start time: 2005-08-05 03:20:34 UTC
  Test finish time: 2005-08-05 03:21:23 UTC
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking active-tests (ACX Series Router)

```

user@host> show services rpm rfc2544-benchmarking active-tests
Test information :
    Test id: 57, Test name: test1, Test type: Back-Back-Frames
    Test mode: Initiate-and-Terminate
    Test packet size: 64 1280
    Test state: RFC2544_TEST_STATE_RUNNING
    Status: Running
    Test start time: 2005-08-05 20:15:41 UTC
    Test finish time: TEST_RUNNING
    Counters last cleared: Never

```

show services rpm rfc2544-benchmarking aborted-tests (MX104 Router)

```

user@host> show services rpm rfc2544-benchmarking aborted-tests
Test information :
    Test id: 1, Test name: prof_tput1, Test type: Reflect
    Test mode: Reflect
    Test packet size: 0
    Test state: TEST_STATE_STOPPED
    Status: Test-intf-ifl-change
    Test start time: 2013-12-16 22:54:27 PST
    Test finish time: 2013-12-16 23:30:28 PST
    Counters last cleared: Never

    Test id: 2, Test name: prof_tput1, Test type: Reflect
    Test mode: Reflect
    Test packet size: 0
    Test state: TEST_STATE_STOPPED
    Status: User-aborted-via-cli
    Test start time: 2013-12-16 23:31:06 PST
    Test finish time: 2013-12-16 23:36:22 PST
    Counters last cleared: Never

    Test id: 3, Test name: prof_tput1, Test type: Reflect
    Test mode: Reflect
    Test packet size: 0
    Test state: TEST_STATE_STOPPED

```

```
Status: User-aborted-via-cli
Test start time: 2013-12-16 23:36:24 PST
Test finish time: 2013-12-17 01:49:24 PST
Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (MX104 Router)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
Test information :
  Test id: 18, Test name: test1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_COMPLETED
  Test start time: 2005-08-05 03:20:34 UTC
  Test finish time: 2005-08-05 03:21:23 UTC
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking active-tests (MX104 Router)

```
user@host> show services rpm rfc2544-benchmarking active-tests
Test information :
  Test id: 4, Test name: prof_tput1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-17 01:49:26 PST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking aborted-tests (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests
Test information :
  Test id: 5, Test name: ts14, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_STOPPED
```



```
Status: User-aborted-via-cli
Test start time: 2020-07-13 16:19:31 CST
Test finish time: 2020-07-13 16:19:37 CST
Counters last cleared: Never
```

show services rpm rfc2544-benchmarking active-tests (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking active-tests
Test information :
  Test id: 1, Test name: ts1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_RUNNING
  Status: Running
  Test start time: 2020-06-15 12:34:52 CST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never

  Test id: 2, Test name: ts14, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_RUNNING
  Status: Running
  Test start time: 2020-07-08 20:46:35 CST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking completed-tests (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking completed-tests
Test information :
  Test id: 6, Test name: ts1, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: TEST_STATE_COMPLETED
  Status: Test-Completed
  Test start time: 2020-07-13 16:28:13 CST
  Test finish time: 2020-07-13 16:30:17 CST
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking aborted-tests detail (SRX300 and SRX550HM)

```
user@host> show services rpm rfc2544-benchmarking aborted-tests detail
```

```
Test information :
```

```
Test id: 1, Test name: ts13, Test type: Reflect
Test mode: Reflect
Test packet size: 0
Test state: TEST_STATE_STOPPED
Status: User-aborted-via-cli
Test start time: 2020-07-10 12:38:26 CST
Test finish time: 2020-07-10 12:49:59 CST
Counters last cleared: Never
```

```
Test Configuration:
```

```
Test mode: Reflect
Duration in seconds: 864000
Test finish wait duration in seconds: 1
Test family: INET
Test iterator pass threshold: 0.50 %
Test receive failure threshold: 0.00 %
Test transmit failure threshold: 0.50 %
Routing Instance Name: default
```

```
Inet family Configuration:
```

```
Egress Interface : ge-0/0/13.0
Destination ipv4 address: 10.0.2.1
Destination udp port: 400
```

Elapsed time	Reflected Packets	Reflected Bytes
692	0	0

```
Test information :
```

```
Test id: 5, Test name: ts14, Test type: Reflect
Test mode: Reflect
Test packet size: 0
Test state: TEST_STATE_STOPPED
Status: User-aborted-via-cli
Test start time: 2020-07-13 16:19:31 CST
Test finish time: 2020-07-13 16:19:37 CST
Counters last cleared: Never
```

Test Configuration:

Test mode: Reflect
 Duration in seconds: 864000
 Test finish wait duration in seconds: 1
 Test family: INET
 Test iterator pass threshold: 0.50 %
 Test receive failure threshold: 0.00 %
 Test transmit failure threshold: 0.50 %
 Routing Instance Name: default

Inet family Configuration:

Egress Interface : ge-0/0/14.0
 Destination ipv4 address: 10.0.3.1
 Destination udp port: 400

Elapsed time	Reflected Packets	Reflected Bytes
5	0	0

Release Information

Command introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)

show services rpm rfc2544-benchmarking test-id**IN THIS SECTION**

- [Syntax | 1801](#)
- [Description | 1801](#)

- Options | 1801
- Required Privilege Level | 1801
- Output Fields | 1801
- Sample Output | 1814
- Release Information | 1825

Syntax

```
show services rpm rfc2544-benchmarking test-id test-id
<brief | detail>
```

Description

Display information about the results of the RFC 2544-based benchmarking test for a specific test ID for each real-time performance monitoring (RPM) instance. The values in the output displayed vary, depending on the state in which the test is passing through, when you issue the command.

Options

- test-id* Display test results for the specified unique identifier.
- brief | detail (Optional) Display the specified level of output.

Required Privilege Level

view

Output Fields

[Table 182 on page 1802](#) lists the output fields for the `show services rpm rfc2544-benchmarking test-id` command. Output fields are listed in the approximate order in which they appear.

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields

Field Name	Field Description	Level of Output
Test information	Details of the performed RFC 2544 benchmarking test.	None specified
Test id	Unique identifier configured for the test.	None specified
Test name	Name configured for the test.	None specified
Test type	The type of actual test run that is collected for the test, based on the configured test type. Throughput-related, latency, frame-loss, or back-to-back frames-related information is displayed for ACX Series routers. Reflected packets-related information is displayed for MX104 Series routers.	None specified
Test mode	<p>Mode configured for the test on the router. Test modes are:</p> <ul style="list-style-type: none"> Initiate-and-Terminate: Test frames are initiated from one end and terminated at the same end. This mode requires a reflector to be configured at the peer end to enable the test frames to be returned to the source. This mode is supported only on ACX Series routers. Reflect: Test frames that originate from one end are reflected back to the originator, such as IPv4 or Ethernet. 	None specified
Test packet size	Size of the test packets in bytes. This field is valid only when the test mode is Initiate-and-Terminate.	None specified
Test state	State of the test that is in progress or active when the output is displayed. For details about the states, see <i>RFC 2544-Based Benchmarking Test States</i> .	None specified
Status	Indicates whether the test is currently in progress or has been terminated.	None specified
Test start time	Time at which the test started in Coordinated Universal Time (UTC) format (YYYY-MM-DD-HH:MM:SS).	None specified

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Test finish time	Time at which the test completed.	None specified
Counters last cleared	Date, time, and how long ago the statistics for the test were cleared. The format is <i>year-month-day hour:minute:second:timezone</i> (<i>hour:minute:second</i> ago). For example, 2010-05-17 07:51:28 PDT (00:04:33 ago). If you did not clear the statistics previously at any point, Never is displayed.	None specified
Test-profile Configuration	(ACX Series routers only) Details of the specified test profile	detail
Test-profile name	(ACX Series routers only) Name of the configured test profile that contains the parameters for the test	detail
Test packet size	(ACX Series routers only) Size of the test packets in bytes	detail
Theoretical max bandwidth	(ACX Series routers only) Theoretical maximum bandwidth configured for the test. This value is typically set to the bandwidth of the server being tested. Valid values are 1 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value tested for this test.	detail
Test Configuration	Details of the configured test ID.	detail
Test mode	Mode configured for the test. Test modes are Initiate-and-Terminate and Reflect.	detail
Duration in seconds	Period in seconds for which the test has been performed.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Test family	The underlying service on which the test is run. Test families are: <ul style="list-style-type: none"> • INET: Indicates that the test is run on a IPV4 service. • CCC: Indicates that the test is run on a circuit cross-connect (CCC) or pseudowire service. 	detail
Routing Instance Name	(ACX Series routers only) Name of the routing instance for the test	detail
Inet family Configuration	Details of the configured inet family for an IPv4 service	detail
Egress Interface	Name of the egress interface from which the test frames are sent	detail
Source ipv4 address	Source IPv4 address used in the IP header of the generated test frame.	detail
Destination ipv4 address	Destination IPv4 address used in the IP header of the generated test frame.	detail
Source udp port	Source UDP port number used in the UDP header of the generated test frame.	detail
Destination udp port	Destination UDP port number used in the UDP header of the generated test frame.	detail
Ccc family Configuration	Details of the configured CCC family for an Ethernet service	detail
Source MAC address	(ACX Series routers only) Source MAC address used in generated test frames for a CCC or Ethernet pseudowire service.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Destination MAC address	(ACX Series routers only) Destination MAC address used in generated test frames for a CCC or Ethernet pseudowire service.	detail
Ivlan-id	(ACX Series routers only) Inner VLAN ID for test-frames.	detail
Ovlan-id	(ACX Series routers only) Outer VLAN ID for test-frames.	detail
Direction egress	Test is run in the egress direction of the interface (NNI)	detail
Direction ingress	Test is run in the ingress direction of the interface (UNI)	detail
Rfc2544 throughput test information	(ACX Series routers only) Details of the throughput test	detail
Initial test load percentage	Percentage of the steady state load for the test.	detail
Test iteration mode	Mode of the test iteration: Binary or step-down.	detail
Test iteration step percent	The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests.	detail
Theoretical max bandwidth	The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested.	detail
Test packet size:	Packet size of the test frames in bytes.	detail
Iteration	Number of the test iteration.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Duration (sec)	Period in seconds for which the test iteration is run	detail
Elapsed time	Amount of time that has passed, in seconds, since the start of the test.	detail
pps	Total count of packets-per-second (pps) transmitted during the test.	detail
Tx Packets	Number of transmitted test packets.	detail
Rx Packets	Number of received test packets.	detail
Tx Bytes	Number of transmitted bytes.	detail
Rx Bytes	Number of received bytes.	detail
Percentage throughput	Percentage of throughput for the test iteration.	detail
Result of the iteration runs (Throughput) :	Results of the completed throughput test for a particular packet size.	detail
Best iteration	Number of the iteration with the highest throughput, among the listed iterations.	detail
Best iteration (pps)	Packets-per-second (pps) count of the iteration with the highest throughput, among the listed iterations.	detail
Best iteration throughput	Percentage of throughput of the iteration with the highest throughput, among the listed iterations.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
RFC2544 Throughput test results summary	Consolidated information of the throughput test.	detail summary
Packet Size	Size of the test packet in bytes.	detail summary
Theoretical rate (pps)	Theoretical frame rate in packets-per-second.	detail summary
Tx Packets	Number of transmitted packets.	detail summary
Rx Packets	Number of received packets.	detail summary
Offered throughput (percentage)	The offered throughput in percentage of the chosen service (such as Layer 3 or Ethernet pseudowire).	detail summary
Measured bandwidth (kbps)	Available bandwidth of the service based on the calculated throughput.	detail summary
Rfc2544 latency test information :	(ACX Series routers only) Details of the latency test	detail
Theoretical max bandwidth	Theoretical maximum bandwidth configured for the test. This value is typically set to the bandwidth of the server being tested. Valid values are 1 Kbps through 1,000,000 Kbps (1 Gbps). The value defined is the highest bandwidth value used for this test.	detail
Initial test load percentage	Percentage of the steady state load for the test.	detail
Duration in seconds	Period in seconds for which the test has been performed.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Test packet size	Size of the test packet in bytes.	detail
Iteration	Number of the test iteration.	detail
Duration (sec)	Period in seconds for which the test iteration is run.	detail
Elapsed time	Amount of time that has passed, in seconds, since the start of the test.	detail
pps	Total count of packets-per-second (pps) transmitted during the test.	detail
Tx Packets	Number of transmitted test packets.	detail
Rx Packets	Number of received test packets.	detail
Latency	Displays the latency parameters.	detail
Min(ns)	Aggregated minimum latency in nanoseconds.	detail
Avg(ns)	Aggregated average latency in nanoseconds.	detail
Max(ns)	Aggregated maximum latency in nanoseconds.	detail
Probe(ns)	Aggregated probe latency in nanoseconds.	detail
Result of the iteration runs (Latency)	Results of the latency test completed for a particular packet size.	detail
Avg (min) Latency	Average of the minimum latency in nanoseconds.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Avg (avg) latency	Average of the average latency in nanoseconds.	detail
Avg (Max) latency	Average of the maximum latency in nanoseconds.	detail
Avg (probe) latency	Average of the probe latency in nanoseconds.	detail
RFC2544 Latency test results summary:	Consolidated statistics of the latency test.	detail summary
Packet Size	Size of the test packet in bytes.	detail summary
Theoretical rate (pps)	Theoretical frame rate in packets-per-second.	detail summary
Tx Packets	Number of transmitted packets.	detail summary
Rx Packets	Number of received packets.	detail summary
Latency	Displays the latency parameters.	detail summary
Min(ns)	Aggregated minimum latency in nanoseconds.	detail summary
Avg(ns)	Aggregated average latency in nanoseconds.	detail summary
Max(ns)	Aggregated maximum latency in nanoseconds.	detail summary
Probe(ns)	Aggregated probe latency in nanoseconds.	detail summary

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Rfc2544 Back-Back test information :	(ACX Series routers only) Details of the back-to-back frames or bursty frames test.	detail
Initial burst length:	Length of the first burst when test frames are sent, as a measure of number of seconds at the rate of Kbps.	detail
Test iteration mode :	Mode of the test iteration: Binary or step-down.	detail
Test iteration step percent	The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests.	detail
Theoretical max bandwidth	The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested.	detail
Test packet size:	Packet size of the test frames in bytes.	detail
Iteration	Number of the test iteration.	detail
Burst Length (Packets)	Number of packets in the burst.	detail
Elapsed time	Amount of time that has passed, in seconds, since the start of the test.	detail
Tx Packets	Number of transmitted test packets.	detail
Rx Packets	Number of received test packets.	detail
Tx Bytes	Number of transmitted bytes.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Rx Bytes	Number of received bytes.	detail
Result of the iteration runs :	Results of the back-to-back frames test completed for a certain packet size.	detail
Best iteration :	Number of the iteration with the longest burst.	detail
Measured burst (num sec)	Time in seconds of the burst of the iteration with the longest burst.	detail
Measured burst (num pkts)	Number of packets during the burst of the iteration with the longest burst.	detail
RFC2544 Back-Back test results summary:	Consolidated statistics of the back-to-back frames test.	detail summary
Packet Size	Size of the test packets in bytes.	detail summary
Measure Burst length (Packets)	Computed burst length in terms of number of packets.	detail summary
Rfc2544 frame-loss test information :	(ACX Series routers only) Details of the frame-loss test.	detail
Initial burst length:	Length of the first burst when test frames are sent, as a measure of number of seconds at the rate of Kbps.	detail
Test iteration mode :	Mode of the test iteration: Binary or step-down.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Test iteration step percent	The test step percentage for tests. If not specified, the default step-percent is 10 percent. This parameter is ignored for all type of tests other than frame-loss tests.	detail
Theoretical max bandwidth	The theoretical limit of the media for the frame size configured for the test. This value is typically set to the bandwidth of the server being tested.	detail
Test packet size	Size of the test packets in bytes.	detail
Iteration	Number of the test iteration.	detail
Duration (sec)	Period, in seconds, for which the test iteration is run.	detail
Offered throughput (percentage)	The offered throughput in percentage of the chosen service (such as Layer 3 or Ethernet pseudowire)	detail
Elapsed time	Amount of time that has passed, in seconds, since the start of the test.	detail
pps	Theoretical frame rate in packets-per-second.	detail
Tx Packets	Number of transmitted test packets.	detail
Rx Packets	Number of received test packets.	detail
Tx Bytes	Number of transmitted bytes.	detail
Rx Bytes	Number of received bytes.	detail

Table 182: show services rpm rfc2544-benchmarking test-id Output Fields (Continued)

Field Name	Field Description	Level of Output
Frame-loss rate %	Percentage of frames that must be forwarded by the router under steady state (constant) load, but were not forwarded due to lack of resources.	detail
Result of the iteration runs :	Results of the frame-loss test completed for a certain packet size.	detail
Frame-loss rate (percent) :	Percentage of dropped frames for the specified packet size	detail
RFC2544 Frame-loss test results summary	Consolidated statistics of the frame-loss test	detail
Packet Size	Size of the test packet in bytes.	detail summary
Theoretical rate (pps)	Theoretical frame rate in packets-per-second.	detail summary
Percentage throughput	Percentage of throughput for the test iteration.	detail summary
Tx Packets	Number of transmitted packets.	detail summary
Rx Packets	Number of received packets.	detail summary
Frame Loss rate percent	Percentage of dropped frames for the specified packet size	detail summary

Sample Output

show services rpm rfc2544-benchmarking test-id detail (Throughput Test on ACX Series Routers)

