



Junos[®] OS

Flow Monitoring Feature Guide



Modified: 2018-07-10

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. and/or its affiliates in the United States and other countries. All other trademarks may be 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 Flow Monitoring Feature Guide

Copyright © 2018 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://www.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xxi
	Documentation and Release Notes	xxi
	Using the Examples in This Manual	xxi
	Merging a Full Example	xxii
	Merging a Snippet	xxii
	Documentation Conventions	xxiii
	Documentation Feedback	xxv
	Requesting Technical Support	xxv
	Self-Help Online Tools and Resources	xxv
	Opening a Case with JTAC	xxvi
Part 1	Overview	
Chapter 1	Understanding Flow Monitoring	3
	Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch	3
	Flow Monitoring Terms and Acronyms	7
	Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches	8
	Configuring Flow-Monitoring Interfaces	9
	Configuring Flow-Monitoring Properties	10
	Directing Traffic to Flow-Monitoring Interfaces	11
	Exporting Flows	11
	Configuring Time Periods When Flow Monitoring Is Active and Inactive	12
	Example: Configuring Flow Monitoring	12
Chapter 2	Understanding Flow Monitoring Output Formats	15
	Flow Monitoring Output Formats	15
	Flow Monitoring Version 5 Format Output Fields	16
	Flow Monitoring Version 8 Format Output Fields	19
	Flow Monitoring Version 9 Format Output Fields	26
Part 2	Passive Flow Monitoring	
Chapter 3	Understanding Passive Flow Monitoring	37
	Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers	37

Chapter 4	System Requirements for Passive Flow Monitoring	39
	Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers	39
	Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers	40
Chapter 5	Configuring Passive Flow Monitoring	43
	Configuring Passive Flow Monitoring on T Series and M Series Routers	44
	Example: Passive Flow Monitoring Configuration on M, MX and T Series Routers	45
	Verifying Your Work	52
	Using a Dynamic Flow Capture Interface on M, MX and T Series Routers to Monitor Traffic On Demand	59
	Configuring a Dynamic Flow Capture Group on M, MX and T Series Routers	60
	Configuring the Content Destination for Dynamic Flow Capture on M, MX and T Series Routers	61
	Configuring the Control Source for Dynamic Flow Capture on M, MX and T Series Routers	61
	Configuring a Dynamic Flow Capture Interface on an M, MX or T Series Router	62
	Configuring System Logging for Dynamic Flow Capture on an M, MX or T Series Router	63
	Configuring M, MX or T Series Router Thresholds for Recording Dynamic Flow Capture Interface System Log Messages	63
	Monitoring a Capture Group Using SNMP or Show Services Commands	64
	Example: Dynamic Flow Capture Configuration on a Router	64
	Verifying Your Work	66
	Router 1	66
	Copying and Redirecting Traffic on M, MX or T Series Routers with Port Mirroring and Filter-Based Forwarding	67
	Specifying Port Mirroring Input and Output on M, MX or T Series Routers	67
	Creating a Firewall Filter on an M, MX or T Series Router to Split the Port-Mirrored Traffic into Different Instances	69
	Configuring a Routing Table Group on an M, MX or T Series Router to Add Interface Routes into the Forwarding Instance	71
	Using IPSec and an ES PIC on an M, MX or T Series Router to Send Encrypted Traffic to a Packet Analyzer	71
	Applying a Firewall Filter Output Interface on an M, MX or T Series Router to Port-mirror Traffic to PICs or Flow Collection Services	73
	Monitoring Traffic on a Router with a VRF Instance and a Monitoring Group	73
	Specifying a Firewall Filter on an M, MX or T Series Router to Select Traffic to Monitor	73
	Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers	74
	Establishing a VRF Instance on an M, MX or T Series Router for Monitored Traffic	77
	Configuring a Monitoring Group on an M, MX or T Series Router to Send Traffic to the Flow Server	78
	Configuring Policy Options on M, MX or T Series Routers	79

	Stripping MPLS Labels on ATM, Ethernet-Based and SONET/SDH Router Interfaces	80
	Using an M, MX or T Series Router Flow Collector Interface to Process and Export Multiple Flow Records	81
	Example: Configuring a Flow Collector Interface on an M, MX or T Series Router	87
	Verifying Your Work	92
Part 3	Active Flow Monitoring	
Chapter 6	Understanding Active Flow Monitoring	101
	Understanding Active Flow Monitoring Versions on M40e, M320 and T Series Routers	101
	When to Use Active Flow Monitoring Applications on M30, M40e, MX Series and T Series Routers	102
Chapter 7	System Requirements for Active Flow Monitoring	105
	Active Flow Monitoring System Requirements for M and T Series Routers	105
	Active Flow Monitoring PIC Specifications for M and T Series Routers	106
Chapter 8	Configuring Active Flow Monitoring	111
	Understanding Active Flow Monitoring PICS and Options on M, MX and T Series Routers	112
	Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250	115
	Configuring Active Flow Monitoring on PTX Series Packet Transport Routers	121
	Sending Packets to a Mediation Device on MX, M and T Series Routers	123
	Understanding Flow-Tap Architecture	124
	Configuring a Flow-Tap Interface on MX, M and T Series Routers	125
	Configuring Flow-Tap Security Properties on MX, M and T Series Routers	126
	Flow-Tap Application Restrictions	127
	Example: Flow-Tap Configuration on T and M Series Routers	127
	Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs	129
	Example: Configuring IPv6 Support for FlowTapLite on an M120 Router With Enhanced III FPCs	131
	Configuring Actively Monitored Interfaces on M, MX and T Series Routers	138
	Collecting Flow Records on M, MX and T Series Routers	139
	Configuring M, MX and T Series Routers for Discard Accounting with an Accounting Group	139
	Configuring M, MX and T Series Routers for Discard Accounting with a Sampling Group	140
	Configuring M, MX and T Series Routers for Discard Accounting with a Template	141
	Defining a Firewall Filter on M, MX and T Series Routers to Select Traffic for Active Flow Monitoring	143
	Processing IPv4 traffic on an M, MX or T Series Router Using Monitoring services, Adaptive services or Multiservices Interfaces	143
	Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers	144

Replicating Version 9 Flow Aggregation From M, MX and T Series Routers to Multiple Flow Servers	145
Configuring Routing Engine-Based Sampling on M, MX and T Series Routers for Export to Multiple Flow Servers	146
Example: Copying Traffic to a PIC While an M, MX or T Series Router Forwards the Packet to the Original Destination	147
Configuring an Aggregate Export Timer on M, MX and T Series Routers for Version 8 Records	168
Rerouting Packets on an M, MX or T Series Router with Port Mirroring	168
Load-balancing Traffic Across PICs for Active Flow Monitoring on M, MX and T Series Routers	169
Sending Port-Mirrored Traffic from an M, MX or T Series Router to Multiple Export Interfaces by Using Next-Hop Groups	169
Example: Configuring Multiple Port Mirroring with Next-Hop Groups on M, MX and T Series Routers	170
Example: Sampling Configuration for M, MX and T Series Routers	174
Verifying Your Work	176
Example: Sampling Instance Configuration on an MX480 Router	178
Example Network Details	178
Example Router Configuration	180
Configuration Commands Used for the Configuration Example	182
Verifying Your Work	183
Example: Sampling and Discard Accounting Configuration on M, MX and T Series Routers	184
Verifying Your Work	188

Part 4

Chapter 9

Configuration Statements and Operational Commands

Configuration Statements for Routers	193
accounting	201
address (Interfaces)	202
address (Services Dynamic Flow Capture)	202
aggregate-export-interval	203
aggregation	204
alarms	205
alarm-mode	206
allowed-destinations	207
analyzer-address	207
analyzer-id	208
archive-sites	208
authentication-mode	209
authentication-key-chain (TWAMP)	210
autonomous-system-type	211
bandwidth-kbps (RFC 2544 Benchmarking)	212
bgp	213
capture-group	214
cflowd (Discard Accounting)	215
cflowd (Flow Monitoring)	216
client	217

client-delegate-probes	218
client-list	219
collector	219
collector (Flow Monitoring Logs for NAT)	220
collector (Flow Template Profiles for NAT)	221
collector-group (Flow Template Profiles for NAT)	222
collector-group (Flow Monitoring Logs for NAT)	223
content-destination	224
control-connection	225
control-source	226
core-dump	227
data-fill	228
data-fill-with zeros	229
data-format	229
data-size	230
delay-factor	231
delegate-probes	232
destination (Interfaces)	233
destination-address (Flow Monitoring Logs for NAT)	234
destination-interface	235
destination-ipv4-address (RFC 2544 Benchmarking)	236
destination-mac-address (RFC2544 Benchmarking)	237
destination-port	238
destination-port (Flow Monitoring Logs for NAT)	239
destination-udp-port (RFC 2544 Benchmarking)	240
destinations	240
direction (RFC2544 Benchmarking)	241
disable (Forwarding Options)	242
disable-signature-check (RFC 2544 Benchmarking)	243
dscp (flow-server)	244
dscp-code-point (Services)	245
duplicates-dropped-periodicity	246
dynamic-flow-capture	247
engine-id (Forwarding Options)	248
engine-type	249
export-format	250
extension-service	251
family (Monitoring)	253
family (Port Mirroring)	254
family (RFC2544 Benchmarking)	255
family (Sampling)	256
file (Sampling)	258
file (Trace Options)	258
file-specification (File Format)	259
file-specification (Interface Mapping)	259
filename	260
filename-prefix	260
files	261
filter	262

flex-flow-sizing	263
flow-active-timeout	264
flow-collector	266
flow-export-destination	267
flow-export-rate	268
flow-inactive-timeout	269
flow-key (Flow Monitoring)	270
flow-monitoring	271
flow-server	273
flow-table-size	275
flow-table-size (Chassis)	276
flow-tap	277
forwarding-class (Sampling)	278
ftp (Flow Collector Files)	279
ftp (Transfer Log Files)	280
g-duplicates-dropped-periodicity	281
g-max-duplicates	282
generate-snmp-traps	283
hard-limit	283
hard-limit-target	284
hardware-timestamp	284
history-size	285
host-outbound media-interface	286
in-service (RFC2544 Benchmarking)	287
inactivity-timeout (Services RPM)	288
inline-jflow	289
input (Port Mirroring)	290
input (Sampling)	290
input-interface-index	291
input-packet-rate-threshold	291
instance (Sampling)	292
interface (Accounting or Sampling)	293
interfaces	294
interface (Services Flow Tap)	294
interface-map	295
interfaces (Services Dynamic Flow Capture)	295
interfaces (Video Monitoring)	296
inet6-options (Services)	299
ip-swap (RFC 2544 Benchmarking)	300
ipv4-flow-table-size	301
ipv4-template	302
ipv6-flow-table-size	303
ipv6-extended-attrib	304
ipv6-template	304
jflow-log (Interfaces)	305
jflow-log (Services)	306
label-position	307
license-server	308
local-dump	309

logical-system	309
match	310
max-connection-duration	310
max-duplicates	311
max-packets-per-second	312
maximum-age	313
maximum-connections	314
maximum-connections-per-client	315
maximum-packet-length	316
maximum-sessions	318
maximum-sessions-per-connection	319
media-loss-rate	320
media-rate-variation	321
message-rate-limit (Flow Monitoring Logs for NAT)	322
minimum-priority	323
mode (RFC 2544 Benchmarking)	323
monitoring	324
moving-average-size	325
mpls-flow-table-size	326
mpls-ipv4-template	327
mpls-ipv6-template	328
mpls-template	329
multiservice-options	330
name-format	331
next-hop (Forwarding Options)	332
next-hop-group (Forwarding Options)	333
next-hop-group (Port Mirroring)	334
nexthop-learning	335
no-filter-check	336
no-remote-trace (Trace Options)	336
no-syslog	337
no-syslog-generation	337
notification-targets	338
observation-domain-id	339
one-way-hardware-timestamp	340
option-refresh-rate	341
options-template-id	342
output (Accounting)	343
output (Monitoring)	344
output (Port Mirroring)	345
output (Sampling)	346
output-interface-index	347
packet-size (RFC 2544 Benchmarking)	348
passive-monitor-mode	349
password (Flow Collector File Servers)	349
password (Transfer Log File Servers)	350
peer-as-billing-template	350
pic-memory-threshold	351
pop-all-labels	352

port (Flow Monitoring)	353
port (RPM)	353
port (TWAMP)	354
port-mirroring	355
post-cli-implicit-firewall	356
pre-rewrite-tos	357
probe	358
probe-count	360
probe-interval	361
probe-limit	362
probe-server	363
probe-type	364
profiles (RFC 2544 Benchmarking)	365
rate (Forwarding Options)	366
receive-options-packets	367
receive-ttl-exceeded	367
refresh-rate (Flow Monitoring Logs for NAT)	368
reflect-mode (RFC2544 Benchmarking)	369
reflect-etype (RFC 2544 Benchmarking)	370
required-depth	371
retry (Services Flow Collector)	372
retry-delay	372
rfc2544-benchmarking	373
rfc6514-compliant-safil29 (Protocols BGP)	374
routing-instance	375
routing-instance (cflowd)	376
routing-instance-list (TWAMP)	377
routing-instances	378
rpm (Interfaces)	379
rpm (Services)	380
rpm-scale	383
run-length	385
sample-once	386
sampling (Forwarding Options)	387
sampling (Interfaces)	390
server	391
server-inactivity-timeout	392
service-port	392
service-type (RFC2544 Benchmarking)	393
services	394
services	395
services-options	396
shared-key	397
size	398
slamon-services	399
soft-limit	400
soft-limit-clear	400
source-address (Forwarding Options)	401
source-address (Services)	402

source-addresses	403
source-id	403
source-ip (Flow Monitoring Logs for NAT)	404
source-ipv4-address (RFC 2544 Benchmarking)	405
source-mac-address (RFC2544 Benchmarking)	406
source-udp-port (RFC 2544 Benchmarking)	407
stamp	407
storm-control	408
syslog	409
target (Services RPM)	410
tcp	411
template (Flow Monitoring IPFIX Version)	412
template (Flow Monitoring Version 9)	413
template (Forwarding Options)	414
template (Forwarding Options Version IPFIX)	414
template-id	415
template-profile (Flow Monitoring Logs for NAT)	416
template-refresh-rate	417
template-type (Flow Monitoring Logs for NAT)	418
templates	419
test	421
tests (RFC 2544 Benchmarking)	423
test-interface (RFC 2544 Benchmarking)	424
test-interval	425
test-name (RFC 2544 Benchmarking)	426
test-profile (RFC 2544 Benchmarking)	427
test-session	428
test-type (RFC 2544 Benchmarking)	429
thresholds	430
traceoptions (Dynamic Flow Capture)	431
traceoptions (Forwarding Options)	432
traceoptions (RPM)	433
transfer	434
transfer-log-archive	435
traps	436
ttl	437
twamp	438
twamp-server	439
trio-flow-offload	440
tunnel-observation	441
udp	442
udp-tcp-port-swap (RFC 2544 Benchmarking)	443
unit	444
use-extended-flow-memory	445
username (Services)	446
variant	446
version	447
version (Flow Monitoring Logs for NAT)	448
version9 (Forwarding Options)	449

	version9 (Flow Monitoring)	450
	version-ipfix (Forwarding Options)	451
	version-ipfix (Services)	452
	video-monitoring	453
	vpls-flow-table-size	455
	vpls-template	456
	world-readable	456
Chapter 10	Basic Flow and Active Flow EX9200 Monitoring Configuration Statements	457
	address (Interfaces)	458
	cflowd (Discard Accounting)	459
	core-dump	460
	destination (Interfaces)	461
	engine-id (Forwarding Options)	462
	engine-type	463
	export-format	464
	family (Monitoring)	465
	filter	466
	flow-active-timeout	467
	flow-export-destination	468
	flow-inactive-timeout	469
	flow-table-size	470
	input-interface-index	471
	interface (Accounting or Sampling)	471
	ipv4-flow-table-size	472
	ipv6-flow-table-size	473
	monitoring	474
	multiservice-options	475
	output-interface-index	475
	output (Monitoring)	476
	port (Flow Monitoring)	477
	sampling (Interfaces)	478
	source-address (Forwarding Options)	479
	syslog	480
	unit	481
Chapter 11	Basic Flow and Active Flow EX9200 Monitoring Operational Commands	483
	show services accounting aggregation	484
	show services accounting aggregation template	488
	show services accounting errors	490
	show services accounting flow	494
	show services accounting flow-detail	499
	show services accounting memory	504
	show services accounting packet-size-distribution	506
	show services accounting status	508
	show services accounting usage	511

Chapter 12	Active Flow Monitoring Commands	513
	show forwarding-options next-hop-group	514
	show forwarding-options port-mirroring	517
	show services accounting aggregation	519
	show services accounting aggregation template	523
	show services accounting errors	525
	show services accounting flow	529
	show services accounting flow-detail	534
	show services accounting memory	539
	show services accounting packet-size-distribution	541
	show services accounting status	543
	show services accounting usage	546
Chapter 13	Dynamic Flow Capture Commands	549
	clear services dynamic-flow-capture	550
	show services dynamic-flow-capture content-destination	551
	show services dynamic-flow-capture control-source	553
	show services dynamic-flow-capture statistics	556
Chapter 14	Flow Collection Commands	559
	clear services flow-collector statistics	560
	request services flow-collector change-destination primary interface	561
	request services flow-collector change-destination secondary interface	562
	request services flow-collector test-file-transfer	563
	show services flow-collector file interface	564
	show services flow-collector input interface	566
	show services flow-collector interface	568
Chapter 15	Passive Flow Monitoring Commands	575
	clear passive-monitoring statistics	576
	show passive-monitoring error	577
	show passive-monitoring flow	579
	show passive-monitoring memory	581
	show passive-monitoring status	583
	show passive-monitoring usage	585

List of Figures

Part 1	Overview	
Chapter 1	Understanding Flow Monitoring	3
	Figure 1: Active Monitoring Configuration Topology	6
Chapter 2	Understanding Flow Monitoring Output Formats	15
	Figure 2: Version 5 Packet Header Format	16
	Figure 3: Version 5 Flow-Export Flow Header Format	17
	Figure 4: Version 8 Template Flow Format	20
	Figure 5: Version 8 AS Aggregation Flow Entry Format	21
	Figure 6: Version 8 Protocol/Port Aggregation Flow Entry Format	21
	Figure 7: Version 8 Prefix Aggregation Flow Entry Format	23
	Figure 8: Version 8 Source Prefix Aggregation Flow Entry Format	24
	Figure 9: Version 8 Destination Prefix Aggregation Flow Entry Format	25
	Figure 10: Version 9 Flow Header Format	28
	Figure 11: Version 9 Template FlowSet Format	29
	Figure 12: Version 9 Data FlowSet Format	31
	Figure 13: Version 9 Options Template Format	32
	Figure 14: Active Flow Monitoring Version 9 Options Data Record Format	33
Part 2	Passive Flow Monitoring	
Chapter 3	Understanding Passive Flow Monitoring	37
	Figure 15: Passive Flow Monitoring Application Topology	37
Chapter 5	Configuring Passive Flow Monitoring	43
	Figure 16: Passive Flow Monitoring—Topology Diagram	45
	Figure 17: Dynamic Flow Capture Topology	60
	Figure 18: Flow Collector Interface Topology Diagram	87
Part 3	Active Flow Monitoring	
Chapter 6	Understanding Active Flow Monitoring	101
	Figure 19: Active Flow Monitoring	103
Chapter 8	Configuring Active Flow Monitoring	111
	Figure 20: Flow-Tap Topology Diagram	125
	Figure 21: FlowTapLite Topology	133
	Figure 22: Routing Engine-Based Sampling Network Topology	148
	Figure 23: Active Flow Monitoring—Multiple Port Mirroring with Next-Hop Groups Topology Diagram	171
	Figure 24: Active Flow Monitoring—Sampling Configuration Topology Diagram	174