```

user@host> show services rpm rfc2544-benchmarking test-id 19 detail
Test information :
    Test id: 19, Test name: test1, Test type: Throughput
    Test mode: Initiate-and-Terminate
    Test packet size: 64 1280
    Test state: RFC2544_TEST_STATE_COMPLETED
    Test start time: 2005-07-29 10:25:00 UTC
    Test finish time: 2005-07-29 10:26:02 UTC
    Counters last cleared: Never

Test-profile Configuration:
    Test-profile name: prof_tput
    Test packet size: 64 1280
    Therotical max bandwidth : 993000 kbps

Test Configuration:
    Test mode: Initiate-and-Terminate
    Duration in seconds: 20
    Test family: INET
    Routing Instance Name: default

Inet family Configuration:
    Egress Interface : ge-0/1/1.0
    Source ipv4 address: 192.0.2.1
    Destination ipv4 address: 192.0.2.2
    Source udp port: 2020
    Destination udp port: 3030

Rfc2544 throughput test information :
    Initial test load percentage : 100.00 %
    Test iteration mode : Binary
    Test iteration step percent : 50.00 %
    Therotical max bandwidth : 993000 kbps

Test packet size: 64
Iteration Duration Elapsed pps      Tx      Rx      Tx      Rx      Percentage
      (sec)   time      Packets Packets Bytes    Bytes throughput

```

1	3	3	134918	404754	404754	27523272	27523272	10.00 %
2	20	20	1349184	26983501	26983501	1834878068	1834878068	100.00 %

Result of the iteration runs : Throughput Test complete for packet size 64

Best iteration : 2, Best iteration (pps) : 1349184

Best iteration throughput : 100.00 %

Test packet size: 1280

Iteration	Duration (sec)	Elapsed time	pps	Tx Packets	Rx Packets	Tx Bytes	Rx Bytes	Percentage throughput
1	3	3	9489	28467	28467	36551628	36551628	10.00 %
2	20	20	94896	1897920	1897920	2436929280	2436929280	100.00 %

Result of the iteration runs : Throughput Test complete for packet size 1280

Best iteration : 2, Best iteration (pps) : 94896

Best iteration throughput : 100.00 %

RFC2544 Throughput test results summary:

Packet Size	Theoretical rate (pps)	Tx Packets	Rx Packets	Offered throughput (percentage)	Measured bandwidth (kbps)
64	1349184	26983501	26983501	100.00 %	993000
1280	94896	1897920	1897920	100.00 %	993000

show services rpm rfc2544-benchmarking test-id detail (Latency Test on ACX Series Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 37 detail
```

Test information :

Test id: 37, Test name: test1, Test type: Latency

Test mode: Initiate-and-Terminate

Test packet size: 64 1280

Test state: RFC2544_TEST_STATE_COMPLETED

Test start time: 2005-07-29 10:26:41 UTC

Test finish time: 2005-07-29 10:36:15 UTC

Counters last cleared: Never

Test-profile Configuration:

Test-profile name: prof_latency

Test packet size: 64 1280

Theretical max bandwidth : 993000 kbps

Test Configuration:

Test mode: Initiate-and-Terminate

Duration in seconds: 10

Test family: INET

Routing Instance Name: default

Inet family Configuration:

Egress Interface : ge-0/1/1.0

Source ipv4 address: 192.0.2.1

Destination ipv4 address: 192.0.2.2

Source udp port: 2020

Destination udp port: 3030

Rfc2544 latency test information :

Theretical max bandwidth : 993000 kbps

Initial test load percentage : 100.00 %

Duration in seconds: 10

Test packet size: 64

Iteration	Duration (sec)	Elapsed time	pps	Tx Packets	Rx Packets
1	3	3	134918	404754	404754
2	10	10	1349184	13491751	13491751
3	10	10	1349184	13491751	13491751
4	10	10	1349184	13491751	13491751
5	10	10	1349184	13491751	13491751
6	10	10	1349184	13491751	13491751
7	10	10	1349184	13491751	13491751
8	10	10	1349184	13491751	13491751
9	10	10	1349184	13491751	13491751
10	10	10	1349184	13491751	13491751
11	10	10	1349184	13491751	13491751
12	10	10	1349184	13491751	13491751
13	10	10	1349184	13491751	13491751
14	10	10	1349184	13491751	13491751
15	10	10	1349184	13491751	13491751
16	10	10	1349184	13491751	13491751
17	10	10	1349184	13491751	13491751
18	10	10	1349184	13491751	13491751
19	10	10	1349184	13491751	13491751

20	10	10	1349184	13491751	13491751
21	10	10	1349184	13491751	13491751

----- Latency -----

Min(ns)	Avg(ns)	Max(ns)	Probe(ns)
17464	18770	18880	18784
17472	18799	20488	18848
17472	18799	20416	18816
17472	18799	20440	18704
17464	18799	20376	18880
17464	18799	20232	18832
17464	18799	20400	18848
17472	18799	20240	18864
17472	18799	20264	18848
17464	18799	20264	18880
17472	18800	20320	18864
17464	18799	20176	18864
17464	18800	20248	18864
17464	18800	20272	18864
17464	18799	20472	18832
17464	18799	20256	18880
17464	18799	20336	18848
17464	18800	20688	18848
17472	18800	20504	18864
17464	18799	20448	18768
17472	18799	20240	18864

Result of the iteration runs : Latency Test complete for packet size 64

Avg (min) Latency	: 17466
Avg (avg) latency	: 18799
Avg (Max) latency	: 20360
Avg (probe) latency	: 18844

Test packet size: 1280

Iteration	Duration (sec)	Elapsed time	pps	Tx Packets	Rx Packets
1	3	3	9489	28467	28467
2	10	10	94896	948960	948960
3	10	10	94896	948960	948960
4	10	10	94896	948960	948960
5	10	10	94896	948960	948960
6	10	10	94896	948960	948960

7	10	10	94896	948960	948960
8	10	10	94896	948960	948960
9	10	10	94896	948960	948960
10	10	10	94896	948960	948960
11	10	10	94896	948960	948960
12	10	10	94896	948960	948960
13	10	10	94896	948960	948960
14	10	10	94896	948960	948960
15	10	10	94896	948960	948960
16	10	10	94896	948960	948960
17	10	10	94896	948960	948960
18	10	10	94896	948960	948960
19	10	10	94896	948960	948960
20	10	10	94896	948960	948960
21	10	10	94896	948960	948960

----- Latency -----

Min(ns)	Avg(ns)	Max(ns)	Probe(ns)
68712	70031	70576	69456
68728	70344	71808	70512
68720	70344	71744	70352
68720	70344	71680	70112
68720	70345	71856	70352
68720	70344	71808	70384
68720	70344	71752	70480
68720	70344	71880	70112
68720	70344	71792	70320
68728	70345	73344	70336
68720	70344	71688	70560
68728	70345	71896	70496
68720	70344	71760	70096
68720	70344	71776	70320
68720	70344	71760	70400
68712	70345	71920	70352
68720	70344	71792	70576
68720	70345	71840	70320
68720	70344	71792	70368
68720	70345	71824	70464
68712	70345	71904	70512

Result of the iteration runs : Latency Test complete for packet size 1280

Avg (min) Latency : 68720

```

Avg (avg) latency           : 70344
Avg (Max) latency           : 71880
Avg (probe) latency         : 70371

```

RFC2544 Latency test results summary:

Packet Size	Theoretical Tx rate (pps)	Rx Packets	Packets	----- Latency -----			
				Min(ns)	Avg(ns)	Max(ns)	Probe(ns)
64	1349184	269835020	269835020	17466	18799	20360	18844
1280	94896	18979200	18979200	68720	70344	71880	70371

show services rpm rfc2544-benchmarking test-id detail (Frame Loss Test on ACX Series Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 73 detail
```

Test information :

```

Test id: 73, Test name: test1, Test type: Frame-Loss
Test mode: Initiate-and-Terminate
Test packet size: 64 1280
Test state: RFC2544_TEST_STATE_COMPLETED
Test start time: 2005-07-29 10:38:41 UTC
Test finish time: 2005-07-29 10:41:19 UTC
Counters last cleared: Never

```

Test-profile Configuration:

```

Test-profile name: prof_fl
Test packet size: 64 1280
Therotical max bandwidth : 993000 kbps

```

Test Configuration:

```

Test mode: Initiate-and-Terminate
Duration in seconds: 20
Test family: INET
Routing Instance Name: default

```

Inet family Configuration:

```

Egress Interface : ge-0/1/1.0
Source ipv4 address: 192.0.2.1
Destination ipv4 address: 192.0.2.2

```

Source udp port: 2020
Destination udp port: 3030

Rfc2544 frame-loss test information :

Initial test load percentage : 100.00 %
Test iteration mode : step-down
Test iteration step percent : 10 %
Theoretical max bandwidth : 993000 kbps

Test packet size: 64

Iteration	Duration	Elapsed	Offered	pps	Tx	Rx	Tx	Rx	Frame- loss
	(sec)	time	throughput%		Packets	Packets	Bytes	Bytes	rate %
1	3	3	10.00 %	134918	404754	404754	27523272	27523272	0.00 %
2	20	20	100.00 %	1349184	26983501	26983501	1834878068	1834878068	0.00 %
3	20	20	100.00 %	1349184	26983501	26983501	1834878068	1834878068	0.00 %
4	20	20	100.00 %	1349184	26983501	26983501	1834878068	1834878068	0.00 %

Result of the iteration runs : Frame-loss test complete for packet size 64

Frame-loss rate (percent) : 0.00 %

Test packet size: 1280

Iteration	Duration	Elapsed	Offered	pps	Tx	Rx	Tx	Rx	Frame- loss
	(sec)	time	throughput%		Packets	Packets	Bytes	Bytes	rate %
1	3	3	10.00 %	9489	404754	28467	36551628	36551628	0.00 %
2	20	20	100.00 %	94896	1897920	1897920	2436929280	2436929280	0.00 %
3	20	20	100.00 %	94896	1897920	1897920	2436929280	2436929280	0.00 %
4	20	20	100.00 %	94896	1897920	1897920	2436929280	2436929280	0.00 %

Result of the iteration runs : Frame-loss test complete for packet size 1280

Frame-loss rate (percent) : 0.00 %

RFC2544 Frame-loss test results summary:

Packet Size	Theoretical rate (pps)	Percentage throughput	Tx Packets	Rx Packets	Frame Loss rate percent
64	1349184	100.00 %	26983501	26983501	0.00 %
1280	94896	100.00 %	1897920	1897920	0.00 %

show services rpm rfc2544-benchmarking test-id detail (Back-to-Back Frames Test on ACX Series Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 55 detail
Test information :
    Test id: 55, Test name: test1, Test type: Back-Back-Frames
    Test mode: Initiate-and-Terminate
    Test packet size: 64 1280
    Test state: RFC2544_TEST_STATE_COMPLETED
    Test start time: 2005-07-29 10:36:54 UTC
    Test finish time: 2005-07-29 10:37:57 UTC
    Counters last cleared: Never

Test-profile Configuration:
    Test-profile name: prof_b2b
    Test packet size: 64 1280
    Therotical max bandwidth : 993000 kbps

Test Configuration:
    Test mode: Initiate-and-Terminate
    Duration in seconds: 20
    Test family: INET
    Routing Instance Name: default

Inet family Configuration:
    Egress Interface : ge-0/1/1.0
    Source ipv4 address: 192.0.2.1
    Destination ipv4 address: 192.0.2.2
    Source udp port: 2020
    Destination udp port: 3030

Rfc2544 Back-Back test information :
    Initial burst length: 20 seconds at 993000 kbps
    Test iteration mode : Binary
    Test iteration step percent : 50.00 %

Test packet size: 64
Iteration   Burst Length   Elapsed        Tx              Rx              Tx              Rx
```


	(Packets)	time	Packets	Packets	Bytes	Bytes
1	404754	3	404754	404754	27523272	27523272
2	26983680	20	26983680	26983680	1834890240	1834890240

Result of the iteration runs : Back-Back-Frames Test complete for packet size 64

Best iteration : 2

Measured burst (num sec) : 20 sec,

Measured burst (num pkts) : 26983680 packets

Result of the iteration runs : Back-Back-Frames Test complete for packet size 64

Best iteration : 2

Measured burst (num sec) : 20 sec,

Measured burst (num pkts) : 26983680 packets

Test packet size: 1280

Iteration	Burst Length	Elapsed	Tx	Rx	Tx	Rx
	(Packets)	time	Packets	Packets	Bytes	Bytes
1	28467	3	28467	28467	36551628	36551628
2	1897920	20	1897920	1897920	2436929280	2436929280

Result of the iteration runs : Back-Back-Frames Test complete for packet size 12

Best iteration : 2

Measured burst (num sec) : 20 sec,

Measured burst (num pkts) : 1897920 packets

RFC2544 Back-Back test results summary:

Packet	Measure Burst
Size	length (Packets)
64	26983680 packets
1280	1897920 packets

show services rpm rfc2544-benchmarking test-id detail (Reflection Test on MX104 Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 1 detail
Test information :
  Test id: 1, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: RFC2544_TEST_STATE_RUNNING
```

```
Status: Running
Test start time: 2013-12-09 16:24:52 IST
Test finish time: TEST_RUNNING
Counters last cleared: Never

Test Configuration:
  Test mode: Reflect
  Duration in seconds: 864000
  Test family: INET
  Routing Instance Name: default

Inet family Configuration:
  Egress Interface : ge-0/3/1.0
  Destination ipv4 address: 198.51.100.2
  Destination udp port: 200
```

Elapsed time	Reflected Packets	Reflected Bytes
176	8977917	9031784502

show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 1 brief
Test information :
  Test id: 1, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: RFC2544_TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-09 16:24:52 IST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never
```

show services rpm rfc2544-benchmarking test-id detail (Reflection Test on MX104 Routers)

```
user@host> show services rpm rfc2544-benchmarking test-id 2 detail
Test information :
  Test id: 2, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
```

```

Test state: RFC2544_TEST_STATE_RUNNING
Status: Running
Test start time: 2013-12-09 16:39:18 IST
Test finish time: TEST_RUNNING
Counters last cleared: Never

```

Test Configuration:

```

Test mode: Reflect
Duration in seconds: 864000
Test family: CCC
Routing Instance Name: default

```

CCC family Configuration:

```

Interface : ge-0/3/2.0
Test direction: Egress

```

Elapsed time	Reflected Packets	Reflected Bytes
23	809137	825319740

show services rpm rfc2544-benchmarking test-id brief (Reflection Test on MX104 Routers)

```

user@host> show services rpm rfc2544-benchmarking test-id 2 brief
Test information :
  Test id: 2, Test name: fort_uni_inet_ref, Test type: Reflect
  Test mode: Reflect
  Test packet size: 0
  Test state: RFC2544_TEST_STATE_RUNNING
  Status: Running
  Test start time: 2013-12-09 16:39:18 IST
  Test finish time: TEST_RUNNING
  Counters last cleared: Never

```

show services rpm rfc2544-benchmarking test-id detail (Reflection Test on SRX300 and SRX550HM)

```

user@host> show services rpm rfc2544-benchmarking test-id 1 detail
Test information :
  Test id: 1, Test name: ts13, Test type: Reflect

```

```

Test mode: Reflect
Test packet size: 0
Test state: TEST_STATE_RUNNING
Status: Running
Test start time: 2020-07-08 19:53:23 CST
Test finish time: TEST_RUNNING
Counters last cleared: Never

Test Configuration:
  Test mode: Reflect
  Duration in seconds: 864000
  Test finish wait duration in seconds: 1
  Test family: INET
  Test iterator pass threshold: 0.50 %
  Test receive failure threshold: 0.00 %
  Test transmit failure threshold: 0.50 %
  Routing Instance Name: default

Inet family Configuration:
  Egress Interface : ge-0/0/13.0
  Destination ipv4 address: 10.0.2.1
  Destination udp port: 400

```

Elapsed	Reflected	Reflected
time	Packets	Bytes
62053	220259916	134445283434

Release Information

Command introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)

show services rpm twamp client connection

IN THIS SECTION

- [Syntax | 1826](#)
- [Description | 1826](#)
- [Options | 1826](#)
- [Required Privilege Level | 1827](#)
- [Output Fields | 1827](#)
- [Sample Output | 1827](#)
- [Release Information | 1828](#)

Syntax

```
show services rpm twamp client connection  
<connection-name>
```

Description

Display information about the connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control-clients. By default, all established sessions are displayed, unless you specify a control-connection name when you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

- | | |
|-------------------------------|--|
| none | Display information about all TWAMP client connection sessions. |
| <i>connection-name</i> | (Optional) Display information about the specified control-connection or TWAMP control-client. |

Required Privilege Level

view

Output Fields

Table 183 on page 1827 lists the output fields for the `show services rpm twamp client connection` command. Output fields are listed in the approximate order in which they appear.

Table 183: show services rpm twamp client connection Output Fields

Field Name	Field Description
Connection Name	Connection name that uniquely identifies the connection between the TWAMP server and a particular client.
Client address	Client IP address.
Client port	Client port number.
Server address	Server IP address.
Server port	Server port number.
Session count	Session count.
Auth mode	Authentication mode.