Figure 25: Active Flow Monitoring—Sampling Instance Configuration Topology Diagram	179
Figure 26: Active Flow Monitoring—Sampling and Discard Accounting Topology Diagram	185

List of Tables

	About the Documentation	xxi
	Table 1: Notice Icons	xxiii
	Table 2: Text and Syntax Conventions	xxiii
Part 1	Overview	
Chapter 2	Understanding Flow Monitoring Output Formats	15
	Table 3: Export Version 5 Packet Header Fields	16
	Table 4: Export Version 5 Flow-Export Flow Header Fields	17
	Table 5: Version 8 Flow Template Fields	20
	Table 6: Version 8 AS Aggregation Flow Entry Fields	21
	Table 7: Version 8 Protocol/Port Aggregation Flow Entry Fields	22
	Table 8: Version 8 Prefix Aggregation Flow Entry Fields	23
	Table 9: Version 8 Source Prefix Aggregation Flow Entry Fields	24
	Table 10: Version 8 Destination Prefix Aggregation Flow Entry Fields	25
	Table 11: Flow Monitoring Version 9 Template Formats	26
	Table 12: Version 9 Flow Header Fields	28
	Table 13: Version 9 Template FlowSet Fields	29
	Table 14: Field Type Definitions Supported in Junos OS	30
	Table 15: Version 9 Data FlowSet Format	31
	Table 16: Version 9 Options Template Format	32
	Table 17: Active Flow Monitoring Version 9 Options Data Record Format	33
Part 2	Passive Flow Monitoring	
Chapter 5	Configuring Passive Flow Monitoring	43
	Table 18: Passive Flow Monitoring PIC Support	44
	Table 19: Output Fields for the show passive-monitoring error Command	53
	Table 20: Output Fields for the show passive-monitoring flow Command	54
	Table 21: Output Fields for the show passive-monitoring memory Command	56
	Table 22: Output Fields for the show passive-monitoring status Command	57
	Table 23: Output Fields for the show passive-monitoring usage Command	58
	Table 24: Name Format Macros	82
	Table 25: Flow Collector Interface Transfer Log Fields	95
	Table 26: Flow Collector Interface File Fields in Order of Appearance	96
Part 3	Active Flow Monitoring	
Chapter 7	System Requirements for Active Flow Monitoring	105
	Table 27: Monitoring Services PIC Specifications	106
	Table 28: Monitoring Services II PIC Specifications	107

	Table 29: Adaptive Services PIC Specifications	107
	Table 30: MultiServices 100 PIC	107
	Table 31: MultiServices 400 PIC	108
	Table 32: MultiServices 500 PIC	108
Chapter 8	Configuring Active Flow Monitoring	111
	Table 33: Passive and Active Flow Monitoring PIC Support	112
Part 4	Configuration Statements and Operational Commands	
Chapter 11	Basic Flow and Active Flow EX9200 Monitoring Operational Commands	483
	Table 34: show services accounting aggregation Output Fields	485
	Table 35: show services accounting aggregation template Output Fields	488
	Table 36: show services accounting errors Output Fields	490
	Table 37: show services accounting flow Output Fields	495
	Table 38: show services accounting flow-detail Output Fields	500
	Table 39: show services accounting memory Output Fields	504
	Table 40: show services accounting packet-size-distribution Output Fields	506
	Table 41: show services accounting status Output Fields	508
	Table 42: show services accounting usage Output Fields	511
Chapter 12	Active Flow Monitoring Commands	513
	Table 43: show forwarding-options next-hop-group Output Fields	514
	Table 44: show forwarding-options port-mirroring Output Fields	517
	Table 45: show services accounting aggregation Output Fields	520
	Table 46: show services accounting aggregation template Output Fields	523
	Table 47: show services accounting errors Output Fields	525
	Table 48: show services accounting flow Output Fields	530
	Table 49: show services accounting flow-detail Output Fields	535
	Table 50: show services accounting memory Output Fields	539
	Table 51: show services accounting packet-size-distribution Output Fields	541
	Table 52: show services accounting status Output Fields	543
	Table 53: show services accounting usage Output Fields	546
Chapter 13	Dynamic Flow Capture Commands	549
	Table 54: show services dynamic-flow-capture content-destination Output Fields	551
	Table 55: show services dynamic-flow-capture control-source Output Fields	553
	Table 56: show services dynamic-flow-capture statistics Output Fields	556
Chapter 14	Flow Collection Commands	559
	Table 57: show services flow-collector file interface Output Fields	564
	Table 58: show services flow-collector input interface Output Fields	566
	Table 59: show services flow-collector interface Output Fields	568
Chapter 15	Passive Flow Monitoring Commands	575
	Table 60: show passive-monitoring error Output Fields	577
	Table 61: show passive-monitoring flow Output Fields	579
	Table 62: show passive-monitoring memory Output Fields	581
	Table 63: show passive-monitoring status Output Fields	583

Table 64: show passive-monitoring usage Output Fields	585
---	-----

About the Documentation

- Documentation and Release Notes on page xxi
- Using the Examples in This Manual on page xxi
- Documentation Conventions on page xxiii
- Documentation Feedback on page xxv
- Requesting Technical Support on page xxv

Documentation and Release Notes

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

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

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

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

```
[edit]
```

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

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







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

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

Documentation Conventions

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

Table 1: Notice Icons

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

[Table 2 on page xxiii](#) defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	<code>user@host> show chassis alarms</code> <code>No alarms currently active</code>
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: <code>[edit]</code> <code>root@# set system domain-name domain-name</code>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the <code>[edit protocols ospf area area-id]</code> hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	<code>stub <default-metric metric>;</code>
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	<code>broadcast multicast</code> <code>(string1 string2 string3)</code>
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	<code>rsvp { # Required for dynamic MPLS only</code>
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	<code>community name members [community-ids]</code>
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	<code>[edit]</code> <code>routing-options {</code> <code> static {</code> <code> route default {</code> <code> nexthop address;</code> <code> retain;</code> <code> }</code> <code> }</code> <code>}</code>
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none"> In the Logical Interfaces box, select All Interfaces. To cancel the configuration, click Cancel.

Table 2: Text and Syntax Conventions (continued)

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

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

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

Requesting Technical Support

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

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://prsearch.juniper.net/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>

- Download the latest versions of software and review release notes:
<https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:
<https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Understanding Flow Monitoring on page 3](#)
- [Understanding Flow Monitoring Output Formats on page 15](#)

CHAPTER 1

Understanding Flow Monitoring

- [Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch on page 3](#)
- [Flow Monitoring Terms and Acronyms on page 7](#)
- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch

A flow monitoring application performs traffic flow monitoring and enables lawful interception of packets transiting between two devices. Traffic flows can either be passively monitored by an offline device or actively monitored by a device participating in the network.

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. That is, intercept unwanted traffic, discard it, and perform accounting on the discarded packets.
- 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* 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:

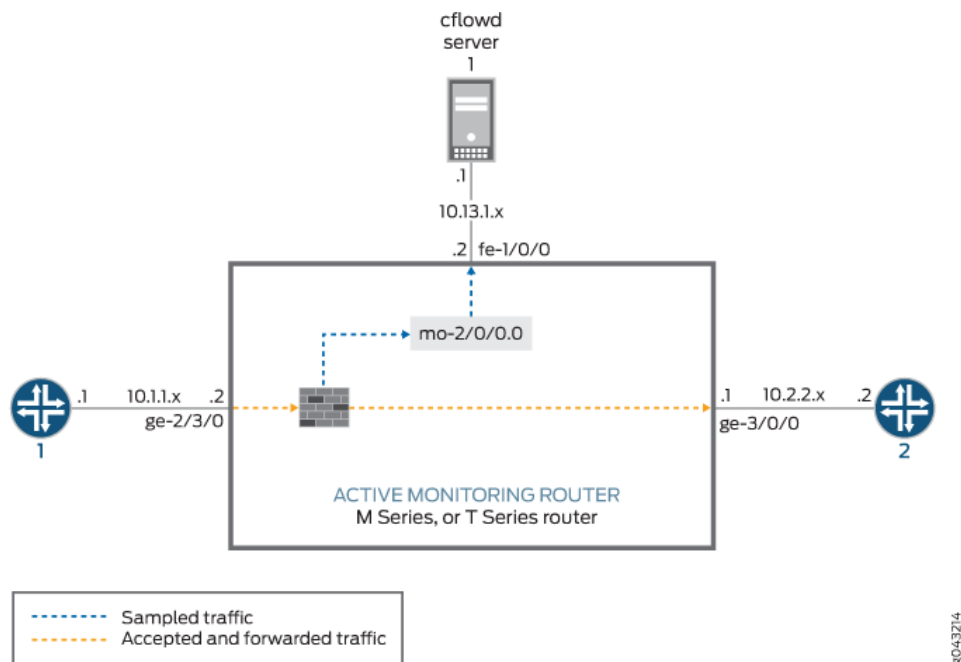
- The router or switch can perform sampling or port mirroring at any one time.
- The router or switch can perform forwarding or discard accounting at any one time.

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 1 on page 6](#) shows a sample topology.

Figure 1: Active Monitoring Configuration Topology



In Figure 1 on page 6, 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.

To learn more about passive flow monitoring, active flow monitoring, cflowd versions 5 and 8, and flow monitoring version 9 see the following:

- Version 9: RFC 3954 at <http://www.faqs.org/rfcs/rfc3954.html>
- Versions 5 and 8: Cooperative Association for Internet Data Analysis (CAIDA) website at <http://www.caida.org>
- *Junos Services Interfaces Configuration Guide*
- *Junos Policy Framework Configuration Guide*
- Internet draft draft-cavuto-dtcp-01.txt, *DTCP: Dynamic Tasking Control Protocol* (expires March 2007)

For more information on IPSec and the ES PIC, see the *Junos System Basics Configuration Guide*.

- | | |
|------------------------------|---|
| Related Documentation | <ul style="list-style-type: none"> • Understanding Active Flow Monitoring Versions on M40e, M320 and T Series Routers on page 101 • Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers on page 37 • Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch on page 3 • <i>Understanding Passive Flow Monitoring Using either a Juniper Networks M Series Multiservice Edge, T Series Core Router, or PICs</i> • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8 • <i>Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers</i> • <i>Configuring Services Interface Redundancy on M and T Series Routers using Flow Monitoring</i> • <i>Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System</i> |
|------------------------------|---|

Flow Monitoring Terms and Acronyms

A

- | | |
|-------------------------------|--|
| 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 <i>Junos Services Interfaces Configuration Guide</i> . |

C

- | | |
|----------------------------|---|
| 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. |

D

DTCP (Dynamic Tasking Control Protocol)	Protocol used to specify filtering criteria in a dynamic flow capture environment.
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.

E

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

F

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.
---------------------------------	--

M

Monitoring Services II PIC	Advanced PIC that handles passive flow monitoring functions.
Monitoring Services III PIC	Advanced PIC that handles dynamic flow capture functions.
Monitoring Services PIC	Original PIC that handles passive and active flow monitoring 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.

P

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 on T Series and M Series Routers and EX9200 Switches

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.

To configure flow monitoring you need to do the following:

- [Configuring Flow-Monitoring Interfaces on page 9](#)
- [Configuring Flow-Monitoring Properties on page 10](#)
- [Example: Configuring Flow Monitoring on page 12](#)

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;
}
```

For more information, see the *Junos OS Administration Library*.

- 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 on page 11](#)
- [Exporting Flows on page 11](#)
- [Configuring Time Periods When Flow Monitoring Is Active and Inactive on page 12](#)

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 [“Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch” on page 3](#).

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 T and M Series Routers*.

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 T and M Series Routers*.

```
[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

- [Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch on page 3](#)

- *Directing Replicated Flows from M and T Series Routers to Multiple Flow Servers*
- *Configuring Services Interface Redundancy on M and T Series Routers using Flow Monitoring*
- *Example: Configuring Active Monitoring on an M, MX or T Series Router's Logical System*

CHAPTER 2

Understanding Flow Monitoring Output Formats

- [Flow Monitoring Output Formats on page 15](#)
- [Flow Monitoring Version 5 Format Output Fields on page 16](#)
- [Flow Monitoring Version 8 Format Output Fields on page 19](#)
- [Flow Monitoring Version 9 Format Output Fields on page 26](#)

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 16](#)

- [Flow Monitoring Version 8 Format Output Fields on page 19](#)
- [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 2 on page 16](#)
- [Table 3 on page 16](#)
- [Figure 3 on page 17](#)
- [Table 4 on page 17](#)

Figure 2: Version 5 Packet Header Format

Byte 3	Byte 2	Byte 1	Byte 0
Version		Count	
sysUptime			
UNIX seconds			
UNIX nanoseconds			
Flow sequence number			
Engine type	Engine ID	Reserved	

g003132

Table 3: 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
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

Table 3: Export Version 5 Packet Header Fields (continued)

Field	Description	Comments
Engine ID	User-configured 8-bit value	—

Figure 3: 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	

9003133

Table 4: 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>

Table 4: Export Version 5 Flow-Export Flow Header Fields (continued)

Field	Description	Comments
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	–
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 4 on page 20](#)
- [Table 5 on page 20](#)
- [Figure 5 on page 21](#)
- [Table 6 on page 21](#)
- [Figure 6 on page 21](#)
- [Table 7 on page 22](#)
- [Figure 7 on page 23](#)
- [Table 8 on page 23](#)
- [Figure 8 on page 24](#)
- [Table 9 on page 24](#)
- [Figure 9 on page 25](#)
- [Table 10 on page 25](#)

Figure 4: Version 8 Template Flow Format*Table 5: Version 8 Flow Template Fields*

Field	Description
Version	8
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
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 5: 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 6: 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
Output interface	SNMP index value for the output interface where the router forwards flows

Figure 6: Version 8 Protocol/Port Aggregation Flow Entry Format

Byte 3	Byte 2	Byte 1	Byte 0
Flows			
Packets			
Bytes			
Start Time of Flow			
End Time of Flow			
IP Protocol	Padding	Reserved	
Source port		Destination port	

9003078

Table 7: 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
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 7: 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	

9003079

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	

9003079

Table 8: Version 8 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
Destination prefix	Destination IP address prefix
Source mask length	Source address network mask length

Table 8: Version 8 Prefix Aggregation Flow Entry Fields (continued)

Field	Description
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 8: 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 9: 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

Table 9: Version 8 Source Prefix Aggregation Flow Entry Fields (continued)

Field	Description
Source AS	AS number of the source address
Input interface	SNMP index value for the input interface where the router receives flows
Reserved	Empty field reserved for future usage

Figure 9: Version 8 Destination Prefix Aggregation Flow Entry Format

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

1803081

Table 10: 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
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

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

- [Table 11 on page 26](#)
- [Figure 10 on page 28](#)
- [Table 12 on page 28](#)
- [Figure 12 on page 31](#)
- [Table 12 on page 28](#)
- [Figure 13 on page 32](#)
- [Table 16 on page 32](#)
- [Figure 14 on page 33](#)
- [Table 17 on page 33](#)

The Junos OS supports the version 9 template formats:

Table 11: Flow Monitoring Version 9 Template Formats

Template	Fields
IPv4	<p>Flow selectors:</p> <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 <p>Flow nonselectors:</p> <ul style="list-style-type: none">• TCP flags• Input and output SNMP• Input bytes• Input packets• Start time• End time

Table 11: 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
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 11: Flow Monitoring Version 9 Template Formats (continued)

Template	Fields
Peer AS billing	<p>Flow selectors:</p> <ul style="list-style-type: none"> IPv4 class of service Ingress interface information BGP peer destination AS number BGP IPv4 next hop address <p>Flow nonselectors</p> <ul style="list-style-type: none"> Input and output SNMP Input bytes Input packets First switch Last switched <p>NOTE: Peer AS billing traffic is not supported for active flow monitoring version 9 configuration on PTX5000 routers tethered to CSE2000.</p>

Figure 10: Version 9 Flow Header Format

Byte 3	Byte 2	Byte 1	Byte 0
Version		Count	
sysUptime			
UNIX seconds			
Flow Sequence Number			
Source ID			

9016765

Table 12: 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.

Table 12: Version 9 Flow Header Fields (continued)

Field	Description
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 11: 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	

98/1786

Table 13: 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 14 on page 30 .
Field Length	Length, in bytes, of the corresponding field type.

Table 14: 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.
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.

Table 14: Field Type Definitions Supported in Junos OS (continued)

Field Type	Description
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.
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 12: Version 9 Data FlowSet Format

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

28/10/2017

Table 15: 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.
Length	FlowSet length. Data FlowSets are fixed in length.

Table 15: Version 9 Data FlowSet Format (continued)

Field	Description
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 13: 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	

9016780

Table 16: 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.
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 14: 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	

9016769

Table 17: 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.

PART 2

Passive Flow Monitoring

- [Understanding Passive Flow Monitoring on page 37](#)
- [System Requirements for Passive Flow Monitoring on page 39](#)
- [Configuring Passive Flow Monitoring on page 43](#)

CHAPTER 3

Understanding Passive Flow Monitoring

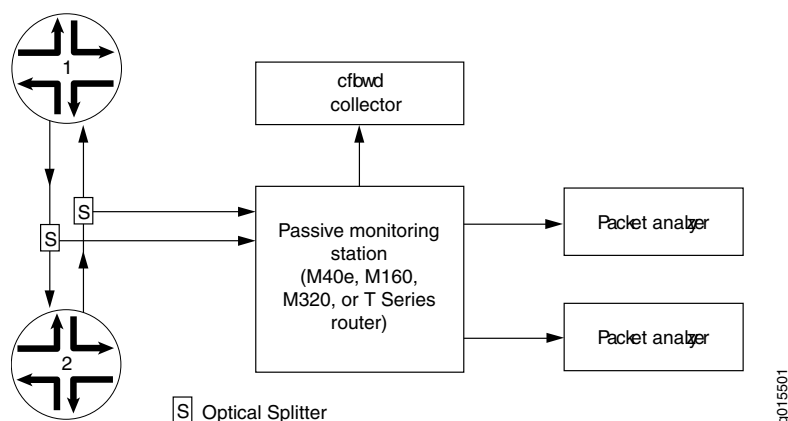
- Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers on page 37

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 15 on page 37](#) shows a typical topology for the passive flow monitoring application.