Sample Output

`show services rpm twamp client connection` (with a managed TWAMP server configured)

```

user@host> show services rpm twamp client connection

```

Connection ID	Client address	Client port	Server address	Server port	Session count	Auth mode
4	192.0.2.1	12345	192.168.219.203	890	16	none
78	198.51.100.55	345	203.0.113.2	89022	5	none

234	192.168.219.203	2345	192.168.22.2	3333	16	none
5	192.168.3.1	82345	192.168.2.2	45909	16	authenticated
1	192.168.1.1	645	192.168.4.23	2394	16	encrypted

show services rpm twamp client connection (with TWAMP light server configured)

```
user@host> show services rpm twamp client connection c1
error: Control-connection c1 does not exist.
```

Release Information

Command introduced in Junos OS Release 15.1.

show services rpm twamp client history-results

IN THIS SECTION

- [Syntax | 1828](#)
- [Description | 1829](#)
- [Options | 1829](#)
- [Required Privilege Level | 1829](#)
- [Output Fields | 1829](#)
- [Sample Output | 1831](#)
- [Release Information | 1833](#)

Syntax

```
show services rpm twamp client history-results
<brief | detail>
<control-connection control-connection-name>
<since time>
<test-session test-session-name>
```

Description

Display standard information about the results of the last 50 probes for each real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) instance. You can also view the historical results of the probes or test packets sent from a TWAMP client to a TWAMP server for a particular control-connection, or a test-session associated with a control-connection.

Options

none	Display the results of the last 50 probes for all RPM TWAMP instances.
brief detail	(Optional) Display the specified level of output. Default: brief
control-connection <i>control-connection-name</i>	(Optional) Display information for the specified control-connection between a TWAMP client and a TWAMP server.
since <i>time</i>	(Optional) Display information from the specified time. Specify time as <i>yyyy-mm-dd.hh.mm.ss</i> .
test-session <i>test-session-name</i>	(Optional) Display information for the specified test session associated with a control-connection between a TWAMP client and a TWAMP server.

Required Privilege Level

view

Output Fields

[Table 184 on page 1829](#) lists the output fields for the `show services rpm twamp client history-results` command. Output fields are listed in the approximate order in which they appear.

Table 184: show services rpm twamp client history-results Output Fields

Field Name	Field Description	Level of Output
Owner	Probe owner or the TWAMP client.	All levels
Test	Name of a test for a TWAMP probe instance.	All levels

Table 184: show services rpm twamp client history-results Output Fields *(Continued)*

Field Name	Field Description	Level of Output
Probe received	Timestamp when the probe result was determined.	All levels
Round trip time	Average ping round-trip time (RTT), in microseconds.	All levels
Probe results	<p>Result of a particular probe performed by a remote host. The following information is contained in the results:</p> <ul style="list-style-type: none"> • Response received—Timestamp when the probe result was determined. • Rtt—Average ping round-trip time (RTT), in microseconds. 	detail
Results over current test	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	detail
Probes sent	Number of probes sent with the current test.	detail
Probes received	Number of probe responses received within the current test.	detail
Loss percentage	Percentage of lost probes for the current test.	detail

Table 184: show services rpm twamp client history-results Output Fields (*Continued*)

Field Name	Field Description	Level of Output
Measurement	<p>Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-ping-timestamp, the egress and ingress delay:</p> <ul style="list-style-type: none"> • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. • Peak to peak—Difference between two peak values of RTT, ingress delay, or egress delay measured over the course of the current test. • Jitter—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test. • Sum—Total round-trip time, in microseconds, measured over the course of the current test. 	detail

Sample Output

show services rpm twamp client history-results

```
user@host> show services rpm twamp client history-results
```

```
Aug 02 19:11:38
```

Owner, Test	Probe Sent	Probe received	Round trip time
c1, t2	Sun Aug 2 18:26:58 2020	Sun Aug 2 18:26:58 2020	6455 usec
c1, t2	Sun Aug 2 18:26:59 2020	Sun Aug 2 18:26:59 2020	6450 usec
c1, t2	Sun Aug 2 18:27:00 2020	Sun Aug 2 18:27:00 2020	6456 usec
c1, t2	Sun Aug 2 18:27:01 2020	Sun Aug 2 18:27:01 2020	6574 usec
c2, t1	Sun Aug 2 18:59:05 2020	Sun Aug 2 18:59:05 2020	879 usec
c2, t1	Sun Aug 2 18:59:06 2020	Sun Aug 2 18:59:06 2020	6582 usec
c2, t1	Sun Aug 2 18:59:07 2020	Sun Aug 2 18:59:07 2020	7211 usec
c2, t1	Sun Aug 2 18:59:10 2020	Sun Aug 2 18:59:10 2020	6551 usec
c2, t1	Sun Aug 2 18:59:11 2020	Sun Aug 2 18:59:11 2020	6547 usec
c2, t2	Sun Aug 2 18:58:59 2020	Sun Aug 2 18:58:59 2020	6585 usec

c2, t2	Sun Aug 2 18:59:00 2020	Sun Aug 2 18:59:00 2020	6581 usec
c2, t2	Sun Aug 2 18:59:01 2020	Sun Aug 2 18:59:01 2020	6586 usec
c2, t2	Sun Aug 2 18:59:02 2020	Sun Aug 2 18:59:02 2020	6592 usec
c2, t2	Sun Aug 2 18:59:03 2020	Sun Aug 2 18:59:03 2020	6591 usec

show services rpm twamp client history-results detail

```
user@host> show services rpm twamp-client history-results detail
```

```
Owner: p, Test: t
```

```
Probe results:
```

```
Response received, Tue Jan 7 05:11:49 2014,
```

```
Rtt: 184 usec, Round trip jitter: -96 usec, Round trip interarrival jitter: 57 usec
```

```
Results over current test:
```

```
Probes sent: 4, Probes received: 4, Loss percentage: 0
```

```
Measurement: Round trip time
```

```
Samples: 4, Minimum: 174 usec, Maximum: 196 usec, Average: 183 usec, Peak to peak: 22 usec, Stddev: 8 usec, Sum: 732 usec
```

```
Measurement: Positive round trip jitter
```

```
Samples: 1, Minimum: 110 usec, Maximum: 110 usec, Average: 110 usec, Peak to peak: 0 usec, Stddev: 0 usec, Sum: 110 usec
```

```
Measurement: Negative round trip jitter
```

```
Samples: 2, Minimum: 96 usec, Maximum: 811 usec, Average: 454 usec, Peak to peak: 715 usec, Stddev: 358 usec, Sum: 907 usec
```

```
Owner: p, Test: t
```

```
Probe results:
```

```
Response received, Tue Jan 7 05:11:50 2014, Rtt: 174 usec, Round trip jitter: -8 usec, Round trip interarrival jitter: 54 usec
```

```
Results over current test:
```

```
Probes sent: 5, Probes received: 5, Loss percentage: 0
```

```
Measurement: Round trip time
```

```
Samples: 5, Minimum: 174 usec, Maximum: 196 usec, Average: 181 usec, Peak to peak: 22 usec, Stddev: 8 usec, Sum: 906 usec
```

```
Measurement: Positive round trip jitter
```

```
Samples: 1, Minimum: 110 usec, Maximum: 110 usec, Average: 110 usec, Peak to peak: 0 usec, Stddev: 0 usec, Sum: 110 usec
```

```
Measurement: Negative round trip jitter
```

```
Samples: 3, Minimum: 8 usec, Maximum: 811 usec, Average: 305 usec, Peak to peak: 803 usec, Stddev: 360 usec, Sum: 915 usec
```

Release Information

Command introduced in Junos OS Release 15.1.

show services rpm twamp client probe-results

IN THIS SECTION

- [Syntax | 1833](#)
- [Description | 1833](#)
- [Options | 1833](#)
- [Required Privilege Level | 1834](#)
- [Output Fields | 1834](#)
- [Sample Output | 1840](#)
- [Release Information | 1844](#)

Syntax

```
show services rpm twamp client probe-results  
<control-connection control-connection-name>  
<test-session test-session-name>
```

Description

Display the results of the most recent real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) probes sent from the TWAMP client to the TWAMP server. You can also view the results of the probes or test packets sent from a TWAMP client to a TWAMP server for a particular control-connection, or a test-session associated with a control-connection.

Options

none	Display all results of the most recent TWAMP probes.
-------------	--

- control-connection** (Optional) Display information for the specified control-connection between a TWAMP client and a TWAMP server.
- control-connection-name**
- test-session** (Optional) Display information for the specified test session associated with a control-connection between a TWAMP client and a TWAMP server.
- test-session-name**

Required Privilege Level

view

Output Fields

Table 185 on page 1834 lists the output fields for the `show services twamp client probe-results` command. Output fields are listed in the approximate order in which they appear.

Table 185: show services twamp client probe-results Output Fields

Field Name	Field Description
Owner	Name of the session-sender or the control-client, which is the TWAMP client. When you configure the <code>control-client-name</code> option at the [edit services rpm twamp client control-connection <code>control-connection-name</code>] hierarchy level, this field displays the configured owner name or the client name.
Test	Name of a test representing a collection of probes. When you configure the test-session <code>test-name</code> statement at the [edit services rpm twamp client control-connection <code>control-connection-name</code>] hierarchy level, the field displays the configured test name.
server-address	Destination address used for the probes.
server-port	Destination port used for the probes.
Client address	Source or TWAMP client address used for the probes.
Client port	Source or TWAMP client port used for the probes.

Table 185: show services twamp client probe-results Output Fields (Continued)

Field Name	Field Description
TWAMP-Server-Status	<p>Possible values:</p> <ul style="list-style-type: none"> • Light: the control-type light statement is configured for the control connection. • Connected or Not Connected: for control connections configured with the control-type managed statement, displays whether or not the connection is up.
Number-of-Retries-with-TWAMP-Server	The number of times the system has tried to connect to the TWAMP server.
Reflector address	Session-reflector or TWAMP server address used for the probes.
Reflector port	Session-reflector or TWAMP server port used for the probes.
Sender address	Session-sender or TWAMP client address used for the probes.
Sender-port	Session-sender or TWAMP client port used for the probes.
Local-link	Egress logical interface for the link-local address
Routing Instance Name	Name configured on the routing-instance statement for the control connection.
Probe type	Protocol configured on the receiving probe server: http-get, http-metadata-get, icmp-ping, icmp-ping-timestamp, tcp-ping, udp-ping, or udp-ping-timestamp.
Test size	Number of probes within a test.
Destination Interface Name	Name of the interface configured on the TWAMP server or the session-reflector on which the TWAMP probe packets sent from the TWAMP client are received.

Table 185: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description
Probe results	<p>Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement:</p> <ul style="list-style-type: none"> • Response received—Timestamp when the probe result was determined. • Client and server hardware timestamps—If timestamps are configured, an entry appears at this point. • Rtt—Average ping round-trip time (RTT), in microseconds. • Egress jitter—Egress jitter, in microseconds. • Ingress jitter—Ingress jitter, in microseconds. • Round trip jitter—Round-trip jitter, in microseconds. • Egress interarrival jitter—Egress interarrival jitter, in microseconds. • Ingress interarrival jitter—Ingress interarrival jitter, in microseconds. • Round trip interarrival jitter—Round-trip interarrival jitter, in microseconds.

Table 185: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description
Results over current test	<p>Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes.</p> <ul style="list-style-type: none"> Probes sent—Number of probes sent within the current test. Probes received—Number of probe responses received within the current test. Loss percentage—Percentage of lost probes for the current test. Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> Samples—Number of probes. Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. Peak to peak—Peak-to-peak difference, in microseconds. Stddev—Standard deviation, in microseconds. Sum—Statistical sum.