Figure 15: 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**
- [Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch on page 3](#)
 - [Understanding Active Flow Monitoring Versions on M40e, M320 and T Series Routers on page 101](#)

CHAPTER 4

System Requirements for Passive Flow Monitoring

- [Passive Flow Monitoring System Requirements for T Series, M Series and MX Series Routers on page 39](#)
- [Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers on page 40](#)

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 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 for M and T Series Routers on page 105](#)
- [Active Flow Monitoring PIC Specifications for M and T Series Routers on page 106](#)

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 [Figure 15 on page 37](#), 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.

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.

CHAPTER 5

Configuring Passive Flow Monitoring

- [Configuring Passive Flow Monitoring on T Series and M Series Routers on page 44](#)
- [Example: Passive Flow Monitoring Configuration on M, MX and T Series Routers on page 45](#)
- [Using a Dynamic Flow Capture Interface on M, MX and T Series Routers to Monitor Traffic On Demand on page 59](#)
- [Configuring a Dynamic Flow Capture Group on M, MX and T Series Routers on page 60](#)
- [Configuring the Content Destination for Dynamic Flow Capture on M, MX and T Series Routers on page 61](#)
- [Configuring the Control Source for Dynamic Flow Capture on M, MX and T Series Routers on page 61](#)
- [Configuring a Dynamic Flow Capture Interface on an M, MX or T Series Router on page 62](#)
- [Configuring System Logging for Dynamic Flow Capture on an M, MX or T Series Router on page 63](#)
- [Configuring M, MX or T Series Router Thresholds for Recording Dynamic Flow Capture Interface System Log Messages on page 63](#)
- [Monitoring a Capture Group Using SNMP or Show Services Commands on page 64](#)
- [Example: Dynamic Flow Capture Configuration on a Router on page 64](#)
- [Copying and Redirecting Traffic on M, MX or T Series Routers with Port Mirroring and Filter-Based Forwarding on page 67](#)
- [Specifying Port Mirroring Input and Output on M, MX or T Series Routers on page 67](#)
- [Creating a Firewall Filter on an M, MX or T Series Router to Split the Port-Mirrored Traffic into Different Instances on page 69](#)
- [Configuring a Routing Table Group on an M, MX or T Series Router to Add Interface Routes into the Forwarding Instance on page 71](#)
- [Using IPsec and an ES PIC on an M, MX or T Series Router to Send Encrypted Traffic to a Packet Analyzer on page 71](#)
- [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 page 73](#)
- [Monitoring Traffic on a Router with a VRF Instance and a Monitoring Group on page 73](#)
- [Specifying a Firewall Filter on an M, MX or T Series Router to Select Traffic to Monitor on page 73](#)

- [Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers on page 74](#)
- [Establishing a VRF Instance on an M, MX or T Series Router for Monitored Traffic on page 77](#)
- [Configuring a Monitoring Group on an M, MX or T Series Router to Send Traffic to the Flow Server on page 78](#)
- [Configuring Policy Options on M, MX or T Series Routers on page 79](#)
- [Stripping MPLS Labels on ATM, Ethernet-Based and SONET/SDH Router Interfaces on page 80](#)
- [Using an M, MX or T Series Router Flow Collector Interface to Process and Export Multiple Flow Records on page 81](#)
- [Example: Configuring a Flow Collector Interface on an M, MX or T Series Router on page 87](#)

Configuring Passive Flow Monitoring on T Series and M Series Routers

Table 18 on page 44 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 18: 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

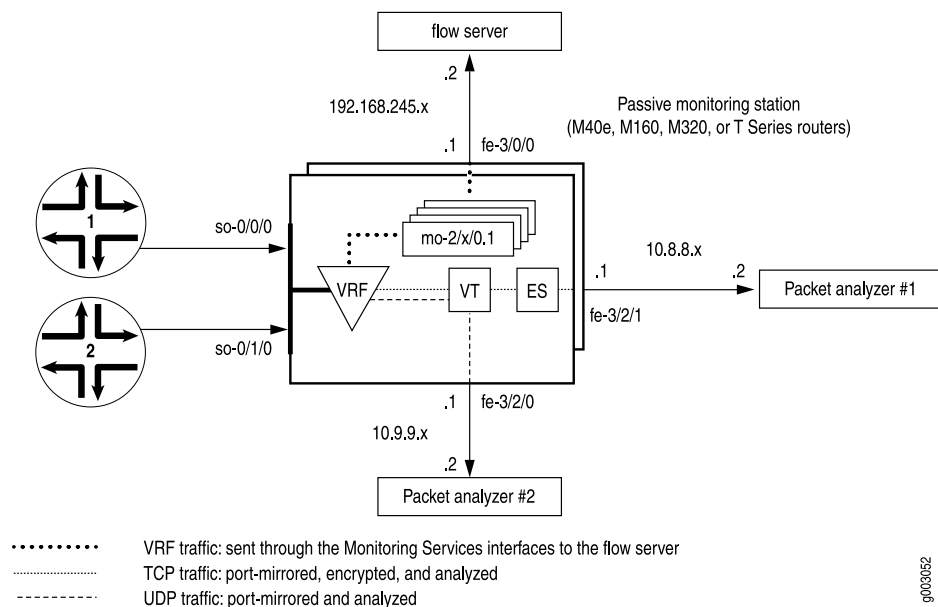
The key configuration hierarchy statement for passive flow monitoring is the **monitoring** statement found at the **[edit forwarding-options]** hierarchy level. At minimum, you must configure a VRF routing instance to direct the traffic to a monitoring services interface for flow processing.

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 on M, MX or T Series Routers with Port Mirroring and Filter-Based Forwarding on page 67](#)
 - [Using an M, MX or T Series Router Flow Collector Interface to Process and Export Multiple Flow Records on page 81](#)
 - [Using a Dynamic Flow Capture Interface on M, MX and T Series Routers to Monitor Traffic On Demand on page 59](#)
 - [Passive Flow Monitoring Router and Software Considerations for T Series, M Series and MX Series Routers on page 40](#)

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

Figure 16: Passive Flow Monitoring—Topology Diagram



In [Figure 16 on page 45](#), 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;
        }
      }
    }
  }
}
```

51

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] 5d17:34:57
                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 19: 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.
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 19: 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 20: 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.

Table 20: Output Fields for the show passive-monitoring flow Command (continued)

Field	Explanation
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 21: 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.
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 22: 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.
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 23: 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%

```

Using a Dynamic Flow Capture Interface on M, MX and T Series Routers to Monitor Traffic On Demand

Dynamic flow capture enables you to capture packet flows based on filtering criteria that you specify in real time. Unlike traditional flow monitoring that requires filtering criteria to be established before operation, dynamic flow capture uses an on demand control protocol that allows you to modify the filtering criteria as network conditions change.

The dynamic flow capture architecture consists of one or more *control sources* that send Dynamic Tasking Control Protocol (DTCP) requests to a *monitoring station*. The requests contain filtering criteria that specify which incoming traffic should be monitored, and the monitoring station forwards any packets that match the filter criteria to a set of one or more *content destinations*.

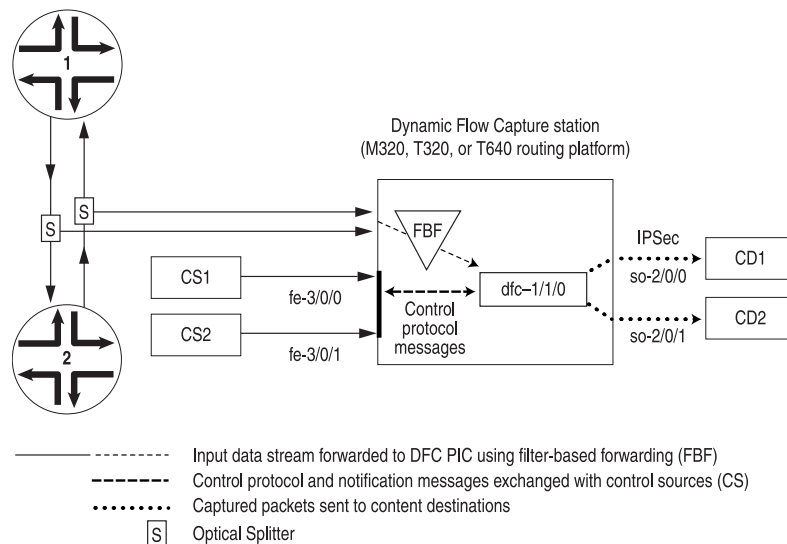
- Control source—A client that wants to monitor electronic data or voice transfer over the network. The control source sends filter requests to the Juniper Networks router using DTCP. The control source is identified by a unique identifier and an optional list of IP addresses.
- Monitoring station—A Juniper Networks T Series or M320 router configured with one or more Monitoring Services III PICs which support dynamic flow capture processing. The monitoring station 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 station. Typically the matched packets are sent using an IPSec tunnel from the monitoring station to another router connected to the content destination. The content destination and the control source can be located on the same host.



NOTE: The DFC PIC forwards the entire packet content to the content destination, rather than just a content record.

Figure 17 on page 60 shows a sample topology that contains control sources, a monitoring station, and content destinations.

Figure 17: Dynamic Flow Capture Topology



g017075

To configure dynamic flow capture, perform the following tasks:

- [Configuring a Dynamic Flow Capture Group on M, MX and T Series Routers on page 60](#)
- [Configuring the Content Destination for Dynamic Flow Capture on M, MX and T Series Routers on page 61](#)
- [Configuring the Control Source for Dynamic Flow Capture on M, MX and T Series Routers on page 61](#)
- [Configuring a Dynamic Flow Capture Interface on an M, MX or T Series Router on page 62](#)
- [Configuring M, MX or T Series Router Thresholds for Recording Dynamic Flow Capture Interface System Log Messages on page 63](#)
- [Configuring System Logging for Dynamic Flow Capture on an M, MX or T Series Router on page 63](#)
- [Monitoring a Capture Group Using SNMP or Show Services Commands on page 64](#)

Configuring a Dynamic Flow Capture Group on M, MX and T Series Routers

A dynamic flow capture group defines a profile of dynamic flow capture 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:

```
[edit services dynamic-flow-capture]
capture-group client-name {
  content-destination identifier {
    address address;
```

```

        ttl hops;
    }
    control-source identifier {
        allowed-destinations [ destination ];
        no-syslog;
        notification-targets [ address address port port-number ];
        service-port port-number;
        shared-key value;
        source-addresses [ address ];
    }
    input-packet-rate-threshold rate;
    interfaces interface-name;
    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 for Dynamic Flow Capture on M, MX and T Series Routers

You must specify a destination for the packets that match dynamic flow capture filter criteria. To configure, include the **content-destination** statement at the [edit services dynamic-flow-capture capture-group *client-name*] hierarchy level:

```

[edit services dynamic-flow-capture capture-group client-name]
content-destination identifier {
    address address;
    ttl hops;
}

```

Assign the **content-destination** a unique **identifier**. In addition, you must specify its IP address, and you can optionally set a time-to-live (TTL) value for the IP-IP header:

- **address**—The dynamic flow capture 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**—By default, the TTL value is 255, with a range from 0 through 255.

Configuring the Control Source for Dynamic Flow Capture on M, MX and T Series Routers

You configure information about the control source, including allowed source addresses, destinations, and authentication key values. To configure the control source information, include the **control-source** statement at the [edit services dynamic-flow-capture] hierarchy level:

```

[edit services dynamic-flow-capture capture-group client-name]
control-source identifier {
    allowed-destinations [ destination-identifier ];
    no-syslog;
    notification-targets [ address address port port-number ];
    service-port port-number;
    shared-key value;
}

```

```
    source-addresses [ address ];  
}
```

Assign the **control-source** statement with 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 matched data to be sent in its control protocol requests. If you do not specify any content destinations, all available destinations are allowed.
- **notification-targets**—One or more destinations to which the dynamic flow capture interface can log information about control protocol-related events and other events such as PIC startup 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 dynamic flow capture interfaces.
- **shared-key**—A 20-byte authentication key value shared between the control source and the dynamic flow capture monitoring station.
- **source-addresses**—One or more allowed IP addresses from which the control source can send control protocol requests to the dynamic flow capture monitoring station. These are /32 addresses.

Configuring a Dynamic Flow Capture Interface on an M, MX or T Series Router

You specify the interface that interacts with the control sources configured in the same dynamic flow 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 dynamic flow capture interface, include the **interfaces** statement at the **[edit services dynamic-flow-capture]** hierarchy level:

```
[edit services dynamic-flow-capture capture-group client-name]  
  interfaces interface-name;
```

You specify dynamic flow capture interfaces using the **dfc-** identifier at the **[edit interfaces]** hierarchy level. Three logical units are required on each dynamic flow capture 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 dynamic flow capture interface:

```
[edit interfaces dfc-0/0/0]  
unit 0 {  
  family inet {  
    address 10.1.0.0/32 { # Address of the Routing Engine for the DFC PIC.
```



```

        destination 10.36.100.1; # Address of DFC PIC; used by the
        # control source to correspond with the monitoring station.
    }
}
unit 1 { # Receives data packets on this logical interface.
    family inet;
}
unit 2 { # Sends copies of matched packets from this logical interface.
    family inet;
}

```

In addition, you must configure the dynamic flow capture application to run on the DFC PIC in the correct chassis location. The following example shows this configuration at the **[edit chassis]** hierarchy level:

```

[edit chassis]
fpc 0 {
    pic 0 {
        monitoring-services application dynamic-flow-capture;
    }
}

```

For more information on configuring chassis properties, see the *Junos System Basics Configuration Guide*.

Configuring System Logging for Dynamic Flow Capture on an M, MX or T Series Router

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:

```

[edit syslog]
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:

```

[edit services dynamic-flow-capture capture-group client-name control-source identifier]
no-syslog;

```

Configuring M, MX or T Series Router Thresholds for Recording Dynamic Flow Capture Interface System Log Messages

You can optionally specify threshold values for situations in which warning messages will be recorded in the system log:

- Input packet rate to the dynamic flow capture interfaces
- Memory usage on the dynamic flow capture interfaces

To configure, include the **input-packet-rate-threshold** or **pic memory-threshold** statements at the **[edit services dynamic-flow-capture capture-group client-name]** hierarchy level:

```
[edit services dynamic-flow-capture capture-group client-name]  
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.

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*.

Example: Dynamic Flow Capture Configuration on a Router

The following example shows a complete dynamic flow capture configuration. On Router 1, configure the dynamic flow capture interface, the interfaces that connect to the control source and content destination, and the interface that receives passively monitored traffic. Then, configure the capture group and specify your control source and content destination requirements. Next, configure filter-based forwarding (FBF) to send monitored traffic to logical unit 1 of the dynamic flow capture interface. Finally, configure a firewall filter and routing table groups to complete the configuration.

```
[edit]  
interfaces {  
  dfc-0/0/0 { # DFC PIC that processes requests from the control source.  
    unit 0 {  
      family inet {  
        address 192.0.2.0/32 { # Address of the Routing Engine for the DFC PIC.  
          destination 10.36.100.1; # Address of DFC PIC; used by  
            } # the control source to communicate with the monitoring station.  
        }  
      }  
    }  
    unit 1 { # This logical interface receives data packets.  
      family inet;  
    }  
    unit 2 { # This logical interface sends out copies of matched packets.  
      family inet;  
    }  
  }  
  fe-4/1/2 { # Interface that receives filtering requests from cs1.  
    unit 0 {  
      family inet {  
        address 10.36.41.2/30;  
      }  
    }  
  }  
  ge-7/0/0 { # Interface that sends monitored packets to cd1.  
    unit 0 {  
      family inet {
```

```

        address 10.36.70.1/30;
    }
}
so-1/2/0 { # Interface that receives traffic to be monitored.
    encapsulation ppp;
    unit 0 {
        passive-monitor-mode; # Enables this interface to be passively monitored.
        family inet {
            filter {
                input catch;
            }
        }
    }
}
}
services {
    dynamic-flow-capture {
        capture-group gl {
            interfaces dfc-0/0/0; # Specifies which interface to use for DFC processing.
            input-packet-rate-threshold 90k; # Traffic threshold for system log messages.
            pic-memory-threshold percentage 80; # Memory threshold for log messages.
            control-source cs1 { # Specifies addresses and ports for the control source.
                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 { # Specifies content destination addresses and TTL.
                address 10.36.70.2;
                ttl 244;
            }
        }
    }
}
firewall {
    filter catch { # Places monitored traffic into the filter-based forwarding instance.
        interface-specific;
        term def {
            then {
                count counter;
                routing-instance fbf_inst;
            }
        }
    }
}
routing-instances {
    fbf_inst { # Sends matching traffic to the DFC PIC for processing.
        instance-type forwarding;
        routing-options {
            static {
                route 0.0.0.0/0 next-hop dfc-0/0/0.1;
            }
        }
    }
}

```

```
    }  
  }  
}  
routing-options {  
  interface-routes {  
    rib-group inet common;  
  }  
  rib-groups {  
    common { # Shares routes between the instance and the main routing table.  
      import-rib [ inet.0 fbf_inst.inet.0 ];  
    }  
  }  
  forwarding-table {  
    export pplb;  
  }  
}
```

Verifying Your Work

To verify that your dynamic flow capture configuration is operating correctly, issue the following command:

```
show services dynamic-flow-capture capture-group group-name control-source  
source-identifier source-id (detail)
```

The following section shows the output of this command when used with the configuration example.

Router 1

```
user@router1> show services dynamic-flow-capture control-source capture-group g1 source-identifier cs2 detail
```

```
Capture group: g1, Control source: cs2  
Criteria added: 1, Criteria add failed: 0  
Active criteria: 2  
Static criteria: 0, Dynamic criteria: 2  
Control protocol requests: 3  
      Add      Delete      List      Refresh      No-op  
Requests      1          0          1          0          1  
Failed        0          0          0          0          0  
  
Add request rate: 0  
Add request peak rate: 1  
Bandwidth across all criteria: 0  
Total notifications: 0  
Restart: 0, Rollover: 0, No-op: 0, Timeout: 0, Congestion: 0, Congestion delete: 0,  
Dups dropped: 0  
Criteria deleted: 0  
Timeout idle: 0, Timeout total: 0, Packets: 0, Bytes: 0  
Sequence number: 242
```

To clear dynamic flow capture criteria belonging to a particular control source, issue the **clear services dynamic-flow-capture** command. For more information on other dynamic flow capture-related operational mode commands, see the *Junos System Basics and Services Command Reference*.

Copying and Redirecting Traffic on M, MX or T Series Routers with Port Mirroring and Filter-Based Forwarding

This section discusses additional techniques you can use with the passive flow monitoring application:

- In addition to flow analysis, you can analyze a copy of the original traffic with a single packet analyzer. To implement this technique, divert traffic with a filter-based forwarding routing instance and send the monitored traffic through a physical interface to the packet analyzer.
- You can cluster the traffic into different groups and redirect this traffic to multiple packet analyzers. For example, you can break traffic flows into TCP groups and UDP groups and send these groups of packets to different analyzers. To accomplish this, you use port mirroring and send a copy of the original traffic to a Tunnel PIC. Then you can apply a firewall filter, split the traffic into your desired groups, and send these groups toward different exit interfaces leading to the packet analyzers. This technique provides maximum flexibility for traffic analysis.
- For secure transmission of the copied or grouped traffic, you can encrypt the diverted traffic with an ES PIC and send this traffic to a packet analyzer over an IP Security (IPSec) tunnel.

To implement the filter-based forwarding enhancement methods, see the following sections:

- [Specifying Port Mirroring Input and Output on M, MX or T Series Routers on page 67](#)
- [Creating a Firewall Filter on an M, MX or T Series Router to Split the Port-Mirrored Traffic into Different Instances on page 69](#)
- [Applying the Firewall Filter to a Tunnel PIC Interface on page 70](#)
- [Using Filter-Based Forwarding to Export Monitored Traffic to Multiple Destinations on page 70](#)
- [Configuring a Routing Table Group on an M, MX or T Series Router to Add Interface Routes into the Forwarding Instance on page 71](#)
- [Using IPSec and an ES PIC on an M, MX or T Series Router to Send Encrypted Traffic to a Packet Analyzer on page 71](#)
- [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 page 73](#)

Specifying Port Mirroring Input and Output on M, MX or T Series Routers

This step works in conjunction with the action specified by the **port-mirror** statement configured at the **[edit firewall family (inet | inet6) filter *filter-name* term *term-name* then]** hierarchy level. At this point, you select input and output statements to determine where the copies of the IPv4 or IPv6 packets are sent. To configure, include the **input** and **output** statements at the **[edit forwarding-options port-mirroring family *family-name*]** hierarchy

level. The traffic to be monitored is copied, port-mirrored, and sent to the packet analyzer for analysis.



NOTE: On M Series routers, you can port-mirror either IPv4 or IPv6 packets at one time. On M120, M320, and T Series routers, you can port-mirror both IPv4 and IPv6 packets simultaneously.



NOTE: On an M320 or T Series router using an Adaptive Services (AS) II PIC or a MultiServices PIC, corrupted IP packets might be sent to the port mirror when traffic passes through an IPSec tunnel. The inbound IP traffic passes through the IPSec tunnel and the `sp` interface is decoded and forwarded to the port mirror correctly, but the return outbound traffic is corrupted and unreadable through the router configured with the port mirror.

The port-mirrored copy of the traffic can travel only to a single next hop. As a result, only one type of analysis can be performed if the packets are sent to a packet analyzer through a physical next hop. If more than one type of analysis is desired, a tunnel interface must be used as the next hop for port mirroring. When the mirrored copy of the traffic arrives at the virtual tunnel interface, it can be filtered, split into groups, and redirected to multiple exit interfaces and packet analyzers.

For your input requirements, include the `rate` and `run-length` statements at the `[edit forwarding-options port-mirroring family family-name input]` hierarchy level. For your output requirements, specify the target interface with the `interface` statement at the `[edit forwarding-options port-mirroring family family-name output]` hierarchy level.

By default, a filter cannot be applied to an interface where port-mirrored traffic is received. To allow the tunnel services interface to be used as a filtered next hop, include the `no-filter-check` statement at the `[edit forwarding-options port-mirroring family family-name output]` hierarchy level.

```
[edit]
forwarding-options {
  port-mirroring {
    family (inet | inet6) {
      input {
        rate 1;
        run-length 5;
      }
      output {
        interface vt-0/2/0.0;
        no-filter-check;
      }
    }
  }
}
```



NOTE: Before Junos OS Release 7.4, you could configure the input and output statements at the [edit forwarding-options port-mirroring] hierarchy level. However, this older syntax has been revised to extend port-mirroring support to IPv6 packets. If you have a configuration that contains the older syntax, we recommend that you update your configuration to the new syntax listed above.

Creating a Firewall Filter on an M, MX or T Series Router to Split the Port-Mirrored Traffic into Different Instances

If you need to split the copy of the monitored traffic into separate groups and send these filtered packets to different analyzers, devise a firewall filter that selects some traffic for sampling and some traffic for discarding. In this case, UDP traffic is sent into one routing instance, TCP traffic is diverted into a second routing instance, and all other traffic is discarded. In a later step, you will define the filter-based forwarding routing instances specified in the **then** statements shown in this filter.

```
[edit]
firewall {
  family inet {
    filter tunnel-interface-filter {
      term tcp {
        from {
          protocol tcp;
        }
        then {
          count tcp;
          routing-instance tcp-routing-table;
        }
      }
      term udp {
        from {
          protocol udp;
        }
        then {
          count udp;
          routing-instance udp-routing-table;
        }
      }
      term rest {
        then {
          count rest;
          discard;
        }
      }
    }
  }
}
```

Applying the Firewall Filter to a Tunnel PIC Interface

Once the firewall filter is defined, apply it as an input filter on a tunnel interface. This is required if the firewall filter defines two or more types of traffic or export interfaces. However, if the firewall filter only specifies one type of traffic and one export interface, you can apply the filter directly to the export interface.

```
[edit]
interfaces {
  vt-0/2/0 {
    unit 0 {
      family inet {
        filter {
          input tunnel-interface-filter;
        }
      }
    }
  }
}
```

Using Filter-Based Forwarding to Export Monitored Traffic to Multiple Destinations

The firewall filter called **tunnel-interface-filter** sends UDP traffic into one filter-based forwarding routing instance called **udp-routing-table**, sends TCP traffic into a second filter-based forwarding routing instance called **tcp-routing-table**, and discards all other packets. Here you will configure the filter-based forwarding instances.

Configure an export interface for each of your routing instances by including a static next hop. To configure, include the **route** statement at the **[edit routing-instances *instance-name* routing-options static]** hierarchy level and specify a next-hop address or interface.

```
[edit]
routing-instances {
  tcp-routing-table {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop es-3/1/0.0;
      }
    }
  }
  udp-routing-table {
    instance-type forwarding;
    routing-options {
      static {
        route 0.0.0.0/0 next-hop 10.9.1.2;
      }
    }
  }
}
```


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 Feature 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 73](#)
- [Configuring Input Interfaces, Monitoring Services Interfaces and Export Interfaces on M, MX or T Series Routers](#)
- [Establishing a VRF Instance on an M, MX or T Series Router for Monitored Traffic on page 77](#)
- [Configuring a Monitoring Group on an M, MX or T Series Router to Send Traffic to the Flow Server](#)
- [Configuring Policy Options on M, MX or T Series Routers](#)
- [Stripping MPLS Labels on ATM, Ethernet-Based and SONET/SDH Router Interfaces](#)

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”](#) on page 70.

```

[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 on M, MX or T Series Routers with Port Mirroring and Filter-Based Forwarding” on page 67.



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 74.

```
[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 24 on page 82](#).

Table 24: 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.

Table 24: Name Format Macros (continued)

Field	Expansion
{day}	01 to 31
{day_abbr}	Sun through Sat
{day_full}	Sunday through Saturday
{generation_number}	Expands to a unique, sequential number for each new file created.
{hour_12}	01 to 12
{hour_24}	00 to 23
{ifalias}	Expands to a description string for the logical interface.
{minute}	00 to 59
{month}	01 to 12
{month_abbr}	Jan through Dec
{month_full}	January through December
{num_zone}	-2359 to +2359
{second}	00 to 60
{time}	Expands to the time the file is created, using the {hour_24}, {minute}, and {second} macros.
{time_zone}	Time zone code name of the locale (gmt, pst, and so on).
{year}	1970, 2008, and so on.
{year_abbr}	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 87.

```
[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

Figure 18: Flow Collector Interface Topology Diagram




Figure 18 on page 87 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-%L_%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      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      Bytes      Bytes
           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 25 on page 95](#) explains the various fields available in the transfer log.

Table 25: 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 file naming 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 26 on page 96](#) explains the various fields available in the flow collector interface file.

Table 26: Flow Collector Interface File Fields in Order of Appearance

Field	Explanation
<code>linkDir</code>	Link directory—A randomly generated number used to identify the record
<code>analyzer-address</code>	Analyzer address

Table 26: Flow Collector Interface File Fields in Order of Appearance (continued)

Field	Explanation
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
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
```


PART 3

Active Flow Monitoring

- [Understanding Active Flow Monitoring on page 101](#)
- [System Requirements for Active Flow Monitoring on page 105](#)
- [Configuring Active Flow Monitoring on page 111](#)

CHAPTER 6

Understanding Active Flow Monitoring

- [Understanding Active Flow Monitoring Versions on M40e, M320 and T Series Routers on page 101](#)
- [When to Use Active Flow Monitoring Applications on M30, M40e, MX Series and T Series Routers on page 102](#)

Understanding Active Flow Monitoring Versions on M40e, M320 and T Series Routers

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

- [Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch on page 3](#)
- [Understanding Passive Flow Monitoring on T Series, M Series and MX Series Routers on page 37](#)

When to Use Active Flow Monitoring Applications on M30, M40e, MX Series and T Series Routers

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.

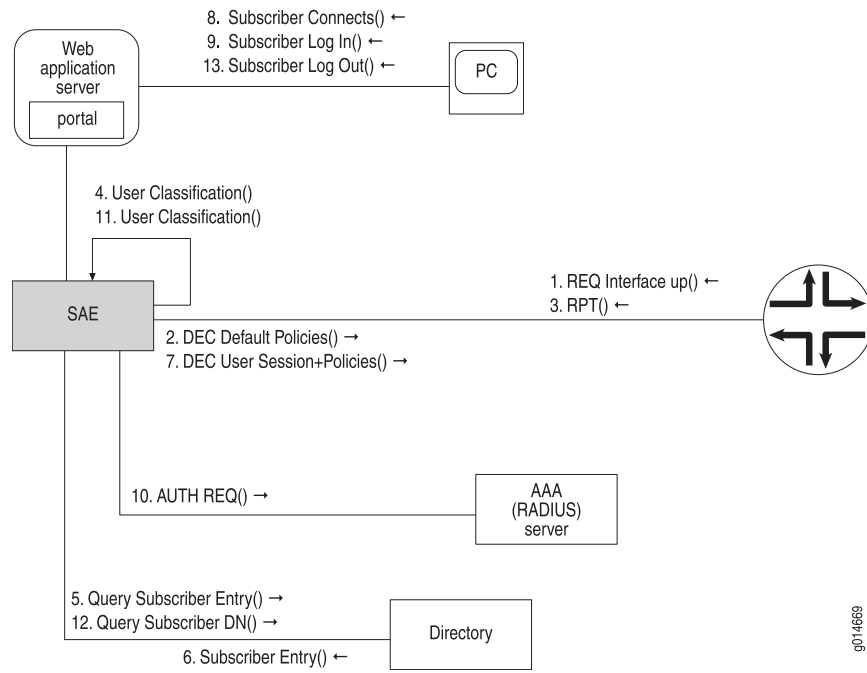
Some examples of the types of things you can use flow monitoring for are:

- 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 19: Active Flow Monitoring**Related Documentation**

- [Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch on page 3](#)
- [Understanding Active Flow Monitoring Versions on M40e, M320 and T Series Routers on page 101](#)

CHAPTER 7

System Requirements for Active Flow Monitoring

- [Active Flow Monitoring System Requirements for M and T Series Routers on page 105](#)
- [Active Flow Monitoring PIC Specifications for M and T Series Routers on page 106](#)

Active Flow Monitoring System Requirements for M and T Series Routers

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 on page 39](#)
- [Active Flow Monitoring PIC Specifications for M and T Series Routers on page 106](#)

Active Flow Monitoring PIC Specifications for M and T Series Routers

For Monitoring Services PIC specifications, see [Table 27 on page 106](#) and [Table 28 on page 107](#). For Adaptive Services PIC specifications, see [Table 29 on page 107](#). For MultiServices PIC specifications, see [Table 30 on page 107](#) and [Table 31 on page 108](#).

Table 27: 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 28: Monitoring Services II PIC Specifications

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	N/A
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 flow collector is not running. • Green—The flow collector is running under acceptable load. • Amber—The flow collector is overloaded.

Table 29: Adaptive Services PIC Specifications

Specification	Description
Physical dimensions	Single-wide PIC that occupies one PIC slot
Connectors	N/A
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 flow collector is not running. • Green—The flow collector is running under acceptable load. • Amber—The flow collector is overloaded.

Table 30: MultiServices 100 PIC

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

Table 30: MultiServices 100 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.

Table 31: 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 32: MultiServices 500 PIC

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

Table 32: 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 on page 39](#)
- [Active Flow Monitoring System Requirements for M and T Series Routers on page 105](#)

CHAPTER 8

Configuring Active Flow Monitoring

- [Understanding Active Flow Monitoring PICS and Options on M, MX and T Series Routers on page 112](#)
- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)
- [Configuring Active Flow Monitoring on PTX Series Packet Transport Routers on page 121](#)
- [Sending Packets to a Mediation Device on MX, M and T Series Routers on page 123](#)
- [Understanding Flow-Tap Architecture on page 124](#)
- [Configuring a Flow-Tap Interface on MX, M and T Series Routers on page 125](#)
- [Configuring Flow-Tap Security Properties on MX, M and T Series Routers on page 126](#)
- [Flow-Tap Application Restrictions on page 127](#)
- [Example: Flow-Tap Configuration on T and M Series Routers on page 127](#)
- [Configuring FlowTapLite on MX Series Routers and M320 Routers with FPCs on page 129](#)
- [Example: Configuring IPv6 Support for FlowTapLite on an M120 Router With Enhanced III FPCs on page 131](#)
- [Configuring Actively Monitored Interfaces on M, MX and T Series Routers on page 138](#)
- [Collecting Flow Records on M, MX and T Series Routers on page 139](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with an Accounting Group on page 139](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with a Sampling Group on page 140](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with a Template on page 141](#)
- [Defining a Firewall Filter on M, MX and T Series Routers to Select Traffic for Active Flow Monitoring on page 143](#)
- [Processing IPv4 traffic on an M, MX or T Series Router Using Monitoring services, Adaptive services or Multiservices Interfaces on page 143](#)
- [Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers on page 144](#)
- [Replicating Version 9 Flow Aggregation From M, MX and T Series Routers to Multiple Flow Servers on page 145](#)

- [Configuring Routing Engine-Based Sampling on M, MX and T Series Routers for Export to Multiple Flow Servers on page 146](#)
- [Example: Copying Traffic to a PIC While an M, MX or T Series Router Forwards the Packet to the Original Destination on page 147](#)
- [Configuring an Aggregate Export Timer on M, MX and T Series Routers for Version 8 Records on page 168](#)
- [Rerouting Packets on an M, MX or T Series Router with Port Mirroring on page 168](#)
- [Load-balancing Traffic Across PICs for Active Flow Monitoring on M, MX and T Series Routers on page 169](#)
- [Sending Port-Mirrored Traffic from an M, MX or T Series Router to Multiple Export Interfaces by Using Next-Hop Groups on page 169](#)
- [Example: Configuring Multiple Port Mirroring with Next-Hop Groups on M, MX and T Series Routers on page 170](#)
- [Example: Sampling Configuration for M, MX and T Series Routers on page 174](#)
- [Example: Sampling Instance Configuration on an MX480 Router on page 178](#)
- [Example: Sampling and Discard Accounting Configuration on M, MX and T Series Routers on page 184](#)

Understanding Active Flow Monitoring PICs and Options on M, MX and T Series Routers

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 33 on page 112](#) 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 33: 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 33: 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 143](#)
- [Configuring Actively Monitored Interfaces on M, MX and T Series Routers on page 138](#)
- [Processing IPv4 traffic on an M, MX or T Series Router Using Monitoring services, Adaptive services or Multiservices Interfaces on page 143](#)
- [Collecting Flow Records on M, MX and T Series Routers on page 139](#)
- [Rerouting Packets on an M, MX or T Series Router with Port Mirroring on page 168](#)
- [Load-balancing Traffic Across PICs for Active Flow Monitoring on M, MX and T Series Routers on page 169](#)
- [Sending Port-Mirrored Traffic from an M, MX or T Series Router to Multiple Export Interfaces by Using Next-Hop Groups on page 169](#)
- [Sending Packets to a Mediation Device on MX, M and T Series Routers on page 123](#)

Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250

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.

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, each family can support only one collector.
- 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 available in four hierarchies 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*). If you are configuring sampling of IPv4 flows, IPv6 flows or VPLS flows, 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, and VPLS flow sampling. 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-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: 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 would trigger 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 the template to specify output properties.

- a. Configure the template name.

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

- b. (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
```

- c. (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
```

- d. (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
```

- e. (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
```

- f. 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** is for version 10 templates only.

- g. (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.

- h. (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.

- 4. 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 **fpc 0**:

```
[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
16.1R1	Starting with Junos OS Release 16.1R1 and 15.1F2, the IPv4 table is assigned a default value of 1024.
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

Related Documentation

- *Configuring Inline Active Flow Monitoring on MX80 and MX104 Routers*
- *Example: Configuring Inline Active Flow Monitoring on MX Series and T4000 Routers*
- [inline-jflow on page 289](#)

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](#)

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 dynamic flow capture, see [“Using a Dynamic Flow Capture Interface on M, MX and T Series Routers to Monitor Traffic On Demand” on page 59](#). 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 124](#)
- [Configuring a Flow-Tap Interface on MX, M and T Series Routers on page 125](#)
- [Configuring Flow-Tap Security Properties on MX, M and T Series Routers on page 126](#)
- [Flow-Tap Application Restrictions on page 127](#)
- [Example: Flow-Tap Configuration on T and M Series Routers on page 127](#)

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;
    }
  }
}
```

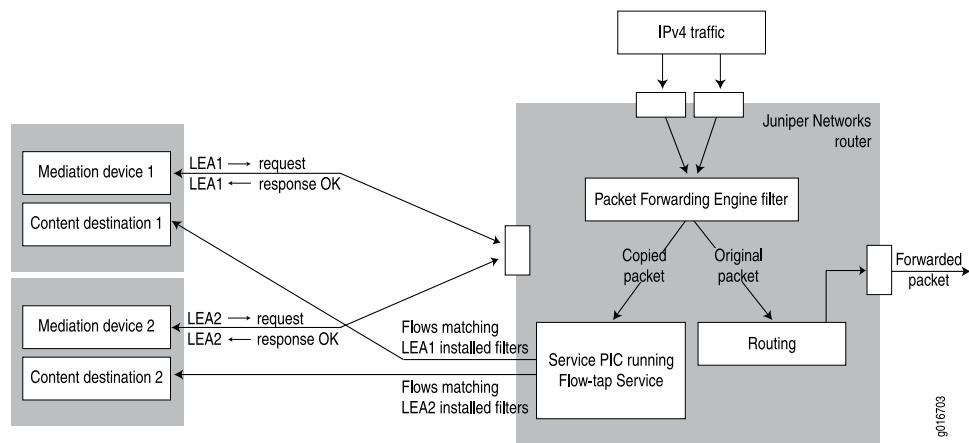
```