Table 185: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description
Results over last test	<p>Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent for the most recently completed test. • Probes received—Number of probe responses received for the most recently completed test. • Loss percentage—Percentage of lost probes for the most recently completed test. • Test completed—Time the most recent test was completed. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type icmp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test. • Maximum—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test. • Average—Average RTT, ingress delay, or egress delay measured for the most recently completed test. • Peak to peak—Peak-to-peak difference, in microseconds. • Stddev—Standard deviation, in microseconds. • Sum—Statistical sum.

Table 185: show services twamp client probe-results Output Fields *(Continued)*

Field Name	Field Description
Results over all tests	<p>Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test.</p> <ul style="list-style-type: none"> • Probes sent—Number of probes sent in all tests. • Probes received—Number of probe responses received in all tests. • Loss percentage—Percentage of lost probes in all tests. • Measurement—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types icmp-ping-timestamp and udp-ping-timestamp, the egress delay and ingress delay. <p>For each measurement type, the following individual calculated results are provided:</p> <ul style="list-style-type: none"> • Samples—Number of probes. • Minimum—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • Maximum—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • Average—Average RTT, ingress delay, or egress delay measured over the course of the current test. • Peak to peak—Peak-to-peak difference, in microseconds. • Stddev—Standard deviation, in microseconds. • Sum—Statistical sum.
Error stats	Error statistics

Sample Output

show services rpm twamp client probe-results

```

user@host> show services rpm twamp client probe-results
Owner: c1, Test: t1
  TWAMP-Server-Status: Light,    Number-Of-Retries-With-TWAMP-Server: 0
  Reflector address: 10.44.6.22, Reflector port: 862,  Sender address: 10.2.2.1, sender-port:
64109
  Routing Instance Name: IN
  Test size: 100 probes
  Probe results:
    Response received
    Probe sent time: Fri Mar  5 09:59:34 2021
    Probe rcvd/timeout time: Fri Mar  5 09:59:34 2021
    Rtt: 1728 usec, Egress jitter: -3082 usec, Ingress jitter: 742 usec, Round trip jitter:
-339 usec
    Egress interarrival jitter: 1047 usec, Ingress interarrival jitter: 826 usec, Round trip
interarrival jitter: 1139 usec
  Results over current test:
    Probes sent: 11, Probes received: 11, Loss percentage: 0.000000
    Measurement: Round trip time
      Samples: 11, Minimum: 1135 usec, Maximum: 4445 usec, Average: 2994 usec, Peak to peak:
3310 usec, Stddev: 1039 usec, Sum: 32938 usec
    Measurement: Positive egress jitter
      Samples: 7, Minimum: 17 usec, Maximum: 3101 usec, Average: 1426 usec, Peak to peak: 3084
usec, Stddev: 958 usec, Sum: 9981 usec
    Measurement: Negative egress jitter
      Samples: 4, Minimum: 1317 usec, Maximum: 3082 usec, Average: 2205 usec, Peak to peak:
1765 usec, Stddev: 868 usec, Sum: 8821 usec
    Measurement: Positive ingress jitter
      Samples: 3, Minimum: 742 usec, Maximum: 3306 usec, Average: 2222 usec, Peak to peak:
2564 usec, Stddev: 1083 usec, Sum: 6665 usec
    Measurement: Negative ingress jitter
      Samples: 8, Minimum: 3 usec, Maximum: 3198 usec, Average: 791 usec, Peak to peak: 3195
usec, Stddev: 1122 usec, Sum: 6325 usec
    Measurement: Positive round trip jitter
      Samples: 4, Minimum: 650 usec, Maximum: 4964 usec, Average: 2580 usec, Peak to peak:
4314 usec, Stddev: 1551 usec, Sum: 10319 usec
    Measurement: Negative round trip jitter
      Samples: 7, Minimum: 230 usec, Maximum: 2952 usec, Average: 919 usec, Peak to peak: 2722
usec, Stddev: 923 usec, Sum: 6430 usec

```

Results over last test:

Probes sent: 100, Probes received: 100, Loss percentage: 0.000000

Test completed on Fri Mar 5 09:58:50 2021

Measurement: Round trip time

Samples: 100, Minimum: 192 usec, Maximum: 5425 usec, Average: 635 usec, Peak to peak: 5233 usec, Stddev: 948 usec, Sum: 63507 usec

Measurement: Positive egress jitter

Samples: 75, Minimum: 9 usec, Maximum: 1810 usec, Average: 223 usec, Peak to peak: 1801 usec, Stddev: 346 usec, Sum: 16735 usec

Measurement: Negative egress jitter

Samples: 25, Minimum: 23 usec, Maximum: 3871 usec, Average: 603 usec, Peak to peak: 3848 usec, Stddev: 805 usec, Sum: 15064 usec

Measurement: Positive ingress jitter

Samples: 19, Minimum: 1 usec, Maximum: 3710 usec, Average: 837 usec, Peak to peak: 3709 usec, Stddev: 922 usec, Sum: 15909 usec

Measurement: Negative ingress jitter

Samples: 81, Minimum: 4 usec, Maximum: 3834 usec, Average: 274 usec, Peak to peak: 3830 usec, Stddev: 577 usec, Sum: 22206 usec

Measurement: Positive round trip jitter

Samples: 48, Minimum: 5 usec, Maximum: 4382 usec, Average: 778 usec, Peak to peak: 4377 usec, Stddev: 1077 usec, Sum: 37365 usec

Measurement: Negative round trip jitter

Samples: 52, Minimum: 3 usec, Maximum: 4915 usec, Average: 787 usec, Peak to peak: 4912 usec, Stddev: 1239 usec, Sum: 40922 usec

Results over all tests:

Probes sent: 38611, Probes received: 38581, Loss percentage: 0.077698

Measurement: Round trip time

Samples: 38556, Minimum: 2 usec, Maximum: 22356 usec, Average: 3275 usec, Peak to peak: 22354 usec, Stddev: 3272 usec, Sum: 126256294 usec

Measurement: Positive egress jitter

Samples: 23770, Minimum: 0 usec, Maximum: 11266 usec, Average: 1054 usec, Peak to peak: 11266 usec, Stddev: 1496 usec, Sum: 25053834 usec

Measurement: Negative egress jitter

Samples: 14781, Minimum: 1 usec, Maximum: 9994 usec, Average: 1533 usec, Peak to peak: 9993 usec, Stddev: 1609 usec, Sum: 22659276 usec

Measurement: Positive ingress jitter

Samples: 12989, Minimum: 0 usec, Maximum: 12438 usec, Average: 1639 usec, Peak to peak: 12438 usec, Stddev: 1673 usec, Sum: 21293431 usec

Measurement: Negative ingress jitter

Samples: 25562, Minimum: 1 usec, Maximum: 13883 usec, Average: 927 usec, Peak to peak: 13882 usec, Stddev: 1458 usec, Sum: 23689224 usec

Measurement: Positive round trip jitter

Samples: 19298, Minimum: 0 usec, Maximum: 21037 usec, Average: 1793 usec, Peak to peak:

```

21037 usec, Stddev: 2146 usec, Sum: 34594283 usec
  Measurement: Negative round trip jitter
    Samples: 19257, Minimum: 1 usec, Maximum: 21288 usec, Average: 1796 usec, Peak to peak:
21287 usec, Stddev: 2148 usec, Sum: 34589725 usec
  Error Stats:
    Invalid client recv timestamp ( $T4 < T1$ ): 1, Invalid client send( $T1$ ), recv( $T4$ ) timestamps: 0
    Invalid server send timestamp ( $T3 < T2$ ): 1, Invalid server processing time ( $T4 - T1$ ) < ( $T3 - T2$ ): 23

```

show services rpm twamp client probe-results (with local-link address information)

```

user@host> show services rpm twamp client probe-results
Owner: tc7, Test: ts1
  TWAMP-Server-Status: Light, Number-Of-Retries-With-TWAMP-Server: 0
  Reflector address: 2001:db8::9f16, Reflector port: 862, Sender address: 2001:db8::167d,
sender-port: 54399
  Local-link: ge-2/2/9.1
  Test size: 1 probes
  Probe results:
    Response received
    Probe sent time: Mon Aug 30 10:30:51 2021
    Probe rcvd/timeout time: Mon Aug 30 10:30:51 2021
    Rtt: 264 usec, Egress jitter: -66 usec, Ingress jitter: -14 usec, Round trip jitter: -59
usec
    Egress interarrival jitter: 73 usec, Ingress interarrival jitter: 45 usec, Round trip
interarrival jitter: 111 usec
  Results over current test:
    Probes sent: 1, Probes received: 1, Loss percentage: 0.000000
    Measurement: Round trip time
      Samples: 1, Minimum: 264 usec, Maximum: 264 usec, Average: 264 usec, Peak to peak: 0
usec, Stddev: 0 usec, Sum: 264 usec
    Measurement: Negative egress jitter
      Samples: 1, Minimum: 66 usec, Maximum: 66 usec, Average: 66 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 66 usec
    Measurement: Negative ingress jitter
      Samples: 1, Minimum: 14 usec, Maximum: 14 usec, Average: 14 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 14 usec
    Measurement: Negative round trip jitter
      Samples: 1, Minimum: 59 usec, Maximum: 59 usec, Average: 59 usec, Peak to peak: 0 usec,
Stddev: 0 usec, Sum: 59 usec

```

Results over last test:

Probes sent: 1, Probes received: 1, Loss percentage: 0.000000

Test completed on Mon Aug 30 10:30:51 2021

Measurement: Round trip time

Samples: 1, Minimum: 264 usec, Maximum: 264 usec, Average: 264 usec, Peak to peak: 0 usec, Stddev: 0 usec, Sum: 264 usec

Measurement: Negative egress jitter

Samples: 1, Minimum: 66 usec, Maximum: 66 usec, Average: 66 usec, Peak to peak: 0 usec, Stddev: 0 usec, Sum: 66 usec

Measurement: Negative ingress jitter

Samples: 1, Minimum: 14 usec, Maximum: 14 usec, Average: 14 usec, Peak to peak: 0 usec, Stddev: 0 usec, Sum: 14 usec

Measurement: Negative round trip jitter

Samples: 1, Minimum: 59 usec, Maximum: 59 usec, Average: 59 usec, Peak to peak: 0 usec, Stddev: 0 usec, Sum: 59 usec

Results over all tests:

Probes sent: 46059, Probes received: 46059, Loss percentage: 0.000000

Measurement: Round trip time

Samples: 46025, Minimum: 209 usec, Maximum: 4730 usec, Average: 361 usec, Peak to peak: 4521 usec, Stddev: 351 usec, Sum: 16593381 usec

Measurement: Positive egress jitter

Samples: 14538, Minimum: 0 usec, Maximum: 4319 usec, Average: 73 usec, Peak to peak: 4319 usec, Stddev: 275 usec, Sum: 1056470 usec

Measurement: Negative egress jitter

Samples: 31486, Minimum: 1 usec, Maximum: 4559 usec, Average: 40 usec, Peak to peak: 4558 usec, Stddev: 191 usec, Sum: 1251936 usec

Measurement: Positive ingress jitter

Samples: 34378, Minimum: 0 usec, Maximum: 4391 usec, Average: 59 usec, Peak to peak: 4391 usec, Stddev: 386 usec, Sum: 2029749 usec

Measurement: Negative ingress jitter

Samples: 11646, Minimum: 1 usec, Maximum: 4396 usec, Average: 158 usec, Peak to peak: 4395 usec, Stddev: 652 usec, Sum: 1834338 usec

Measurement: Positive round trip jitter

Samples: 23383, Minimum: 0 usec, Maximum: 10709 usec, Average: 276 usec, Peak to peak: 10709 usec, Stddev: 772 usec, Sum: 6454130 usec

Measurement: Negative round trip jitter

Samples: 22641, Minimum: 1 usec, Maximum: 11233 usec, Average: 285 usec, Peak to peak: 11232 usec, Stddev: 788 usec, Sum: 6454183 usec

Error Stats:

Invalid client recv timestamp ($T4 < T1$): 11, Invalid client send($T1$), recv($T4$) timestamps:

0

Invalid server send timestamp ($T3 < T2$): 3, Invalid server processing time ($T4 - T1 < (T3 - T2)$): 20

show services rpm twamp client probe-results control-connection

```
user@host> show services rpm twamp client probe-results control-connection c2
Owner: c2, Test: t1
server-address: 192.0.2.14, server-port: 862, Client address: 192.0.2.13,
Client port: 57170
Reflector address: 192.0.2.14, Reflector port: 10010,
Sender address: 192.0.2.13, sender-port: 10010
Destination interface name: si-1/1/0.10
Test size: 500 probes
Probe results:
  Request timed out, Fri Feb 13 00:07:14 2015
Results over current test:
  Probes sent: 188, Probes received: 0, Loss percentage: 100.00000
Results over last test:
  Probes sent: 500, Probes received: 0, Loss percentage: 100.00000
Results over all tests:
  Probes sent: 3688, Probes received: 0, Loss percentage: 100.00000
```

Release Information

Command introduced in Junos OS Release 15.1.