    }
}

```

Figure 20 on page 125 shows a sample topology that uses two mediation devices and two content destinations.

Figure 20: Flow-Tap Topology Diagram



Related Documentation

- [Configuring a Flow-Tap Interface on MX, M and T Series Routers on page 125](#)
- [Configuring Flow-Tap Security Properties on MX, M and T Series Routers on page 126](#)
- [Flow-Tap Application Restrictions on page 127](#)
- [Example: Flow-Tap Configuration on T and M Series Routers on page 127](#)

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 on page 124](#)
 - [Configuring Flow-Tap Security Properties on MX, M and T Series Routers on page 126](#)
 - [Flow-Tap Application Restrictions on page 127](#)
 - [Example: Flow-Tap Configuration on T and M Series Routers on page 127](#)

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 on page 124](#)
 - [Configuring a Flow-Tap Interface on MX, M and T Series Routers on page 125](#)
 - [Flow-Tap Application Restrictions on page 127](#)
 - [Example: Flow-Tap Configuration on T and M Series Routers on page 127](#)

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 on page 124](#)
- [Configuring a Flow-Tap Interface on MX, M and T Series Routers on page 125](#)
- [Configuring Flow-Tap Security Properties on MX, M and T Series Routers on page 126](#)
- [Example: Flow-Tap Configuration on T and M Series Routers on page 127](#)

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 on page 124](#)
- [Configuring a Flow-Tap Interface on MX, M and T Series Routers on page 125](#)
- [Configuring Flow-Tap Security Properties on MX, M and T Series Routers on page 126](#)
- [Flow-Tap Application Restrictions on page 127](#)

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.



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.

For more information about this configuration, see the *Junos OS Administration Library*.

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
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 on MX, M and T Series Routers*
- *Configuring Junos Packet Vision on MX, M and T Series Routers*
- *Examples: Configuring Junos Packet Vision on M, T, and MX Series Routers*
- *Subscriber Secure Policy Overview*

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

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: 2001:0DB8:ABCD:EF12:3456:78AB:ABC8:1235/112
Dest-Address: afte::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.

This example describes how to configure IPv6 support for FlowTapLite on an M120 router:

- [Requirements on page 132](#)
- [Overview and Topology on page 133](#)
- [Configuration on page 133](#)
- [Verification on page 136](#)

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

Figure 21 on page 133 shows the FlowTapLite configuration for one M120 router to lawfully intercept packets.

Figure 21: 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

To configure IPv6 FlowTapLite on an M120 router, perform these tasks:

- [Configuring User Credentials on page 134](#)
- [Configuring the Tunnel Interface for FlowTapLite on page 134](#)
- [Configuring the Logical Tunnel Interface on page 134](#)
- [Configuring FlowTapLite on page 135](#)
- [Results on page 135](#)

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 "xxxxxx"
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 password used by the mediation device:

```
[edit system]
user@router# set login user ftap authentication encrypted-password xxxxxx
```

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 "xxxxxx"; ## 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

To confirm that the configuration is working properly, perform the following tasks:

- [Verifying That the Router Received the Filter Request on page 136](#)
- [Checking That Filters Are Installed and Working on the Router on page 136](#)
- [Sending a List Request on page 137](#)

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

Action Use the **show filter** and the **show filter index** commands to check that filters are installed:

```

ADPC2(diving vty)# 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_

```

```
ADPC2(diving vty)# 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 on page 129](#)
- [flow-tap on page 277](#)
- [Tunnel Interface Configuration on MX Series Routers Overview](#)

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 on M, MX and T Series Routers

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 **version 9 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 15.

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 140](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with an Accounting Group on page 139](#)
- [Configuring M, MX and T Series Routers for Discard Accounting with a Template on page 141](#)
- [Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers on page 144](#)
- [Replicating Version 9 Flow Aggregation From M, MX and T Series Routers to Multiple Flow Servers on page 145](#)
- [Configuring Routing Engine-Based Sampling on M, MX and T Series Routers for Export to Multiple Flow Servers on page 146](#)
- [Configuring an Aggregate Export Timer on M, MX and T Series Routers for Version 8 Records on page 168](#)

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
          option-refresh-rate seconds 10; # The default is 60 seconds and the range
            # is 10–600 seconds.
          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

- [Understanding Active Flow Monitoring on an M Series Multiservice Edge, T Series Core Router or EX9200 Switch on page 3](#)
- [Understanding Active Flow Monitoring Versions on M40e, M320 and T Series Routers on page 101](#)
- [When to Use Active Flow Monitoring Applications on M30, M40e, MX Series and T Series Routers on page 102](#)
- [Replicating M, MX and T Series Routing Engine-Based Sampling to Multiple Flow Servers on page 144](#)

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

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 **sampld** 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 on page 147](#)
- [Overview and Topology on page 148](#)
- [Configuration on page 148](#)
- [Verification on page 164](#)

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

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).

Figure 22: Routing Engine-Based Sampling Network Topology



Configuration

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 on page 148](#)
- [Configuring the PE1 Router on page 150](#)
- [Configuring the PE2 Router on page 156](#)
- [Configuring the CE2 Router on page 162](#)

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 sampled
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 **sampled** log file in the **/var/log** directory to record traffic sampling:

```
[edit forwarding-options]
user@router-pe2# set sampling traceoptions file sampled
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;
        port bgp;
      }
      then accept;
    }
    term 2 {
      then {
        sample;
        accept;
      }
    }
  }
}
}
routing-instances {
  vrf1 {
    instance-type vrf;
    interface ge-3/1/2.0;
    route-distinguisher 10.255.19.12:1;
    vrf-import vpna-import;
    vrf-export vpna-export;
    protocols {
      bgp {
        group R3-R4 {
          type external;
          peer-as 65000;
          as-override;
          neighbor 10.10.10.1;
        }
      }
    }
  }
}
}

```

Configuring the CE2 Router

Step-by-Step Procedure

In this step, you configure interfaces, routing options, protocols, and policy options for the CE2 router. To configure the CE2 router:

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

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 on page 165](#)
- [Verifying Sampled Traffic on page 165](#)
- [Cross Verifying Sampled Traffic on page 166](#)

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 M and T Series Routers](#)” and “[Example: Configuring a Flow Collector Interface on an M, MX or T Series Router](#)” on [page 87](#).”

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 165](#) 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.13vpn.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 165](#) 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 165](#) task, indicating that the intended traffic is being sampled.

Related Documentation

- [Configuring Traffic Sampling on MX, M and T Series Routers](#)
-

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;
    }
  }
}
```

Rerouting Packets on an M, MX or T Series Router with Port Mirroring

You can copy packets and reroute them to another interface by using port mirroring. To send packet copies to an interface, include the **interface** statement at the **[edit forwarding-options port-mirroring family *family-name* output]** hierarchy level and specify the interface to receive the traffic.

You can even send port-mirrored traffic to a monitoring services or adaptive services interface. If you choose this option, accepted traffic is copied and the packet copies are sent to the services interface for flow processing.

To configure how often packets are copied from the monitored traffic, include the **rate** statement at the **[edit forwarding-options port-mirroring family *family-name* input]** hierarchy level. A rate of 1 port-mirrors every packet, while a rate of 10 port-mirrors every tenth packet.

```
[edit]
forwarding-options {
  port-mirroring {
    family (inet | inet6) {
      input {
        rate 1;
      }
      output {
        interface sp-2/0/0.0;
      }
    }
  }
}
```

Load-balancing Traffic Across PICs for Active Flow Monitoring on M, MX and T Series Routers

For active flow monitoring, you can load-balance traffic across multiple Monitoring Services PICs using the same method as passive flow monitoring. The only difference is that you do not configure the input interface with the **passive-monitor-mode** statement at the **[edit interfaces *interface-name*]** hierarchy level.

To load-balance traffic for active flow monitoring, port-mirror the incoming packets to a tunnel services interface. Redirect this copy of the traffic to a filter-based forwarding instance by applying a firewall filter to the tunnel services interface. Configure the instance to send the traffic to a group of monitoring services interfaces. Finally, use a monitoring group to send flow records from the monitoring services interfaces to a flow server.



NOTE: When you load-balance port-mirrored traffic across several Monitoring Services interfaces, there are some limitations:

- The original Monitoring Services PIC supports this method. You cannot use a Monitoring Services II PIC.
- You must use the suite of **show passive-monitoring** commands to monitor traffic. The **show services accounting** commands are not supported.
- Because load-balanced traffic is routed through the Tunnel Services PIC, the total throughput of the load-balanced traffic coming from the Monitoring Services PICs cannot exceed the bandwidth of the tunnel interface.

For detailed information on this method, see “[Copying and Redirecting Traffic on M, MX or T Series Routers with Port Mirroring and Filter-Based Forwarding](#)” on page 67.

Sending Port-Mirrored Traffic from an M, MX or T Series Router to Multiple Export Interfaces by Using Next-Hop Groups

To send port-mirrored traffic to multiple flow servers or packet analyzers, you can use the **next-hop-group** statement. The router can make up to 16 copies of traffic per group and send the traffic to the next-hop group members you configure. A maximum of 30 groups can be configured on a router at any given time. The port-mirrored traffic can be sent to any interface, except aggregated SONET/SDH, aggregated Ethernet, loopback (**lo0**), or administrative (**fxp0**) interfaces. To configure multiple port mirroring with next-hop groups, include the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level.

You must port-mirror the initial traffic to a tunnel interface so that it can be filtered and duplicated. Also, you need configure only the interface names for point-to-point interfaces, but you must configure the interface names and a next hop for multipoint interfaces (such as Ethernet).

[edit]

```
forwarding-options {
  port-mirroring {
    family inet {
      input {
        rate 1;
      }
      output {
        interface vt-3/3/0.1;
        no-filter-check;
      }
    }
  }
  next-hop-group ftp-traffic {
    interface so-4/3/0.0;
    interface so-0/3/0.0;
  }
  next-hop-group http-traffic {
    interface ge-1/1/0.0 {
      next-hop 10.12.1.2;
    }
    interface ge-1/2/0.0 {
      next-hop 10.13.1.2;
    }
  }
  next-hop-group default-collect {
    interface so-7/0/0.0;
    interface so-7/0/1.0;
  }
}
```

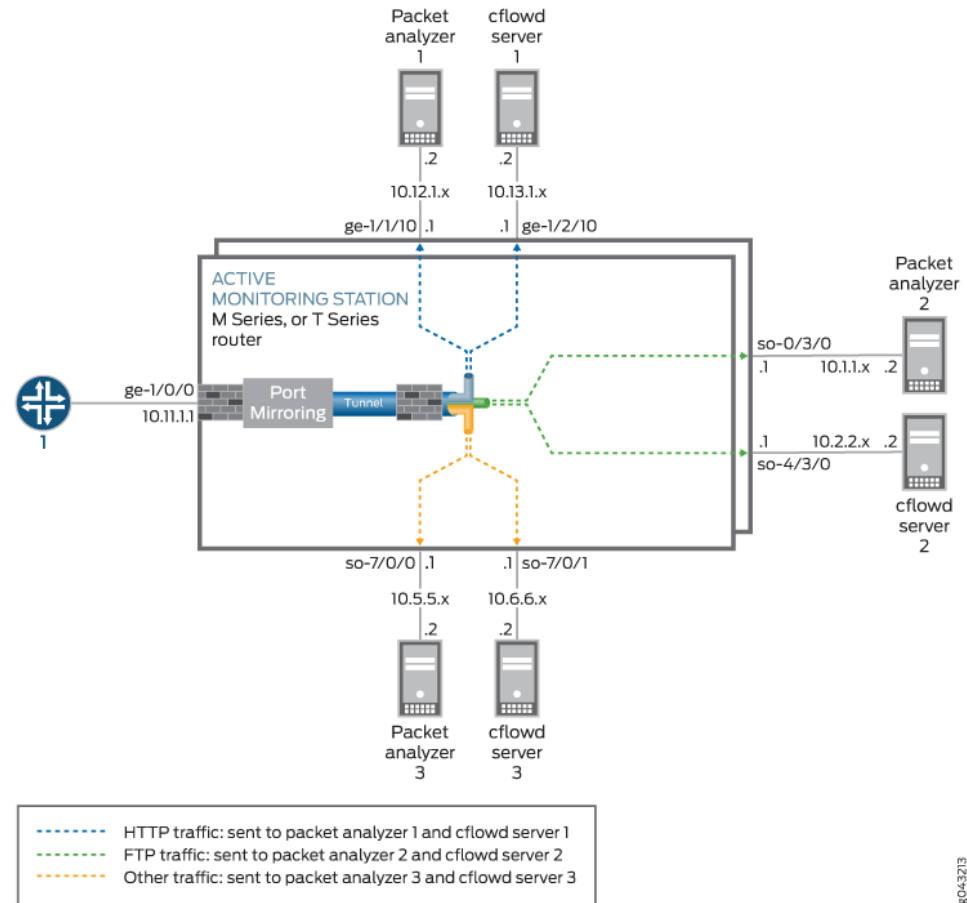


NOTE: Next-hop groups are supported on M Series routers only, except the M120 router and the M320 router.

Example: Configuring Multiple Port Mirroring with Next-Hop Groups on M, MX and T Series Routers

When you need to analyze traffic containing more than one packet type, or you wish to perform multiple types of analysis on a single type of traffic, you can implement multiple port mirroring and next-hop groups. You can make up to 16 copies of traffic per group and send the traffic to next-hop group members. A maximum of 30 groups can be configured on a router at any given time. The port-mirrored traffic can be sent to any interface, except aggregated SONET/SDH, aggregated Ethernet, loopback (lo0), or administrative (**fxp0**) interfaces. To send port-mirrored traffic to multiple flow servers or packet analyzers, you can use the **next-hop-group** statement at the **[edit forwarding-options]** hierarchy level.

Figure 23: Active Flow Monitoring—Multiple Port Mirroring with Next-Hop Groups Topology Diagram



“Active Flow Monitoring” on page 99 shows an example of how to configure multiple port mirroring with next-hop groups. All traffic enters the monitoring router at interface ge-1/0/0. A firewall filter counts and port-mirrors all incoming packets to a Tunnel Services PIC. A second filter is applied to the tunnel interface and splits the traffic into three categories: HTTP traffic, FTP traffic, and all other traffic. The three types of traffic are assigned to three separate next-hop groups. Each next-hop group contains a unique pair of exit interfaces that lead to different groups of packet analyzers and flow servers.



NOTE: Instances enabled to mirror packets to different destinations from the same PFE, also use different sampling parameters for each instance. When we configure Layer2 Port-mirroring with both global port-mirroring and instance based port-mirroring, PIC level instances will override FPC level and the FPC level will override the Global instance.

```
[edit]
interfaces {
```

```
ge-1/0/0 { # This is the input interface where packets enter the router.
  unit 0 {
    family inet {
      filter {
        input mirror_pkts; # Here is where you apply the first filter.
      }
      address 10.11.1.1/24;
    }
  }
}
ge-1/1/0 { # This is an exit interface for HTTP packets.
  unit 0 {
    family inet {
      address 10.12.1.1/24;
    }
  }
}
ge-1/2/0 { # This is an exit interface for HTTP packets.
  unit 0 {
    family inet {
      address 10.13.1.1/24;
    }
  }
}
so-0/3/0 { # This is an exit interface for FTP packets.
  unit 0 {
    family inet {
      address 10.1.1.1/30;
    }
  }
}
so-4/3/0 { # This is an exit interface for FTP packets.
  unit 0 {
    family inet {
      address 10.2.2.1/30;
    }
  }
}
so-7/0/0 { # This is an exit interface for all remaining packets.
  unit 0 {
    family inet {
      address 10.5.5.1/30;
    }
  }
}
so-7/0/1 { # This is an exit interface for all remaining packets.
  unit 0 {
    family inet {
      address 10.6.6.1/30;
    }
  }
}
vt-3/3/0 { # The tunnel interface is where you send the port-mirrored traffic.
  unit 0 {
    family inet;
  }
}
```

```

    unit 1 {
        family inet {
            filter {
                input collect_pkts; # This is where you apply the second firewall filter.
            }
        }
    }
}
forwarding-options {
    port-mirroring { # This is required when you configure next-hop groups.
        family inet {
            input {
                rate 1; # This port-mirrors all packets (one copy for every packet received).
            }
            output { # Sends traffic to a tunnel interface to enable multipoint mirroring.
                interface vt-3/3/0.1;
                no-filter-check;
            }
        }
    }
}
next-hop-group ftp-traffic { # Point-to-point interfaces require you to specify the
    interface so-4/3/0.0; # interface name.
    interface so-0/3/0.0;
}
next-hop-group http-traffic { # Configure a next hop for all multipoint interfaces.
    interface ge-1/1/0.0 {
        next-hop 10.12.1.2;
    }
    interface ge-1/2/0.0 {
        next-hop 10.13.1.2;
    }
}
next-hop-group default-collect {
    interface so-7/0/0.0;
    interface so-7/0/1.0;
}
}
firewall {
    family inet {
        filter mirror_pkts { # Apply this filter to the input interface.
            term catch_all {
                then {
                    count input_mirror_pkts;
                    port-mirror; # This action sends traffic to be copied and port-mirrored.
                }
            }
        }
        filter collect_pkts { # Apply this filter to the tunnel interface.
            term ftp-term { # This term sends FTP traffic to an FTP next-hop group.
                from {
                    protocol ftp;
                }
                then next-hop-group ftp-traffic;
            }
            term http-term { # This term sends HTTP traffic to an HTTP next-hop group.

```

```

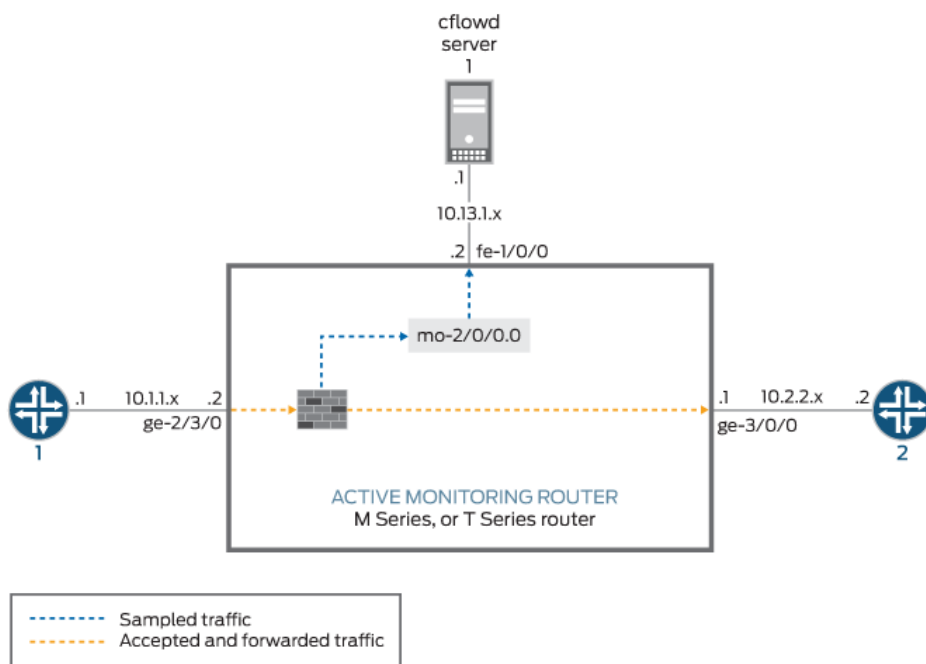
        from {
            protocol http;
        }
        then next-hop-group http-traffic;
    }
    term default { # This sends all remaining traffic to a final next-hop group.
        then next-hop-group default-collectors;
    }
}
}
}

```

- Related Documentation**
- *Understanding Port Mirroring on Routers with an Internet Processor II ASIC or T Series Internet Processor*
 - *Configuring Port Mirroring on M, T MX, and PTX Series Routers*

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

Figure 24: Active Flow Monitoring—Sampling Configuration Topology Diagram



In Figure 24 on page 174, 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%
```

Example: Sampling Instance Configuration on an MX480 Router

You can configure active sampling using a sampling instance and associate that sampling instance to a particular FPC, MPC, or DPC. In addition, you can define multiple sampling instances associated with multiple destinations and protocol families per sampling instance destination.



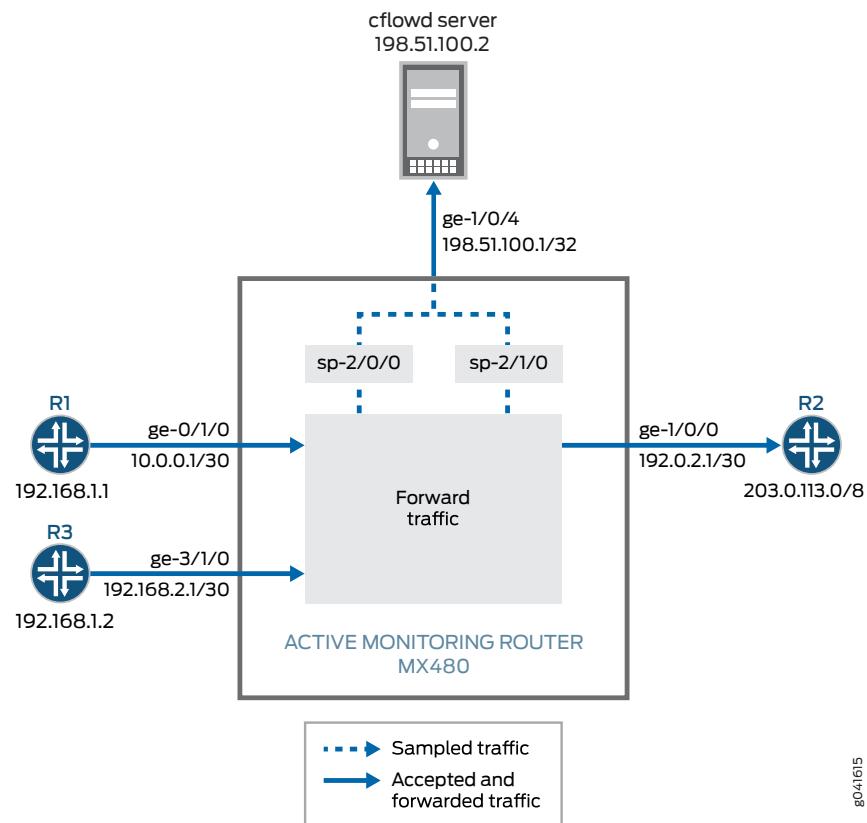
NOTE: This procedure also works on T Series routers, M320, M120, EX Series switches, and other MX Series routers.

- [Example Network Details on page 178](#)
- [Example Router Configuration on page 180](#)
- [Configuration Commands Used for the Configuration Example on page 182](#)
- [Verifying Your Work on page 183](#)

Example Network Details

The following example shows the configuration of two sampling instances on an MX480 router running Junos OS Release 9.6.

Figure 25: Active Flow Monitoring—Sampling Instance Configuration Topology Diagram



In Figure 25 on page 179, 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
```

Source Address	Destination Address	Src Dst		Proto	TOS	Packet Count
		Port/ ICMP Type	Port/ ICMP Code			
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
```

Source Address	Destination Address	Src Dst		Proto	TOS	Packet Count
		Port/ ICMP Type	Port/ ICMP Code			
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
```

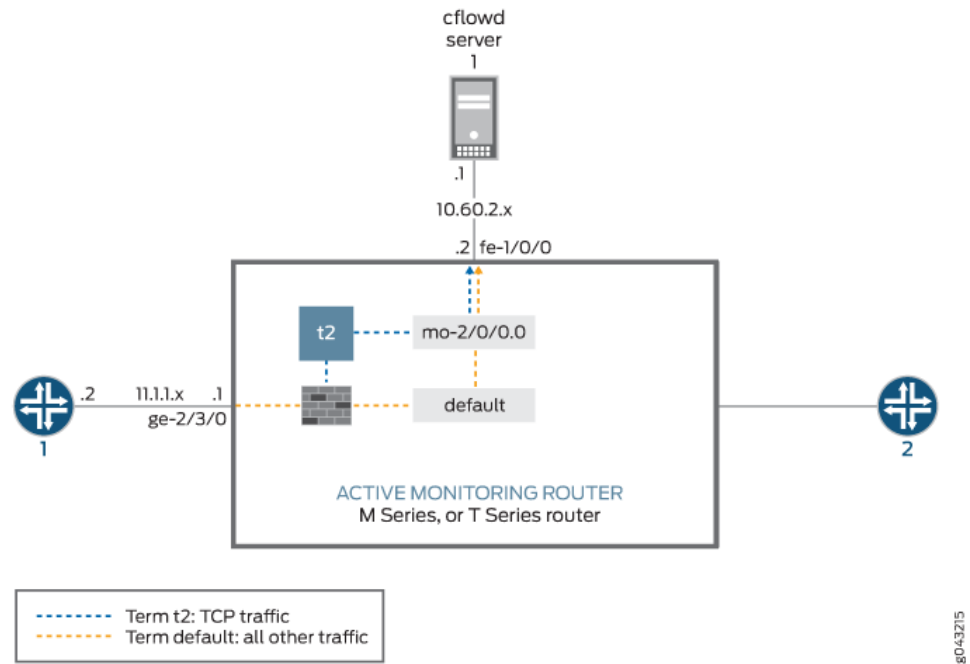
Related Documentation

- *Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches*
- [Understanding Active Flow Monitoring PICS and Options on M, MX and T Series Routers on page 112](#)
- *sampling-instance*

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

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 26: Active Flow Monitoring—Sampling and Discard Accounting Topology Diagram



8043215

In [Figure 26 on page 185](#), 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 10.1.1.1/20;
        }
    }
}
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

PART 4

Configuration Statements and Operational Commands

- [Configuration Statements for Routers on page 193](#)
- [Basic Flow and Active Flow EX9200 Monitoring Configuration Statements on page 457](#)
- [Basic Flow and Active Flow EX9200 Monitoring Operational Commands on page 483](#)
- [Active Flow Monitoring Commands on page 513](#)
- [Dynamic Flow Capture Commands on page 549](#)
- [Flow Collection Commands on page 559](#)
- [Passive Flow Monitoring Commands on page 575](#)

CHAPTER 9

Configuration Statements for Routers

- [accounting](#) on page 201
- [address \(Interfaces\)](#) on page 202
- [address \(Services Dynamic Flow Capture\)](#) on page 202
- [aggregate-export-interval](#) on page 203
- [aggregation](#) on page 204
- [alarms](#) on page 205
- [alarm-mode](#) on page 206
- [allowed-destinations](#) on page 207
- [analyzer-address](#) on page 207
- [analyzer-id](#) on page 208
- [archive-sites](#) on page 208
- [authentication-mode](#) on page 209
- [authentication-key-chain \(TWAMP\)](#) on page 210
- [autonomous-system-type](#) on page 211
- [bandwidth-kbps \(RFC 2544 Benchmarking\)](#) on page 212
- [bgp](#) on page 213
- [capture-group](#) on page 214
- [cflowd \(Discard Accounting\)](#) on page 215
- [cflowd \(Flow Monitoring\)](#) on page 216
- [client](#) on page 217
- [client-delegate-probes](#) on page 218
- [client-list](#) on page 219
- [collector](#) on page 219
- [collector \(Flow Monitoring Logs for NAT\)](#) on page 220
- [collector \(Flow Template Profiles for NAT\)](#) on page 221
- [collector-group \(Flow Template Profiles for NAT\)](#) on page 222
- [collector-group \(Flow Monitoring Logs for NAT\)](#) on page 223
- [content-destination](#) on page 224

- [control-connection](#) on page 225
- [control-source](#) on page 226
- [core-dump](#) on page 227
- [data-fill](#) on page 228
- [data-fill-with zeros](#) on page 229
- [data-format](#) on page 229
- [data-size](#) on page 230
- [delay-factor](#) on page 231
- [delegate-probes](#) on page 232
- [destination \(Interfaces\)](#) on page 233
- [destination-address \(Flow Monitoring Logs for NAT\)](#) on page 234
- [destination-interface](#) on page 235
- [destination-ipv4-address \(RFC 2544 Benchmarking\)](#) on page 236
- [destination-mac-address \(RFC2544 Benchmarking\)](#) on page 237
- [destination-port](#) on page 238
- [destination-port \(Flow Monitoring Logs for NAT\)](#) on page 239
- [destination-udp-port \(RFC 2544 Benchmarking\)](#) on page 240
- [destinations](#) on page 240
- [direction \(RFC2544 Benchmarking\)](#) on page 241
- [disable \(Forwarding Options\)](#) on page 242
- [disable-signature-check \(RFC 2544 Benchmarking\)](#) on page 243
- [dscp \(flow-server\)](#) on page 244
- [dscp-code-point \(Services\)](#) on page 245
- [duplicates-dropped-periodicity](#) on page 246
- [dynamic-flow-capture](#) on page 247
- [engine-id \(Forwarding Options\)](#) on page 248
- [engine-type](#) on page 249
- [export-format](#) on page 250
- [extension-service](#) on page 251
- [family \(Monitoring\)](#) on page 253
- [family \(Port Mirroring\)](#) on page 254
- [family \(RFC2544 Benchmarking\)](#) on page 255
- [family \(Sampling\)](#) on page 256
- [file \(Sampling\)](#) on page 258
- [file \(Trace Options\)](#) on page 258
- [file-specification \(File Format\)](#) on page 259
- [file-specification \(Interface Mapping\)](#) on page 259

- [filename on page 260](#)
- [filename-prefix on page 260](#)
- [files on page 261](#)
- [filter on page 262](#)
- [flex-flow-sizing on page 263](#)
- [flow-active-timeout on page 264](#)
- [flow-collector on page 266](#)
- [flow-export-destination on page 267](#)
- [flow-export-rate on page 268](#)
- [flow-inactive-timeout on page 269](#)
- [flow-key \(Flow Monitoring\) on page 270](#)
- [flow-monitoring on page 271](#)
- [flow-server on page 273](#)
- [flow-table-size on page 275](#)
- [flow-table-size \(Chassis\) on page 276](#)
- [flow-tap on page 277](#)
- [forwarding-class \(Sampling\) on page 278](#)
- [ftp \(Flow Collector Files\) on page 279](#)
- [ftp \(Transfer Log Files\) on page 280](#)
- [g-duplicates-dropped-periodicity on page 281](#)
- [g-max-duplicates on page 282](#)
- [generate-snmp-traps on page 283](#)
- [hard-limit on page 283](#)
- [hard-limit-target on page 284](#)
- [hardware-timestamp on page 284](#)
- [history-size on page 285](#)
- [host-outbound media-interface on page 286](#)
- [in-service \(RFC2544 Benchmarking\) on page 287](#)
- [inactivity-timeout \(Services RPM\) on page 288](#)
- [inline-jflow on page 289](#)
- [input \(Port Mirroring\) on page 290](#)
- [input \(Sampling\) on page 290](#)
- [input-interface-index on page 291](#)
- [input-packet-rate-threshold on page 291](#)
- [instance \(Sampling\) on page 292](#)
- [interface \(Accounting or Sampling\) on page 293](#)
- [interfaces on page 294](#)

- [interface \(Services Flow Tap\) on page 294](#)
- [interface-map on page 295](#)
- [interfaces \(Services Dynamic Flow Capture\) on page 295](#)
- [interfaces \(Video Monitoring\) on page 296](#)
- [inet6-options \(Services\) on page 299](#)
- [ip-swap \(RFC 2544 Benchmarking\) on page 300](#)
- [ipv4-flow-table-size on page 301](#)
- [ipv4-template on page 302](#)
- [ipv6-flow-table-size on page 303](#)
- [ipv6-extended-attr on page 304](#)
- [ipv6-template on page 304](#)
- [jflow-log \(Interfaces\) on page 305](#)
- [jflow-log \(Services\) on page 306](#)
- [label-position on page 307](#)
- [license-server on page 308](#)
- [local-dump on page 309](#)
- [logical-system on page 309](#)
- [match on page 310](#)
- [max-connection-duration on page 310](#)
- [max-duplicates on page 311](#)
- [max-packets-per-second on page 312](#)
- [maximum-age on page 313](#)
- [maximum-connections on page 314](#)
- [maximum-connections-per-client on page 315](#)
- [maximum-packet-length on page 316](#)
- [maximum-sessions on page 318](#)
- [maximum-sessions-per-connection on page 319](#)
- [media-loss-rate on page 320](#)
- [media-rate-variation on page 321](#)
- [message-rate-limit \(Flow Monitoring Logs for NAT\) on page 322](#)
- [minimum-priority on page 323](#)
- [mode \(RFC 2544 Benchmarking\) on page 323](#)
- [monitoring on page 324](#)
- [moving-average-size on page 325](#)
- [mpls-flow-table-size on page 326](#)
- [mpls-ipv4-template on page 327](#)
- [mpls-ipvx-template on page 328](#)

- [mpls-template](#) on page 329
- [multiservice-options](#) on page 330
- [name-format](#) on page 331
- [next-hop \(Forwarding Options\)](#) on page 332
- [next-hop-group \(Forwarding Options\)](#) on page 333
- [next-hop-group \(Port Mirroring\)](#) on page 334
- [nexthop-learning](#) on page 335
- [no-filter-check](#) on page 336
- [no-remote-trace \(Trace Options\)](#) on page 336
- [no-syslog](#) on page 337
- [no-syslog-generation](#) on page 337
- [notification-targets](#) on page 338
- [observation-domain-id](#) on page 339
- [one-way-hardware-timestamp](#) on page 340
- [option-refresh-rate](#) on page 341
- [options-template-id](#) on page 342
- [output \(Accounting\)](#) on page 343
- [output \(Monitoring\)](#) on page 344
- [output \(Port Mirroring\)](#) on page 345
- [output \(Sampling\)](#) on page 346
- [output-interface-index](#) on page 347
- [packet-size \(RFC 2544 Benchmarking\)](#) on page 348
- [passive-monitor-mode](#) on page 349
- [password \(Flow Collector File Servers\)](#) on page 349
- [password \(Transfer Log File Servers\)](#) on page 350
- [peer-as-billing-template](#) on page 350
- [pic-memory-threshold](#) on page 351
- [pop-all-labels](#) on page 352
- [port \(Flow Monitoring\)](#) on page 353
- [port \(RPM\)](#) on page 353
- [port \(TWAMP\)](#) on page 354
- [port-mirroring](#) on page 355
- [post-cli-implicit-firewall](#) on page 356
- [pre-rewrite-tos](#) on page 357
- [probe](#) on page 358
- [probe-count](#) on page 360
- [probe-interval](#) on page 361

- [probe-limit](#) on page 362
- [probe-server](#) on page 363
- [probe-type](#) on page 364
- [profiles \(RFC 2544 Benchmarking\)](#) on page 365
- [rate \(Forwarding Options\)](#) on page 366
- [receive-options-packets](#) on page 367
- [receive-ttl-exceeded](#) on page 367
- [refresh-rate \(Flow Monitoring Logs for NAT\)](#) on page 368
- [reflect-mode \(RFC2544 Benchmarking\)](#) on page 369
- [reflect-etype \(RFC 2544 Benchmarking\)](#) on page 370
- [required-depth](#) on page 371
- [retry \(Services Flow Collector\)](#) on page 372
- [retry-delay](#) on page 372
- [rfc2544-benchmarking](#) on page 373
- [rfc6514-compliant-safi129 \(Protocols BGP\)](#) on page 374
- [routing-instance](#) on page 375
- [routing-instance \(cflowd\)](#) on page 376
- [routing-instance-list \(TWAMP\)](#) on page 377
- [routing-instances](#) on page 378
- [rpm \(Interfaces\)](#) on page 379
- [rpm \(Services\)](#) on page 380
- [rpm-scale](#) on page 383
- [run-length](#) on page 385
- [sample-once](#) on page 386
- [sampling \(Forwarding Options\)](#) on page 387
- [sampling \(Interfaces\)](#) on page 390
- [server](#) on page 391
- [server-inactivity-timeout](#) on page 392
- [service-port](#) on page 392
- [service-type \(RFC2544 Benchmarking\)](#) on page 393
- [services](#) on page 394
- [services](#) on page 395
- [services-options](#) on page 396
- [shared-key](#) on page 397
- [size](#) on page 398
- [slamon-services](#) on page 399
- [soft-limit](#) on page 400

- [soft-limit-clear](#) on page 400
- [source-address](#) (Forwarding Options) on page 401
- [source-address](#) (Services) on page 402
- [source-addresses](#) on page 403
- [source-id](#) on page 403
- [source-ip](#) (Flow Monitoring Logs for NAT) on page 404
- [source-ipv4-address](#) (RFC 2544 Benchmarking) on page 405
- [source-mac-address](#) (RFC2544 Benchmarking) on page 406
- [source-udp-port](#) (RFC 2544 Benchmarking) on page 407
- [stamp](#) on page 407
- [storm-control](#) on page 408
- [syslog](#) on page 409
- [target](#) (Services RPM) on page 410
- [tcp](#) on page 411
- [template](#) (Flow Monitoring IPFIX Version) on page 412
- [template](#) (Flow Monitoring Version 9) on page 413
- [template](#) (Forwarding Options) on page 414
- [template](#) (Forwarding Options Version IPFIX) on page 414
- [template-id](#) on page 415
- [template-profile](#) (Flow Monitoring Logs for NAT) on page 416
- [template-refresh-rate](#) on page 417
- [template-type](#) (Flow Monitoring Logs for NAT) on page 418
- [templates](#) on page 419
- [test](#) on page 421
- [tests](#) (RFC 2544 Benchmarking) on page 423
- [test-interface](#) (RFC 2544 Benchmarking) on page 424
- [test-interval](#) on page 425
- [test-name](#) (RFC 2544 Benchmarking) on page 426
- [test-profile](#) (RFC 2544 Benchmarking) on page 427
- [test-session](#) on page 428
- [test-type](#) (RFC 2544 Benchmarking) on page 429
- [thresholds](#) on page 430
- [traceoptions](#) (Dynamic Flow Capture) on page 431
- [traceoptions](#) (Forwarding Options) on page 432
- [traceoptions](#) (RPM) on page 433
- [transfer](#) on page 434
- [transfer-log-archive](#) on page 435

- [traps on page 436](#)
- [ttl on page 437](#)
- [twamp on page 438](#)
- [twamp-server on page 439](#)
- [trio-flow-offload on page 440](#)
- [tunnel-observation on page 441](#)
- [udp on page 442](#)
- [udp-tcp-port-swap \(RFC 2544 Benchmarking\) on page 443](#)
- [unit on page 444](#)
- [use-extended-flow-memory on page 445](#)
- [username \(Services\) on page 446](#)
- [variant on page 446](#)
- [version on page 447](#)
- [version \(Flow Monitoring Logs for NAT\) on page 448](#)
- [version9 \(Forwarding Options\) on page 449](#)
- [version9 \(Flow Monitoring\) on page 450](#)
- [version-ipfix \(Forwarding Options\) on page 451](#)
- [version-ipfix \(Services\) on page 452](#)
- [video-monitoring on page 453](#)
- [vpls-flow-table-size on page 455](#)
- [vpls-template on page 456](#)
- [world-readable on page 456](#)

accounting

Syntax	<pre> accounting <i>name</i> { output { aggregate-export-interval <i>seconds</i>; cflowd <i>hostname</i> { aggregation { autonomous-system; destination-prefix; protocol-port; source-destination-prefix { caida-compliant; } source-prefix; } autonomous-system-type (origin peer); port <i>port-number</i>; version <i>format</i>; } flow-active-timeout <i>seconds</i>; flow-inactive-timeout <i>seconds</i>; interface <i>interface-name</i> { engine-id <i>number</i>; engine-type <i>number</i>; source-address <i>address</i>; } } } </pre>
Hierarchy Level	[edit forwarding-options]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Specify the discard accounting instance name and options.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Discard Accounting on M and T Series Routers</i>

address (Interfaces)

Syntax	<code>address address { destination address; }</code>
Hierarchy Level	[edit <code>interfaces interface-name unit logical-unit-number family family</code>]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other options not associated with flow monitoring.• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

address (Services Dynamic Flow Capture)

Syntax	<code>address address;</code>
Hierarchy Level	[edit <code>services dynamic-flow-capture capture-group client-name content-destination identifier</code>]
Release Information	Statement introduced in Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Content Destination</i>

aggregate-export-interval

Syntax	<code>aggregate-export-interval <i>seconds</i>;</code>
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]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the duration, in seconds, of the interval for exporting aggregate accounting information.
Options	<i>seconds</i> —Duration.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Discard Accounting on M and T Series Routers</i>

aggregation

Syntax	<pre>aggregation { autonomous-system; destination-prefix; protocol-port; source-destination-prefix { caida-compliant; } source-prefix; }</pre>
Hierarchy Level	[edit forwarding-options accounting output cflowd hostname], [edit forwarding-options sampling instance instance-name family (inet inet6 impls) output flow-server hostname], [edit forwarding-options sampling family (inet inet6 impls) output flow-server hostname]
Release Information	Statement introduced before Junos OS Release 7.4.
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	<p>autonomous-system—Aggregate by autonomous system (AS) number.</p> <p>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 <i>cflowd Configuration Guide</i>, dated August 30, 1999.</p> <p>destination-prefix—Aggregate by destination prefix.</p> <p>protocol-port—Aggregate by protocol and port number.</p> <p>source-destination-prefix—Aggregate by source and destination prefix.</p> <p>source-prefix—Aggregate by source prefix.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Enabling Flow Aggregation on T and M Series Routers</i>

alarms