Local-link field added in Junos OS Release 21.4R1.

show services rpm twamp client session

IN THIS SECTION

- [Syntax | 1845](#)
- [Description | 1845](#)
- [Options | 1845](#)

- [Required Privilege Level | 1845](#)
- [Output Fields | 1846](#)
- [Sample Output | 1846](#)
- [Release Information | 1847](#)

Syntax

```
show services rpm twamp client session
<control-connection control-connection-name>
<test-session test-session-name>
```

Description

Display information about the sessions established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients for control packets and data packets. By default, all established control-connection and data-connection or test sessions are displayed, unless you specify a control-connection name or a test-session name when you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

none	Display information about all established connections and sessions.
control-connection <i>control-connection-name</i>	(Optional) Display information about the specified control-connection, which is established for control-packets exchanged between a TWAMP client and a TWAMP server.
test-session <i>test-session-name</i>	(Optional) Display information about the specified test session, which is established for data packets transmitted between a TWAMP client and a TWAMP server, associated with a control-connection..

Required Privilege Level

view

Output Fields

[Table 186 on page 1846](#) lists the output fields for the `show services rpm twamp client session` command. Output fields are listed in the approximate order in which they appear.

Table 186: show services rpm twamp client session Output Fields

Field Name	Field Description
Connection Name	Name of the control connection that uniquely identifies the connection between the TWAMP server and the TWAMP client.
Session Name	Name of the test session that uniquely identifies the data-session between the TWAMP server and the TWAMP client.
Sender address	Sender IP address.
Sender port	Sender port number.
Reflector address	Reflector IP address.
Reflector port	Reflector port number.

Sample Output

show services rpm twamp client session

```
user@host> show services rpm twamp client session
```

Connection	Session	Sender	Sender	Reflector	Reflector
Name	Name	address	port	address	port
cs1	ts1	198.51.100.1	41998	198.51.100.2	5008
cs2	ts1	198.51.100.1	53710	198.51.100.2	5009

show services rpm twamp client session control-connection

```
user@host> show services rpm twamp client session control-connection c2
```

Connection	Session	Sender	Sender	Reflector	Reflector
Name	Name	address	port	address	port
c2	t1	192.0.2.13	10008	192.0.2.14	10008

Release Information

Command introduced in Junos OS Release 15.1.

show services rpm twamp server connection

IN THIS SECTION

- [Syntax | 1847](#)
- [Description | 1848](#)
- [Options | 1848](#)
- [Required Privilege Level | 1848](#)
- [Output Fields | 1848](#)
- [Sample Output | 1849](#)
- [Release Information | 1849](#)

Syntax

```
show services rpm twamp server connection  
<connection-id>
```

Description

Display information about the connections established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control-clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command. Because TWAMP light servers are stateless, information about them is not included in the output of this command; only information about managed servers is included.

Options

- none** Display connection information about all established sessions.
- connection-id*** (Optional) Identifier of the connection that you want to display information about.

Required Privilege Level

view

Output Fields

[Table 187 on page 1848](#) lists the output fields for the `show services rpm twamp server connection` command. Output fields are listed in the approximate order in which they appear.

Table 187: show services rpm twamp server connection Output Fields

Field Name	Field Description
Connection ID	Connection ID that uniquely identifies the connection between the TWAMP server and a particular client.
Client address	Client IP address.
Client port	Client port number.
Server address	Server IP address.
Server port	Server port number.

Table 187: show services rpm twamp server connection Output Fields *(Continued)*

Field Name	Field Description
Session count	Session count.
Auth mode	Authentication mode.

Sample Output

show services rpm twamp server connection

```
user@host> show services rpm twamp server connection
```

Connection ID	Client address	Client port	Server address	Server port	Session count	Auth mode
4	192.0.2.1	12345	192.168.219.203	890	16	none
78	198.51.100.55	345	203.0.113.2	89022	5	none
234	192.168.219.203	2345	192.168.22.2	3333	16	none
5	192.168.3.1	82345	192.168.2.2	45909	16	authenticated
1	192.168.1.1	645	192.168.4.23	2394	16	encrypted

Release Information

Command introduced in Junos OS Release 9.3.

show services rpm twamp server session

IN THIS SECTION

- [Syntax | 1850](#)
- [Description | 1850](#)
- [Options | 1850](#)
- [Required Privilege Level | 1850](#)

- [Output Fields | 1850](#)
- [Sample Output | 1851](#)
- [Release Information | 1851](#)

Syntax

```
show services rpm twamp server session
<session-id>
```

Description

Display information about the sessions established between the real-time performance monitoring (RPM) Two-Way Active Measurement Protocol (TWAMP) server and control clients. By default, all established sessions are displayed, unless you specify a session ID when you issue the command.

Options

- none** Display information about all established sessions.
- session-id*** (Optional) Identifier of the session that you want to display information about.

Required Privilege Level

view

Output Fields

[Table 188 on page 1851](#) lists the output fields for the `show services rpm twamp server session` command. Output fields are listed in the approximate order in which they appear.

Table 188: show services rpm twamp server session Output Fields

Field Name	Field Description
Session ID	Session ID that uniquely identifies the session between the TWAMP server and a particular client.
Connection ID	Connection ID that uniquely identifies the connection between the TWAMP server and a particular client.
Sender address	Sender IP address.
Sender port	Sender port number.
Reflector address	Reflector IP address.
Reflector port	Reflector port number.

Sample Output

show services rpm twamp server session

```

user@host> show services rpm twamp server session
  Session  Connection  Sender      Sender  Reflector  Reflector
  ID       ID          address     port    address    port
  4        44        192.0.2.1   12345   192.168.219.203  890
  78       44        198.51.100.55  345    203.0.113.2    89022
  234     423     192.168.219.203 2345   192.168.22.2    3333
  5       423     192.168.3.1    82345   192.168.2.2     45909
  1       423     192.168.1.1    645     192.168.4.23    2394

```

Release Information

Command introduced in Junos OS Release 9.3.

show services service-sets statistics jflow-log

IN THIS SECTION

- [Syntax | 1852](#)
- [Description | 1852](#)
- [Options | 1852](#)
- [Required Privilege Level | 1853](#)
- [Output Fields | 1853](#)
- [Sample Output | 1856](#)
- [Release Information | 1862](#)

Syntax

```
show services service-sets statistics jflow-log
<interface interface-name>
<service-set service-set-name>
<brief | detail>
```

Description

Display statistical information on the logs or records generated in flow monitoring format with optional filtering by interface and service set name..

Options

none	Display the statistical details on flow monitoring logs for NAT events for all services interfaces and all service sets.
brief	(Default) (Optional) Display abbreviated flow monitoring log statistics.
detail	(Optional) Display detailed flow monitoring log statistics.

- interface**
interface-name (Optional) Display the flow monitoring log statistics for the specified adaptive service interface. On M Series and T Series routers, *interface-name* can be *ms-fpc/pic/port*. It is supported only on MS-MICs and MS-MPCs.
- service-set**
service-set name (Optional) Display the flow monitoring log statistics for the specified named service-set.

Required Privilege Level

view

Output Fields

[Table 189 on page 1853](#) lists the output fields for the `show services service-sets statistics jflow-log` command. Output fields are listed in the approximate order in which they appear.

Table 189: show services service-sets statistics jflow-log Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a services interface.	all
Rate limit	Maximum number of NAT error events for which records in flow monitoring format must be recorded.	all
Template records	Details of the template records in flow monitoring log messages.	all
Sent	Number of template records sent to a collector	all
Messages dropped	Number of template records dropped while transmission to a collector.	all
Data records	Details of the data records in flow monitoring log messages.	all
Sent	Number of data records sent to a collector.	all

Table 189: show services service-sets statistics jflow-log Output Fields (Continued)

Field Name	Field Description	Level of Output
Dropped	Number of data records dropped while transmission to a collector	all
Service set	Name of a service set.	all
Unresolvable collectors	Number of collectors that cannot be traced and be reached to export records for NAT events.	all

Table 189: show services service-sets statistics jflow-log Output Fields (Continued)

Field Name	Field Description	Level of Output
class name	<p>Logs created for events for each of the following classes:</p> <ul style="list-style-type: none"> • NAT44 Session logs—Details of logs created for NAT44 sessions • NAT64 Session logs—Details of logs created for NAT64 sessions • NAT44 BIB logs—Details of logs created for NAT44 binding information bases, which is a table of bindings kept by a NAT44. Each NAT44 has a BIB for each translated protocol. • NAT64 BIB logs—Details of logs created for NAT44 binding information bases, which is a table of bindings kept by a NAT64. Each NAT64 has a BIB for each translated protocol. • NAT Address Exhausted logs—Details of logs created for exhaustion of NAT addresses • NAT Port Exhausted logs—Details of logs created for exhaustion of NAT pool • NAT44 Quota Exceeded logs—Details of logs created when allocated quota is exceeded for NAT44 events • NAT64 Quota Exceeded logs—Details of logs created when allocated quota is exceeded for NAT64 events • NAT44 Address Bind logs—Details of logs generated for address bindings for NAT44 events • NAT64 Address Bind logs—Details of logs generated for address bindings for NAT64 events • NAT44 PBA logs—Details of logs generated for NAT44 port block allocation events • NAT64 PBA logs—Details of logs generated for NAT64 port block allocation events <p>The following information is displayed for flow monitoring log messages for each class of event that is logged:</p>	detail

Table 189: show services service-sets statistics jflow-log Output Fields (*Continued*)

Field Name	Field Description	Level of Output
	<ul style="list-style-type: none"> • Template records—Details of the template records in flow monitoring log messages • Sent—Number of template records sent to a collector • Dropped—Number of template records dropped while transmission to a collector • Data records—Details of the data records in flow monitoring log messages • Sent—Number of data records sent to a collector • Dropped—Number of data records dropped while transmission to a collector. Counts are provided for the drop reasons • socket send error—Number of times a socket was not opened for data transmission • no memory—Number of messages dropped because of insufficient memory • invalid data—Number of messages dropped because of invalid data in the records • above rate limit—The maximum number of flow monitoring log messages per second was exceeded. 	

Sample Output

show services service-sets statistics jflow-log brief

```

user@host> show services service-sets statistics jflow-log brief
Interface: ms-5/0/0
Rate limit: 1000
Template records:
  Sent: 36
  Dropped: 0
Data records:

```

```

Sent: 2
Dropped: 0

Service-set: sset_44
  Unresolvable collectors: 0
  Template records:
    Sent: 36
    Dropped: 0
  Data records:
    Sent: 2
    Dropped: 0

```

show services service-sets statistics jflow-log detail

```

user@host> show services service-sets statistics jflow-log detail
Interface: ms-5/0/0
  Rate limit: 1000
  Template records:
    Sent: 48
    Dropped: 0
  Data records:
    Sent: 4
    Dropped: 0

Service-set: sset_44
  Unresolvable collectors: 0
  Template records:
    Sent: 48
    Dropped: 0
  Data records:
    Sent: 4
    Dropped: 0
  NAT44 Session logs:
    Template records:
      Sent: 4
      Dropped: 0 (socket send error: 0, no memory: 0)
    Data records:
      Sent: 4
      Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
  NAT64 Session logs:

```

```
Template records:
  Sent: 4
  Dropped: 0 (socket send error: 0, no memory: 0)
Data records:
  Sent: 0
  Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 BIB logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 BIB logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Address Exhausted logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Port Exhausted logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Quota Exceeded logs:
  Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 Quota Exceeded logs:
  Template records:
```

```

    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Address Bind logs:
Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 Address Bind logs:
Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 PBA logs:
Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 PBA logs:
Template records:
    Sent: 4
    Dropped: 0 (socket send error: 0, no memory: 0)
Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)

```

show services service-sets statistics jflow-log service-set

```

user@host> show services service-sets statistics jflow-log service-set sset_44
Interface: ms-5/0/0

Service-set: sset_44

```

```

Unresolvable collectors: 0
Template records:
  Sent: 72
  Dropped: 0
Data records:
  Sent: 4
  Dropped: 0

```

show services service-sets statistics jflow-log service-set detail

```

user@host> show services service-sets statistics jflow-log service-set sset_44 detail
Interface: ms-5/0/0

```

```

Service-set: sset_44
  Unresolvable collectors: 0
  Template records:
    Sent: 84
    Dropped: 0
  Data records:
    Sent: 4
    Dropped: 0
  NAT44 Session logs:
    Template records:
      Sent: 7
      Dropped: 0 (socket send error: 0, no memory: 0)
    Data records:
      Sent: 4
      Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
  NAT64 Session logs:
    Template records:
      Sent: 7
      Dropped: 0 (socket send error: 0, no memory: 0)
    Data records:
      Sent: 0
      Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
  NAT44 BIB logs:
    Template records:
      Sent: 7
      Dropped: 0 (socket send error: 0, no memory: 0)
    Data records:
      Sent: 0

```

```
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 BIB logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Address Exhausted logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT Port Exhausted logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Quota Exceeded logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 Quota Exceeded logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 Address Bind logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
```



```

NAT64 Address Bind logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT44 PBA logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)
NAT64 PBA logs:
  Template records:
    Sent: 7
    Dropped: 0 (socket send error: 0, no memory: 0)
  Data records:
    Sent: 0
    Dropped: 0 (invalid data: 0, no memory: 0, above rate limit: 0)

```

Release Information

Command introduced in Junos OS Release 15.1.

RELATED DOCUMENTATION

clear services service-sets statistics syslog

show services video-monitoring mdi errors fpc-slot

IN THIS SECTION

- [Syntax | 1863](#)
- [Description | 1863](#)

- [Options | 1863](#)
- [Required Privilege Level | 1863](#)
- [Output Fields | 1863](#)
- [Sample Output | 1864](#)
- [Release Information | 1864](#)

Syntax

```
show services video-monitoring mdi errors fpc-slot fpc-slot
```

Description

Display video monitoring error statistics.

Options

fpc-slot Number of the fpc slot for which statistics are displayed.

Required Privilege Level

view

Output Fields

[Table 190 on page 1863](#) lists the output fields for the `show services video-monitoring mdi errors fpc-slot` command. Output fields are listed in the approximate order in which they appear.

Table 190: show services video-monitoring mdi errors fpc-slot Output Fields

Field Name	Field Description
FPC slot	Slot number of the monitored FPC.

Table 190: show services video-monitoring mdi errors fpc-slot Output Fields (Continued)

Field Name	Field Description
Flow Insert Error	Number of errors during new flow insert operations.
Flow Policer Drops	<p>Number of packets dropped by flow policer process.</p> <p>NOTE: New flows usually arrive within a very short time interval (1.5 microseconds). These errors do not represent the loss of entire flows, because subsequent packets in the flow can establish the flow. All packets are monitored after a flow has been established. Packet forwarding occurs independently of the video monitoring, and packets are not dropped due to video monitoring errors.</p>
Unsupported Media Packets Count	Number of packets dropped because they are not media packets or they are unsupported media packets.
PID Limit Exceeded	<p>Number of packets unmonitored because the process identifier (PID) limit exceeded has been exceeded.</p> <p>NOTE: The current PID limit is 6.</p>

Sample Output

show services video-monitoring mdi errors fpc-slot

```

user@host> show services video-monitoring mdi errors fpc-slot 2
MDI Errors Information
  FPC Slot: 2
  Flow Insert Error: 0, Flow Policer Drops: 0
  Unsupported Media Packets Count: 0, PID Limit Exceeded: 202995

```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers](#) | 927

show services video-monitoring mdi flows fpc-slot

IN THIS SECTION

- [Syntax](#) | 1865
- [Description](#) | 1866
- [Options](#) | 1866
- [Required Privilege Level](#) | 1867
- [Output Fields](#) | 1867
- [Sample Output](#) | 1869
- [Sample Output](#) | 1869
- [Sample Output](#) | 1870
- [Release Information](#) | 1872

Syntax

```
show services video-monitoring mdi flows fpc-slot fpc-slot  
<brief>  
<count>  
<destination-address>  
<destination-port>  
<detail>  
<df-mlr-split-view>  
<flow-over-ipv4 | flow-over-ipv4-over-mpls | flow-over ipv6 | flow-over-ipv6-mpls>  
<input>  
<interface-name>  
<output>  
<rtp>  
<source-address>  
<source-port>
```

```
<template-name>
<udp>
```

Description

Display inline video monitoring flow statistics.

Options

<i>fpc-slot</i>	Number of the slot for which flows are reported.
brief	(Optional) Display brief output (default).
count	(Optional) Display the number of flows.
destination-address	(Optional) Filter output by destination address.
destination-port	(Optional) Filter output by destination port.
detail	(Optional) Display output in detailed format including media delivery index records.
df-mlr-split-view	(Optional) Display detailed/brief flow output with DF and MLR split.
flow-over-ipv4 flow-over-ipv4-over-mpls flow-over-ipv6 flow-over-ipv6-mpls	(Optional) Display only IPv4 flows, only IPv4-over-MPLS flows, only IPv6 flows, or only IPv6-over-MPLS flows.
input	(Optional) Filter output by flow direction input.
interface-name	(Optional) Filter output by logical interface name.
output	(Optional) Filter output by flow direction output.
rtp	(Optional) Filter output by flow type rtp.
source-address	(Optional) Filter output by source IP address.
source-port	(Optional) Filter output by source port.
template-name	(Optional) Filter output by media delivery index template name.
udp	(Optional) Filter output by flow type MPEG-TS.
df-mlr-split-view	(Optional) Display detailed/brief flow output with DF and MLR split.

Required Privilege Level

view

Output Fields

Table 191 on page 1867 lists the output fields for the show services video-monitoring mdi flows fpc-slot command. Output fields are listed in the approximate order in which they appear.

Table 191: show services video-monitoring mdi flows fpc-slot Output Fields

Field Name	Field Description
SIP	Source IP address.
DIP	Destination IP address.
SP	Source port.
DP	Destination port.
Di	Direction of flow (I=Input, O=Output).
Ty	Type of flow.
Last DF:MLR	<p>Delay factor and media loss rate value of last media delivery index record.</p> <p>NOTE: If you choose df-mlr-split-view option, then the DF and MLR values will be displayed in split format as:</p> <p>Last DF: <xxx>, Last MLR: <xxx>, Avg DF: <xxx>, Avg MLR: <xxx></p>
Avg DF:MLR	Average value of delay factor and media loss rate.
Last MRV	Media rate variation value of last media delivery index record.
Avg MRV	Average value of media rate variation.

Table 191: show services video-monitoring mdi flows fpc-slot Output Fields (Continued)

Field Name	Field Description
IFL	Interface name on which flow is received.
Template Name	Name of template associated with flow.
Flow Identifier	Identifier for the flow.
Source Address	Source IP address.
Destination Address	Destination IP address.
Interface Name	Interface name on which flow is received.
Flow Direction	Direction of flow (I=Input, O=Output).
Flow Type	<p>Flow type that is monitored. The value is one of the following:</p> <ul style="list-style-type: none"> • RTP • RTP over IPv6 • RTP over IPv4-over-MPLS • RTP over IPv6-over-MPLS • UDP • UDP over IPv6 • UDP over IPv4-over-MPLS • UDP over IPv6-over-MPLS
Rec No	Record number.
PID	Process identifier.

Table 191: show services video-monitoring mdi flows fpc-slot Output Fields (Continued)

Field Name	Field Description
MLR	Media loss rate.

Sample Output

show services video-monitoring mdi flows fpc-slot brief

```
user@host> show services video-monitoring mdi flows fpc-slot 2 brief
-----
Sno |SIP      |SP |DIP      |DP |Di|Ty |Last DF:MLR |Avg DF:MLR |Last MRV |Avg MRV |
IFL      |Template Name |Flow Identifier
-----
1   192.0.2.2 1024 198.51.100.2 2048 I  UDP  70.90:1      92.15:8205  -7.09     -9.36
xe-2/2/1.0      t1          16777216
```

Sample Output

show services video-monitoring mdi flows fpc-slot 0 detail

```
user@host> show services video-monitoring mdi flows fpc-slot 0 detail
Source Address: 192.0.2.22, Source Port: 1024
Destination Address: 203.0.113.1, Destination Port: 2060
Last DF:MLR: 3.56:0, Avg DF:MLR: 3.60:0
Last MRV: 0.00, Avg MRV: 0.00
Interface Name: ge-0/3/4.0, Template Name: t1
Flow Direction: Input, Flow Type: RTP
MDI Records Count: 10
Flow Identifier: 16777216
-----+
Rec No|      DF|      MLR|      MRV|
-----+
```


1	3.61	0	0.00
2	3.64	0	0.00
3	3.58	0	0.00
4	3.62	0	0.00
5	3.57	0	0.00
6	3.60	0	0.00
7	3.63	0	0.00
8	3.58	0	0.00
9	3.61	0	0.00
10	3.56	0	0.00

Source Address: 192.0.2.22, Source Port: 1024

Destination Address: 203.0.113.1, Destination Port: 2060

Last DF:MLR: 3.57:0, Avg DF:MLR: 3.60:0

Last MRV: 0.00, Avg MRV: -0.04

Interface Name: ge-0/2/2.0, Template Name: t1

Flow Direction: Output, Flow Type: RTP over IPv4-over-MPLS

MPLS Labels: (299776,16,0)

MDI Records Count: 10

Flow Identifier: 16777217

Rec No	DF	MLR	MRV
1	3.59	0	0.00
2	3.62	0	0.00
3	3.57	0	0.00
4	3.60	0	0.00
5	3.64	0	0.00
6	3.58	0	0.00
7	3.62	0	0.00
8	3.57	0	-0.35
9	3.62	0	0.00
10	3.57	0	0.00

Sample Output

show services video-monitoring mdi flows fpc-slot 2 detail

```
user@host> show services video-monitoring flows fpc-slot 2 detail count 19
```

+-----+										
Rec No	DF	MLR	MRV	PID-0		PID-1		PID-2		
PID-3		PID-4		PID-5						

+-----+										
					Val	MLR	Val	MLR	Val	MLR
Val	MLR	Val	MLR	Val	MLR					

-----+										

1	3.63	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
2	3.59	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
3	3.64	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
4	3.60	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
5	3.64	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
6	3.61	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
7	3.57	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
8	3.62	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
9	3.58	40977	0.00	0x1f40	40977	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				
10	3.63	0	0.00	0x1f40	0	0x1f41	0	0x12	0
0x1f54	0	0x11	0	0x1020	0				

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

Understanding Inline Video Monitoring on MX Series Routers | 927

show services video-monitoring mdi stats fpc-slot

IN THIS SECTION

- Syntax | 1873
- Description | 1873
- Options | 1873

- [Required Privilege Level | 1873](#)
- [Output Fields | 1873](#)
- [Sample Output | 1875](#)
- [Release Information | 1875](#)

Syntax

```
show services video-monitoring mdi stats fpc-slot fpc-slot
```

Description

Display inline video monitoring statistics.