```
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]

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- *Understanding Inline Video Monitoring on MX Series Routers*
- [delay-factor on page 231](#)
- [no-syslog-generation on page 337](#)
- [generate-snmp-traps on page 283](#)
- [storm-control on page 408](#)
- [alarm-mode on page 206](#)
- [media-rate-variation on page 321](#)
- [media-loss-rate on page 320](#)

alarm-mode

Syntax alarm-mode {
 mdi-records-count *number*;
 average;
}

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- *Understanding Inline Video Monitoring on MX Series Routers*
- [alarms on page 205](#)

allowed-destinations

Syntax	<code>allowed-destinations [<i>destinations</i>];</code>
Hierarchy Level	[edit services dynamic-flow-capture <code>capture-group</code> <i>client-name</i> <code>control-source</code> <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Identify flow capture destinations that are allowed in messages sent from this control source.
Options	<i>destinations</i> —Allowed content destination name.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Control Source</i>

analyzer-address

Syntax	<code>analyzer-address <i>address</i>;</code>
Hierarchy Level	[edit services flow-collector]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an IP address for the packet analyzer that overrides the default value.
Options	<i>address</i> —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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring a Packet Analyzer</i>

analyzer-id

Syntax	<code>analyzer-id name;</code>
Hierarchy Level	[edit services flow-collector]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure an identifier for the packet analyzer that overrides the default value.
Options	<i>name</i> —Identifier for packet analyzer.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring a Packet Analyzer</i>

archive-sites

Syntax	<pre>archive-sites { ftp:url { password "password"; username username; } }</pre>
Hierarchy Level	[edit services flow-collector transfer-log-archive]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Transfer Logs</i>

authentication-mode

Syntax	<code>authentication-mode (authenticated encrypted none);</code>
Hierarchy Level	[edit services rpm twamp server], [edit services rpm twamp client control-connection <i>control-client-name</i>]
Release Information	Statement introduced in Junos OS Release 9.5. Support at the [edit <code>services rpm twamp client control-connection control-client-name</code>] hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i> • <i>Understanding Two-Way Active Measurement Protocol on Routers</i>

authentication-key-chain (TWAMP)

Syntax	<pre>authentication-key-chain <i>identifier</i> { key-id <i>identifier</i> { secret <i>password-string</i>; } }</pre>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 9.5.
Description	Apply and enable an authentication keychain 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	<p><i>identifier</i>—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").</p> <p><i>password-string</i>—Authentication key, consisting of 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>


autonomous-system-type

Syntax	<code>autonomous-system-type (origin peer);</code>
Hierarchy Level	[edit forwarding-options <code>sampling instance</code> <i>instance-name</i> <code>family</code> (inet inet6 mpls) <code>output flow-server</code> <i>hostname</i>], [edit forwarding-options <code>sampling family</code> (inet inet6 mpls) <code>output flow-server</code> <i>hostname</i>]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.
Description	Specify the type of AS numbers that cflowd exports.
Default	<code>origin</code>
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	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Flow Aggregation on T and M Series Routers</i>

bandwidth-kbps (RFC 2544 Benchmarking)

Syntax	<code>bandwidth-kbps <i>kbps</i>;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking profile test-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
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	<i>kbps</i> —Bandwidth limit, in kilobits per second (kbps). Range: 1,000 kbps through 1,000,000 Kbps.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>RFC 2544-Based Benchmarking Tests Overview</i>• <i>Configuring RFC 2544-Based Benchmarking Tests</i>• rfc2544-benchmarking on page 373

bgp

Syntax	<pre> 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-safil29; test-interval interval; } </pre>
Hierarchy Level	<pre> [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] </pre>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure BGP neighbor discovery through Real-Time Performance Monitoring (RPM).
Options	<p>bgp—Define properties for configuring BGP neighbor discovery.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
<div>  <p>NOTE: On MX Series routers, you can configure all the statements. On M Series and T Series routers, you can configure only the <code>logical-system</code> and <code>routing-instances</code> statements.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i>

capture-group

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]

Release Information Statement introduced in Junos OS Release 7.4.

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.

Related Documentation • *Configuring the Capture Group*

cflowd (Discard Accounting)

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**]

Release Information Statement introduced before Junos OS Release 7.4.

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 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.

Related Documentation

- *Enabling Flow Aggregation on T and M Series Routers*

cflowd (Flow Monitoring)

Syntax	<code>cflowd hostname { port port-number; }</code>
Hierarchy Level	[edit forwarding-options monitoring name inet output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility <code>cfcollect</code>.</p> <p>You can configure up to eight version 5 flow formats at the [edit forwarding-options monitoring name output] hierarchy level. Version 8 flow formats are not supported for flow-monitoring applications.</p>
Options	<p>hostname—IP address or identifier of the host system (the workstation running the <code>cflowd</code> utility).</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Enabling Flow Aggregation on T and M Series Routers</i>

client

```
Syntax  client {
        control-connection control-client- name {
            authentication-mode
            destination-interface interface-name;
            destination-port port;
            history-size size;
            moving-average-size number;
            routing-instance instance-name;
            target (url url | address address);
            test-interval interval;
            traps traps;
            data-fill-with zeros
            data-size size;
            dscp-code-point (Services) dscp-bits;
            probe-count count;
            probe-interval seconds;
            thresholds thresholds;
            test-session session-name{
                data-fill-with zeros data;
                data-size size;
                dscp-code-point (Services) dscp-bits;
                probe-count count;
                probe-interval seconds;
                target (url url | address address);
            }
        }
    }
```

Hierarchy Level [edit services rpm twamp]

Release Information Statement introduced in Junos OS Release 15.1.

Description Specify the TWAMP client configuration settings.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Understanding Two-Way Active Measurement Protocol on Routers*

client-delegate-probes

Syntax	<code>rpm client-delegate-probes;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	Statement introduced in Junos OS Release 17.3R1 on MX Series routers.
Description	<p>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.</p> <p>The destination-interface statement must be configured at the <code>[edit services rpm probe <i>owner</i> test <i>test-name</i>]</code> 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 <code>[edit services rpm probe <i>owner</i>]</code> hierarchy level to complete the configuration.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>

client-list

Syntax	<code>client-list <i>list-name</i> { address <i>address</i>; }</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 9.3.
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	<p><i>list-name</i>—Name of client address list.</p> <p><i>address</i>—Address and mask for an allowed client.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

collector

Syntax	<code>collector <i>interface-name</i>;</code>
Hierarchy Level	[edit services flow-collector interface-map]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the default flow collector interface for interface mapping.
Options	<i>interface-name</i> —Default flow collector interface.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Interface Mappings</i>

collector (Flow Monitoring Logs for NAT)

Syntax	<pre>collector <i>collector-name</i> { source-ip <i>address</i>; destination-address <i>address</i>; destination-port <i>port-number</i>; }</pre>
Hierarchy Level	[edit services jflow-log]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>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.</p>
Options	<p><i>collector-name</i>—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_]</p> <p>The remaining statements are described separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

collector (Flow Template Profiles for NAT)

Syntax	<code>collector <i>collector-name</i>;</code>
Hierarchy Level	<code>[edit services jflow-log template-profile <i>template-profile-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 15.1.
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 <code>[edit services jflow-log]</code> 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 <code>[a-zA-Z0-9_]</code>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i> • <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i> • <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i> • <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

collector-group (Flow Template Profiles for NAT)

Syntax	<code>collector-group collector-group-name;</code>
Hierarchy Level	[edit services jflow-log template-profile <i>template-profile-name</i>]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>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.</p>
Options	<p>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_]</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

collector-group (Flow Monitoring Logs for NAT)

Syntax	<code>collector-group collector-group-name { [collector-name1 collector-name2]; }</code>
Hierarchy Level	[edit services jflow-log]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>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.</p> <p>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.</p>
Options	<p>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_]</p> <p>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_]</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i> • <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i> • <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i> • <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

content-destination

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]`

Release Information Statement introduced in Junos OS Release 7.4.

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.

Related Documentation

- *Configuring the Content Destination*

control-connection

Syntax

```
control-connection control-client-name {
  authentication-mode
  destination-interface interface-name;
  destination-port port;
  history-size size;
  moving-average-size number;
  routing-instance instance-name;
  target (url url | address address);
  test-interval interval;
  traps traps;
  data-fill-with zeros data;
  data-size size;
  dscp-code-point (Services) dscp-bits;
  probe-count count;
  probe-interval seconds;
  thresholds thresholds;
  test-session session-name{
    data-fill-with zeros data;
    data-size size;
    dscp-code-point (Services) dscp-bits;
    probe-count count;
    probe-interval seconds;
    target (url url | address address);
  }
}
```

Hierarchy Level [edit services rpm twamp client]

Release Information Statement introduced in Junos OS Release 15.1.

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. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.


Related Documentation

- *Understanding Two-Way Active Measurement Protocol on Routers*

control-source

Syntax	<pre>control-source <i>identifier</i> { allowed-destinations [<i>destinations</i>]; minimum-priority <i>value</i>; no-syslog; notification-targets <i>address</i> port <i>port-number</i>; service-port <i>port-number</i>; shared-key <i>value</i>; source-addresses [<i>addresses</i>]; }</pre>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Identify the source of the dynamic flow capture request.
Options	<p><i>identifier</i>—Name of control source.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Control Source</i>

core-dump

Syntax	(core-dump no-core-dump);
Hierarchy Level	[edit interfaces mo-fpc/pic/port multiservice-options]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>A useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory <code>/var/tmp</code> 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):</p> <hr/> <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p>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.</p> </div> </div> <hr/> <ul style="list-style-type: none"> • core-dump—Enable the core dumping operation. • no-core-dump—Disable the core dumping operation.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

data-fill

Syntax	<code>data-fill data;</code> <code>data-fill-with-zeros data;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test test-name], [edit services rpm twamp client control-connection control-client-name]
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 9.3 for PTX Series Packet Transport routers. 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 Release 18.1 for QFX Series switches.
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. For TWAMP client, if this knob is set, then fill the test packet with zeros, if the knob is not set then the data content is random value as indicated in RFC.
Options	data —A hexadecimal value; for example, 0-9 , A-F .
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i>• <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>


data-fill-with-zeros

Syntax	<code>data-fill-with-zeros;</code>
Hierarchy Level	[edit services rpm twamp client control-connection <i>control-client-name</i> test-session <i>session-name</i>]
Release Information	Statement introduced in Junos OS Release 15.1 for MX Series routers.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Two-Way Active Measurement Protocol on Routers</i>

data-format

Syntax	<code>data-format <i>format</i>;</code>
Hierarchy Level	[edit services flow-collector file-specification variant <i>variant-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the data format for a specific file format variant.
Options	<i>format</i> —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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring File Formats</i>


data-size

Syntax	<code>data-size size;</code>
Hierarchy Level	<code>[edit services rpm bgp],</code> <code>[edit services rpm probe owner test test-name],</code> <code>[edit services rpm twamp client control-connection <i>control-client-name</i> test-session session-name]</code>
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 <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> 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.
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.
<div> NOTE: If you configure the hardware timestamp feature (see <i>Configuring RPM Timestamping on MX, M and T Series Routers and EX Series Switches</i>):</div> <ul style="list-style-type: none">• 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. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i>

delay-factor

Syntax	<pre> delay-factor { no-syslog-generation; generate-snmp-traps; storm-control { count <i>number</i>; interval <i>number</i>; } alarm-mode { mdi-records-count <i>number</i>; average; } } </pre>
Hierarchy Level	[edit services]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>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.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Understanding Inline Video Monitoring on MX Series Routers</i> alarms on page 205

delegate-probes

Syntax	delegate-probes;
Hierarchy Level	[edit services rpm probe owner]
Release Information	Statement introduced in Junos OS Release 17.3R1 on MX Series routers.
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>To avoid packet bursts in the network due to RPM, probes will be distributed in a better way.</p> <p>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.</p>
	<div> 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.</div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>• client-delegate-probes on page 218

destination (Interfaces)

Syntax	<code>destination address;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For CoS on ATM interfaces, specify the remote address of the connection.</p> <p>For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.</p> <p>For tunnel and encryption interfaces, specify the remote address of the tunnel.</p>
Options	address —Address of the remote side of the connection.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Linear RED Profiles on ATM Interfaces</i> • <i>Multilink and Link Services Logical Interface Configuration Overview</i> • <i>Configuring Encryption Interfaces</i> • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i> • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8 • <i>Configuring Unicast Tunnels</i>

destination-address (Flow Monitoring Logs for NAT)

Syntax	<code>destination-address <i>address</i>;</code>
Hierarchy Level	<code>[edit services jflow-log collector <i>collector-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>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.</p>
Options	<i>address</i> —Destination hostname, or IPv4 or IPv6 address of the collector.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

destination-interface

Syntax	<code>destination-interface <i>interface-name</i>;</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>],</code> <code>[edit services rpm probe-server (tcp udp)],</code> <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 7.5. Support at the <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.
Description	<p>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.</p> <p>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.</p> <p>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.</p> <p>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 <code>[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]</code> hierarchy level. For the extension-provider package, package-name in the package package-name statement is jservices-rpm.</p> <p>Starting in Junos OS Release 17.3R1, you can use destination-interface <i>interface-name.logical-unit-number</i> at the <code>[edit services rpm probe owner test <i>test-name</i>]</code> 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 <code>[edit services rpm probe owner]</code> hierarchy level and the rpm client-delegate-probes and the family (inet inet6) address address statements at the <code>[edit interfaces interface-name unit logical-unit-number]</code> hierarchy level.</p>
Options	<i>interface-name</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Timestamping on MX, M and T Series Routers and EX Series Switches</i>• <i>Configuring RPM Receiver Servers on M, MX, T and PTX Series Routers and EX Series Switches</i>• <i>Configuring RPM Timestamping on MX, M and T Series Routers and EX Series Switches</i>• hardware-timestamp on page 284• rpm (Interfaces) on page 379• <i>Enabling RPM on MX, M and T Series Routers and SRX Firewalls for the Services SDK</i>


destination-ipv4-address (RFC 2544 Benchmarking)

Syntax	destination-ipv4-address <i>address</i> ;
Hierarchy Level	[edit services rpm rfc2544-benchmarkingtests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.
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 cccas the family.
Options	address —Valid IPv4 address. Default: If you do not configure the destination IPv4 address, the default value of 192.168.1.20 is used.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>• <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>• rfc2544-benchmarking on page 373

destination-mac-address (RFC2544 Benchmarking)

Syntax	<code>destination-mac-address mac-address;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers.
Description	Specify the destination MAC address used in the generated test frames. This is a mandatory parameter for family bridge .
Options	mac-address —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> —for example, 0000:5e00:5355 or 00:00:5e:00:53:55 .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • rfc2544-benchmarking on page 373 • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

destination-port

Syntax	<code>destination-port <i>port</i>;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test <i>test-name</i>], [edit services rpm twamp client control-connection <i>control-client-name</i>]
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 <i>control-client-name</i>] 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.
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 .
Options	Default: The default value for the port is 862 to which the TWAMP client establishes control connection. port —Port number 7 or from 49,160 through 65,535.
<div>  NOTE: The specified port numbers are recommended for RPM only. </div>	
Required Privilege Level	system—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i> <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>

destination-port (Flow Monitoring Logs for NAT)

Syntax	<code>destination-port <i>port-number</i>;</code>
Hierarchy Level	<code>[edit services jflow-log collector <i>collector-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 15.1.
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	<p><i>port-number</i>—UDP port number for the test frames.</p> <p>Default: 4041</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i> • <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i> • <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i> • <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

destination-udp-port (RFC 2544 Benchmarking)

Syntax	<code>destination-udp-port <i>port-number</i>;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name <i>test-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.
Description	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.
Options	<i>port-number</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>• <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>• rfc2544-benchmarking on page 373



destinations

Syntax	<pre>destinations { ftp:url { password "<i>password</i>"; } }</pre>
Hierarchy Level	[edit services flow-collector]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Destination FTP Servers for Flow Records</i>

direction (RFC2544 Benchmarking)

Syntax	<code>direction (egress ingress);</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.
Description	Specify the direction of the interface on which the test must be run. This parameter is valid only for a ccc family and 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	<p>egress—Run the test in the egress direction of the interface (network-to-network interface (NNI)). This option is applicable for a ccc and bridge family.</p> <p>ingress—Run the test in the ingress direction of the interface (user-to-network interface (UNI)). You cannot configure this option for a bridge family.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • rfc2544-benchmarking on page 373 • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

disable (Forwarding Options)