Options

fpc-slot Number of the fpc slot for which statistics are displayed.

Required Privilege Level

view

Output Fields

[Table 192 on page 1873](#) lists the output fields for the `show services video-monitoring mdi stats fpc-slot` command. Output fields are listed in the approximate order in which they appear.

Table 192: show services video-monitoring mdi stats fpc-slot Output Fields

Field Name	Field Description
FPC Slot	Slot number of the monitored FPC

Table 192: show services video-monitoring mdi stats fpc-slot Output Fields (Continued)

Field Name	Field Description
Active Flows	Number of active flows currently monitored. active flows = inserted flows - deleted flows.
Total Inserted Flows	Number of flows initiated under video monitoring.
Total Deleted Flows	Number of flows deleted due to inactivity timeout.
Total Packets Count	Number of total packets monitored.
Total Bytes Count	Number of total bytes monitored.
DF Alarm Count	Number of delay factor alarms at each of the following levels: <ul style="list-style-type: none"> • Info level • Warning level • Critical level
MLR Alarm Count	Number of media loss rate (MLR) alarms at each of the following levels: <ul style="list-style-type: none"> • Info level • Warning level • Critical level
MRV alarm count	Number of media rate variation (MRV) alarms at each of the following levels: <ul style="list-style-type: none"> • Info level • Warning level • Critical level

Sample Output

show services video-monitoring mdi stats fpc-slot

```
user@host> show services video-monitoring mdi stats fpc-slot 2

May 13 22:51:38 blurr fpc2 user.crit afdt-trio: [MDI] MRV: -28.90, exceeded critical threshold
for flow (flow_id:16779635 src:4.4.4.2 dst:7.7.7.2 sport:1000 dport:9000) egressing from
interface index(345) with template mpls_ipv4_layer.
May 13 22:51:42 blurr fpc2 user.crit afdt-trio: [MDI] DF: 292.58 ms, exceeded critical
threshold for flow(flow_id:16777352 src:4.4.4.3 dst:7.7.7.2 sport:1000 dport:9000) egressing
from interface index(345) with template mpls_ipv4_layer.

May 13 23:30:18 blurr fpc2 user.warning afdt-trio: [MDI] MRV: -28.60, exceeded warning
threshold for flow (flow_id:16779635 src:4.4.4.2 dst:7.7.7.2 sport:1000 dport:9000) egressing
from interface index(345) with template mpls_ipv4_layer.
May 13 23:30:23 blurr fpc2 user.warning afdt-trio: [MDI] DF: 289.57 ms, exceeded warning
threshold for flow(flow_id:16777352 src:4.4.4.3 dst:7.7.7.2 sport:1000 dport:9000) egressing
from interface index(345) with template mpls_ipv4_layer.

May 13 23:32:58 blurr fpc2 user.info afdt-trio: [MDI] MRV: -29.30, exceeded info threshold for
flow (flow_id:16779635 src:4.4.4.2 dst:7.7.7.2 sport:1000 dport:9000) egressing from interface
index(345) with template mpls_ipv4_layer.
May 13 23:33:03 blurr fpc2 user.info afdt-trio: [MDI] DF: 291.26 ms, exceeded info threshold
for flow(flow_id:16777352 src:4.4.4.3 dst:7.7.7.2 sport:1000 dport:9000) egressing from
interface index(345) with template mpls_ipv4_layer.
```

Release Information

Command introduced in Junos OS Release 14.1.

RELATED DOCUMENTATION

[Understanding Inline Video Monitoring on MX Series Routers](#) | 927

test services monitoring rfc2544

IN THIS SECTION

- [Syntax | 1876](#)
- [Description | 1876](#)
- [Options | 1877](#)
- [Additional Information | 1877](#)
- [Required Privilege Level | 1877](#)
- [Sample Output | 1877](#)
- [Release Information | 1879](#)

Syntax

```
test services monitoring rfc2544
(test test-name | test-id id)
<routing-instance <routing-instance-name>>
(clear-counters | start | stop)
```

Description

Start or stop an RFC 2544-based benchmarking test, or clear the statistical counters for a test. You can start or stop a specific test name, or you can stop a test based on its test identifier. When you trigger an RFC 2544-based benchmarking test, it passes through a series of states. These states are displayed in the Test state field of the `show services monitoring rfc2544` command.

NOTE: The RFC 2544 test is stopped at the initiator automatically after the test successfully completes all of the test steps. You need not explicitly enter the `test services monitoring rfc2544-benchmarking (test test-name | test-id test-id) stop` command. However, at the reflector, you must explicitly enter this command to stop the test after the test is completed at the initiator.

Options

(test <i>test-name</i> test-id <i>id</i>)	Name of the test to be started or stopped or the unique identifier of the test to be stopped.
routing-instance <i>routing-instance-name</i>	(Optional) Name of the routing instance for the test.
(clear-counters start stop)	Clear the statistics associated with a test, or start or stop a test.

Additional Information

Required Privilege Level

view

Sample Output

test services monitoring rfc2544 test *test-name* start

```
user@host> test services monitoring rfc2544 test mytest start
```

If successful, the identifier of the started test is displayed:

```
Test mytestid 1 started
```

If unsuccessful, an error message is displayed:

```
Error starting test mytest: error-description
```

Some examples:

- If you try to start a test that is already running:

```
Error starting test mytest: previous session with id 1 still active
```


- If you try to start an INET reflector test that has an incorrect address configured on the destination-ipv4-address statement:

```
Error starting test mytest: inet address '192.0.2.2' not local for default routing-instance
```

- If you try to start a test for a test name that has not been configured on the test-name statement:

```
Error starting test mytest: configuration entry for test 'mytest' not found
```

test services monitoring rfc2544 test *test-name* routing-instance *instance-name* start

```
user@host> test services monitoring rfc2544 test mytest routing-instance lg01 start
```

If successful, the identifier of the started test is displayed:

```
Test mytest id 1 started
```

If unsuccessful, an error message is displayed:

```
Error starting test mytest: error-description
```

Some examples:

- If you try to start an INET reflector test in a routing instance that has an address configured on the destination-ipv4-address statement that is not in the given routing-instance:

```
Error starting test mytest: inet address '192.0.2.2' not local in routing-instance lg01
```

- If you try to start a test for a routing instance that has not been configured on the routing-instance statement:

```
Error starting test mytest: routing instance 'lg01' not found
```

test services monitoring rfc2544 test-id *test-id* stop

```
user@host> test services monitoring rfc2544 test-id 1 stop
```

If successful, the identifier of the started test is displayed:

```
Test test-name id 1 stopped
```

If unsuccessful, an error message is displayed:

```
Error stopping test with id 1: error-description
```

For example, if the test identifier you specified when you issued the command doesn't exist:

```
Error stopping test test-name: no active sessions found
```

test services monitoring rfc2544 test *test-name* clear-counters

```
user@host> test services monitoring rfc2544 test mytest clear-counters
```

If successful, a message is displayed describing which counters were cleared.

```
Test mytest id test-id clearing counters
```

If unsuccessful, an error message is displayed:

```
Error clearing counters for test mytest: error-description
```

For example, if you try to clear counters for a test that does not have active sessions:

```
Error clearing counters for test mytest: no active sessions found
```

Release Information

Command introduced for Junos OS Evolved Release 21.1R1.

test services rpm rfc2544-benchmarking test

IN THIS SECTION

- [Syntax \(ACX Series\) | 1880](#)
- [Syntax \(MX104 Router\) | 1880](#)
- [Syntax \(SRX300 and SRX550HM\) | 1881](#)
- [Description | 1881](#)
- [Options | 1881](#)
- [Additional Information | 1882](#)
- [Required Privilege Level | 1882](#)
- [Output Fields | 1882](#)
- [Sample Output | 1882](#)
- [Release Information | 1882](#)

Syntax (ACX Series)

```
test services rpm rfc2544-benchmarking test
<clear-counters>
<routing-instance routing-instance-name>
<test-name>
<test-id>
<start | stop>
```

Syntax (MX104 Router)

```
test services rpm rfc2544-benchmarking test
<test-name>
<test-id>
<start | stop>
```

Syntax (SRX300 and SRX550HM)

```
test services rpm rfc2544-benchmarking test
<clear-counters>
<test-name>
<test-id>
<start | stop>
```

Description

Start or stop an RFC 2544-based benchmarking test. You can start or stop all of the test names that are defined on a device, or start or stop a specific test name. You can also stop a test based on its test identifier. You can also clear the statistical counters associated with the test. When you trigger an RFC 2544-based benchmarking test, it passes through a series of states. These states are displayed in the Test state field in the brief or displayed output information of the `show services rpm rfc2544-benchmarking` command.

NOTE: The RFC 2544 test is stopped at the initiator automatically after the test successfully completes all of the test steps. You need not explicitly enter the `test services rpm rfc2544-benchmarking test <test-name / test-id> stop` command. However, at the reflector, you must explicitly enter this command to stop the test after the test is completed at the initiator.

When a Layer 2 circuit pseudowire is not up, you cannot start the RFC 2544-based benchmarking test in reflection test mode by entering the `test services rpm rfc2544-benchmarking test test-name start` command. If you attempt to start the reflection test mode, a message is displayed explaining the reason for the failure to commence the test.

Options

start	Start the RFC 2544-based benchmarking test
stop	Terminate the RFC 2544-based benchmarking test
clear-counters	(ACX Series routers, SRX300, and SRX550HM devices only) Clear the statistics associated with the benchmarking test that was run.
routing-instance	(ACX Series routers only) (Optional) Name of the routing instance for the test.
test-name	(Optional) Name of the benchmarking test that must be started or stopped.

test-id (Optional) Unique identifier of the test that must be stopped. You can stop a test based on the test identifier. You can use the *test-id* option with only the test `services rpm rfc2544-benchmarking stop` command.

Additional Information

The test session is supported in out-of-service mode for the underlying service. You must not transmit any traffic to the UNI port, configured as a generator or a reflector, that is being tested during the duration of the test.

Required Privilege Level

view

Output Fields

To display the results of the benchmarking test, use the `show services rpm rfc2544-benchmarking test` command.

Sample Output

test services rpm rfc2544-benchmarking test start

```
user@host> test services rpm rfc2544-benchmarking test test1 start
Test "test1" id 56 started
```

The response specifies that a test has been started with test id 56. The test ID can be further used in `show` commands to view test output.

Release Information

Command introduced in Junos OS Release 12.3X52.

RELATED DOCUMENTATION

[Configuring an RFC 2544-Based Benchmarking Test | 739](#)

[Understanding RFC 2544-Based Benchmarking Tests on MX Series Routers and SRX Devices | 728](#)

[rfc2544-benchmarking | 1347](#)