Syntax	disable;
Hierarchy Level	[edit forwarding-options port-mirror], [edit forwarding-options port-mirror instance <i>instance-name</i>], [edit forwarding-options sampling], [edit forwarding-options sampling instance <i>instance-name</i>], [edit forwarding-options sampling family (inet inet6 mpls vpls)], [edit forwarding-options sampling family (inet inet6 mpls vpls) output file]
Release Information	Statement introduced before Junos OS Release 7.4. Statement added to port-mirror hierarchy in Junos OS Release 9.6.
	<div>  <p>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 <i>instance-name</i> 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 <i>instance-name</i> family (inet inet6 mpls) command to prevent sampling.</p> </div>
Description	Disable traffic accounting, port mirroring, or sampling.
	<div>  <p>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 <i>instance-name</i>] hierarchy level.</p> </div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Disabling Traffic Sampling <ul style="list-style-type: none"> Configuring Traffic Sampling on MX, M and T Series Routers Configuring Port Mirroring on M, T MX, and PTX Series Routers

disable-signature-check (RFC 2544 Benchmarking)

Syntax	disable-signature-check;
Hierarchy Level	[edit services rpm rfc2544-benchmarkingtests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 15.1 for MX104 3D Universal Edge routers.
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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Supported RFC2544-Based Benchmarking Statements on MX Series Routers</i> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

dscp (flow-server)

Syntax	<code>dscp <i>dscp-value</i></code>
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6) output flow-server <i>hostname</i>]
Release Information	Statement introduced in Junos OS Release 15.1F4 for the PTX Series. Statement introduced in Junos OS Release 16.1 for the MX Series.
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	<p>dscp <i>dscp-value</i>—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.</p> <p>The <i>dscp-value</i> is overwritten by the CoS DSCP value if you configure dscp under the [edit class-of-service] hierarchy.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Inline Active Flow Monitoring on PTX Series Routers</i>• <i>Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers</i>• <i>Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches</i>

dscp-code-point (Services)

Syntax	<code>dscp-code-point <i>dscp-bits</i>;</code>
Hierarchy Level	<code>[edit services rpm probe owner test test-name],</code> <code>[edit services rpm twamp client control-connection control-client-name test-session session-name]</code>
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 <code>[edit services rpm twamp client control-connection control-client-name]</code> 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. Support for also setting the IEEE 802.1p code point on EX2300, EX3400, and EX4300 switches introduced in Junos OS Release 18.2R1.
Description	Specify the value of the Differentiated Services (DiffServ) field within the IP header of host-generated RPM packets. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern. Starting in Junos 18.2R1 for EX2300, EX3400, and EX4300 switches, the first three bits of the 6-bit pattern also set the IEEE 802.1p code point of host-generated RPM packets.
Options	<p><i>dscp-bits</i>—A valid 6-bit pattern; for example, 001111, or one of the following configured DSCP aliases:</p> <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

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches*
- *Understanding Two-Way Active Measurement Protocol on Routers*

duplicates-dropped-periodicity

Syntax duplicates-dropped-periodicity *seconds*;

Hierarchy Level [edit services dynamic-flow-capture [capture-group](#) *client-name*]

Release Information Statement introduced in Junos OS Release 9.2.

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.

Related Documentation

- [g-duplicates-dropped-periodicity on page 281](#)
- *Limiting the Number of Duplicates of a Packet*
- [max-duplicates on page 311](#)

dynamic-flow-capture

```
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]

Release Information Statement introduced in Junos OS Release 7.4.

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.


Related Documentation

- *Understanding Using Junos Capture Vision on M and T Series Routers*

engine-id (Forwarding Options)

Syntax	<code>engine-id <i>number</i>;</code>
Hierarchy Level	<code>[edit forwarding-options accounting name output interface <i>interface-name</i>],</code> <code>[edit forwarding-options monitoring name output interface <i>interface-name</i>],</code> <code>[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output</code> <code>interface <i>interface-name</i>],</code> <code>[edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the engine ID number for flow monitoring and accounting services.
Options	<i>number</i> —Identity of accounting interface.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Traffic Sampling on MX, M and T Series Routers• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8• Configuring Discard Accounting on M and T Series Routers

engine-type

Syntax	<code>engine-type <i>number</i>;</code>
Hierarchy Level	<p>[edit forwarding-options accounting name output interface <i>interface-name</i>], [edit forwarding-options monitoring name output interface <i>interface-name</i>], [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output interface <i>interface-name</i>], [edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>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.</p>
<div>  <p>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.</p> </div>	
Options	<i>number</i> —Platform-specific accounting interface type.
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Traffic Sampling on MX, M and T Series Routers • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8 • Configuring Discard Accounting on M and T Series Routers

export-format

Syntax	<code>export-format <i>format</i>;</code>
Hierarchy Level	[edit forwarding-options monitoring name output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Flow monitoring export format.
Options	<i>format</i> —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.
Related Documentation	<ul style="list-style-type: none">• version on page 447• Exporting Flows on page 11

extension-service

Syntax

```
extension-service {
  service-name {
    provider-specific rules;
  }
  application {
    argument argument-names;
    checksum number;
    daemonize;
    max-datasize datasize;
  }
  max-datasize datasize;
  traceoptions {
    file filename;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit forwarding-options **sampling instance** *instance-name* **family** (inet |inet6) **output**],
[edit forwarding-options **sampling family** (inet |inet6) **output**],
[edit services service-set *service-set-name*],
{edit system services}

Release Information Statement introduced in Junos OS Release 9.0.

Description Define a customer specific sampling configuration.

Define a service set or traffic monitoring for applications using application-specific configuration guidelines.



NOTE: If the **extension-service** statement is specified while configuring a service set, the **service-order** statement is mandatory.

Define configuration parameters for an application.

Options **argument** *argument-names*—Use the specified command line arguments to the JET application

checksum *number*—Checksum of the script.

daemonize—Run the application as a background process.

file *filename*—Use the specified name of the file to receive the output of the tracing operation. All files are placed in the directory /var/log.

flag *flag*—Use the specified tracing operation to perform:

- **all**—Trace everything.
- **config**—Trace configuration events.
- **general**—Trace general events.
- **notification**—Trace notification events.
- **routing-socket**—Trace routing socket calls.
- **thriftv**—Trace thrift server events.
- **timeouts**—Trace timeouts.
- **timer**—Trace internal timer events.

max-datasize *datasize*—Maximum data segment size allowed for application execution (23068672..1073741824 bytes).

no-remote-trace—Disable remote tracing.

provider-specific rules—Provider-specific subhierarchy for services and service sets. See the application-specific documentation for details.

service-name—Use the specified name of the service.

Required Privilege Level	system —To view this statement in the configuration. system-control —To add this statement to the configuration.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• service-order• sampling on page 387
------------------------------	---

family (Monitoring)

```
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](#)]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

family (Port Mirroring)

List of Syntax	MX, M, T Series Routers, EX Series Switches and SRX Series Firewalls on page 254 Syntax: QFX Series Switches, EX4600 and NFX Series Devices on page 254
MX, M, T Series Routers, EX Series Switches and SRX Series Firewalls	<pre> family (inet inet6) { output { interface <i>interface-name</i> { next-hop <i>address</i>; } no-filter-check; } } </pre>
Syntax: QFX Series Switches, EX4600 and NFX Series Devices	<pre> family ethernet-switching { output { interface <i>interface-name</i> { } no-filter-check; } vlan <i>vlan-name</i> { no-tag; } } inet output { ip-address <i>address</i> { } routing-instance <i>instance-name</i> { ip-address <i>address</i> { } } } } </pre>
Hierarchy Level	[edit forwarding-options port-mirroring]
Release Information	Statement introduced before Junos OS Release 7.4 for MX, M, T Series routers, EX Series switches and SRX Series firewalls. Statement introduced in Junos OS Release 13.2 for the QFX Series and EX4600.
Description	Specify the type of interface that will be used to forward port mirrored packet to an analyzer device. Configure the protocol family to be sampled. Only IPv4 (inet) and IPv6 (inet6) are 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. routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring Port Mirroring on M, T MX, and PTX Series Routers*
 - *Understanding Port Mirroring*
 - *Configuring Port Mirroring*
 - *Examples: Configuring Port Mirroring for Local Analysis*

family (RFC2544 Benchmarking)

Syntax	family (bridge ccc inet);
Hierarchy Level	[edit services rpm rfc2544-benchmarkingtests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers. bridge option introduced in Junos OS Release 12.3X53 for ACX Series routers. bridge option introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers.
Description	Configure the address type family for the benchmarking test.
Options	<p>bridge—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.</p> <p>ccc—Run the test on a circuit cross-connect (CCC) or Ethernet pseudowire service. You can run the RFC2544-based benchmarking test either in the egress or ingress direction.</p> <p>inet—Run the test on an IPv4 service.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • rfc2544-benchmarking on page 373 • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>

family (Sampling)

```
Syntax family (inet | inet6 | mpls | vpls | bridge) {
    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;
        }
        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*]

- 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.
- 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.
- Related Documentation**
- *Configuring Traffic Sampling on MX, M and T Series Routers*

file (Sampling)

Syntax	<pre>file { disable; filename filename; files number; size bytes; (stamp no-stamp); (world-readable no-world-readable); }</pre>
Hierarchy Level	[edit forwarding-options sampling family inet output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Collect the traffic samples in a file.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

file (Trace Options)

Syntax	<pre>file filename <files number <size bytes> <world-readable no-world-readable>;</pre>
Hierarchy Level	[edit forwarding-options port-mirroring traceoptions], [edit forwarding-options sampling traceoptions]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure information about the files that contain trace logging information.
Options	<p>filename—Name of the file containing the trace information.</p> <p>Default: /var/log/sampled</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Tracing Traffic Sampling Operations</i>

file-specification (File Format)

Syntax	<pre>file-specification { variant <i>variant-number</i> { data-format <i>format</i>; name-format <i>format</i>; transfer { record-level <i>number</i>; timeout <i>seconds</i>; } } }</pre>
Hierarchy Level	[edit services flow-collector]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure the file format for the flow collection files.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring File Formats

file-specification (Interface Mapping)

Syntax	<pre>file-specification { variant <i>variant-number</i>; }</pre>
Hierarchy Level	[edit services flow-collector interface-map]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the default file specification for interface mapping.
Options	<i>variant-number</i> —Default file format variant.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

filename

Syntax	<code>filename <i>filename</i>;</code>
Hierarchy Level	[edit forwarding-options sampling family (inet inet6 mpls) output file]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the name of the output file.
Options	<i>filename</i> —Name of the file in which to place the traffic samples. All files are placed in the directory <code>/var/tmp</code> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

filename-prefix

Syntax	<code>filename-prefix <i>prefix</i>;</code>
Hierarchy Level	[edit services flow-collector transfer-log-archive]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the filename prefix for log files.
Options	<i>prefix</i> —Filename identifier.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Transfer Logs</i>


files

Syntax	<code>files <i>number</i>;</code>
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]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the total number of files to be saved with samples or trace data.
Options	<p><i>number</i>—Maximum number of traffic sampling or trace log files. When a file named <i>sampling-file</i> reaches its maximum size, it is renamed <i>sampling-file.0</i>, then <i>sampling-file.1</i>, and so on, until the maximum number of traffic sampling files is reached. Then the oldest sampling file is overwritten.</p> <p>Range: 1 through 100 files</p> <p>Default: 5 files for sampling output; 10 files for trace log information</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i> • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>



filter

Syntax	<pre>filter { input <i>filter-name</i>; output <i>filter-name</i>; group <i>filter-group-number</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a firewall filter to an interface. You can also use filters for encrypted traffic.
Options	<p>group <i>filter-group-number</i>—Use the specified interface to be part of a filter group. The default filter group number is 0.</p> <p>input <i>filter-name</i>—Use the specified filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Use the specified filter to evaluate when packets are transmitted on the interface.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide</i> or the <i>Junos OS Administration Library</i>• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

flex-flow-sizing

Syntax	flex-flow-sizing;
Hierarchy Level	[edit chassis fpc slot-number inline-services]
Release Information	Statement introduced in Junos OS Release 15.1F5 for MX Series routers.
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.
	<div>  <p>NOTE: You cannot configure the explicit flow-table-sizes because flex-flow-sizing and explicit flow-table-sizes are mutually exclusive.</p> <p>You need not perform fpc reboot to change from flex to per family configuration.</p> </div>
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.
Related Documentation	<ul style="list-style-type: none"> • Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115 • Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers

flow-active-timeout

Syntax	<code>flow-active-timeout <i>seconds</i>;</code>
Hierarchy Level	<p>[edit forwarding-options accounting <i>name</i> output], [edit forwarding-options monitoring <i>name</i> output], [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls vpls) output], [edit forwarding-options sampling family (inet inet6 mpls vpls) output], [edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>], [edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2.</p> <p>Support at the [edit services flow-monitoring version9 template <i>template-name</i>] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.</p>
Description	Set the interval after which an active flow is exported.
<div>  <p>NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.</p> </div>	
Options	<p>seconds—Duration of the timeout period.</p> <p>Range: 60 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)</p> <p>Default: 1800 seconds (for forwarding-options configurations); 60 seconds (for services configurations)</p>
<div>  <p>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.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- [Configuring Time Periods When Flow Monitoring Is Active and Inactive on page 12](#)
 - *Configuring the Version 9 Template Properties*
 - *Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250*

flow-collector

```
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]

Release Information Statement introduced before Junos OS Release 7.4.

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.

- Related Documentation**
- [Understanding Flow Collection Using Monitoring Services II or Multiservices 400 PIC Interface on Routers](#)


flow-export-destination

Syntax	<pre>flow-export-destination { (cflowd-collector collector-pic); }</pre>
Hierarchy Level	[edit forwarding-options monitoring group-name family inet output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure flow collection.
Options	<p>cflowd-collector—Use the cflowd collector.</p> <p>collector-pic—Use the collector PIC.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Exporting Flows on page 11


flow-export-rate

Syntax	flow-export-rate <i>rate</i> ;
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family inet output inline-jflow]
Release Information	Statement introduced in Junos OS Release 10.2.
Description	Specify the flow export rate of monitored packets in kpps. The rate specified here is applied at the level of line card, not PFE. In other words, 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.
Options	rate —Flow export rate of monitored packets in kpps (from 1 through 400). Default: 1 kpps (applies to all PFEs on the FPC)
	<div> NOTE: The maximum rate per PFE is 100, so for a FPC with four PFEs (such as AS cards) you can set a maximum flow-export-rate of 400. For a FPC with two PFEs (such as the MPC2), the maximum flow-export-rate is 200. For a FPC with one PFE (such as the MPC5), the maximum is 100. The Junos CLI may accept as valid any value within the range of 1 to 400, but when applied the value may trigger an error message such as The configured flow export rate is higher than supported value/chip in the Junos message log.</div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Discard Accounting on M and T Series Routers</i>• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

flow-inactive-timeout

Syntax	<code>flow-inactive-timeout <i>seconds</i>;</code>
Hierarchy Level	<p>[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],</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level added in Junos OS Release 10.2.</p> <p>Support at the [edit services flow-monitoring version9 template template-name] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.</p>
Description	Set the interval of inactivity that marks a flow inactive.
<div>  <p>NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.</p> </div>	
Options	<p><i>seconds</i>—Duration of the timeout period.</p> <p>Range: 15 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)</p> <p>Default: 60 seconds (for forwarding-options configurations); 60 seconds (for services configurations)</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Time Periods When Flow Monitoring Is Active and Inactive on page 12 • Configuring the Version 9 Template Properties • Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250

flow-key (Flow Monitoring)

Syntax	<pre>flow-key { flow-direction; vlan-id; output-interface; }</pre>
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 15.2. output-interface option added in Junos OS Release 18.2R1.
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	<p>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.</p> <p>vlan-id—Include VLAN IDs in both the ingress and egress directions in the flow key.</p> <p>output-interface—Configure the output-interface field as part of flow-key for bridge or VPLS family.</p>
<div> 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.</div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115• <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i>• <i>Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250</i>

flow-monitoring

```

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 {
                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](#)]

Release Information Statement introduced in Junos OS Release 8.3.

Description Specify the active monitoring properties for flow aggregation version 9 or IPFIX.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

- Related Documentation**
- *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*
 - *Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250*

flow-server

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**]

Release Information Statement introduced before Junos OS Release 7.4.
version9 statement introduced in Junos OS Release 8.3.
 Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.
 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**].

Description Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility cfdcollect. 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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

flow-table-size

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;
 bridge-flow-table-size units;
 }`

Hierarchy Level `[edit chassis fpc slot-number inline-services]`

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.

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
```



NOTE: If you are using a Junos release prior to Junos OS Release 15.1F2, this command initiates an automatic reboot of the FPC, and we recommend you run this command during a maintenance window.

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.

Related Documentation

- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)
- [Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers](#)

flow-table-size (Chassis)

Syntax	flow-table-size <i>size</i> ;
Hierarchy Level	[edit chassis fpc slot inline-video-monitoring]
Release Information	Statement introduced in Junos OS Release 16.1 on the MX Series.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Inline Video Monitoring on MX Series Routers</i>

flow-tap

Syntax	<pre>flow-tap { (interface <i>interface-name</i> tunnel-interface <i>interface-name</i>); family (inet inet6); }</pre>
Hierarchy Level	[edit services]
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>ccc option introduced in Junos OS Release 17.2.</p>
Description	<p>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.</p> <p>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.</p> <p>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.</p>
Options	<p>interface <i>interface-name</i>—Use the specified interface for the flow-tap application.</p> <p>tunnel-interface <i>interface-name</i>—Use the specified tunnel interface for the FlowTapLite application.</p> <p>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.</p> <ul style="list-style-type: none"> • inet—IPv4 traffic. • inet6—IPv6 traffic.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

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 (Sampling)

Syntax	<code>forwarding-class <i>class-name</i>;</code>
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6) output flow-server <i>hostname</i>]
Release Information	Statement introduced in Junos OS Release 16.1 on the MX Series and the PTX Series.
Description	Specify the forwarding class to which exported packets for inline active flow monitoring are sent.
Default	If you do not include the <i>forwarding-class</i> statement, exported packets are sent to the best effort queue.
Options	forwarding-class <i>class-name</i> —Name of the forwarding class: <ul style="list-style-type: none">• 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.
Related Documentation	• <i>Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches</i>

ftp (Flow Collector Files)

Syntax	<code>ftp:url;</code>
Hierarchy Level	[edit services flow-collector destination]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the primary and secondary destination FTP server addresses.
Options	<p>url—FTP server address. The URL can include the following macros, typed in braces:</p> <ul style="list-style-type: none"> • {%D}—Date • {%T}—Time when the file is created • {%I}—Description string for the logical interface configured using the collector <i>interface-name</i> 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_abbr}—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_abbrev}**—From 00 through 99

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Destination FTP Servers for Flow Records*

ftp (Transfer Log Files)

Syntax ftp:url;

Hierarchy Level [edit services flow-collector [transfer-log-archive](#) archive-sites]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- *Configuring Transfer Logs*


g-duplicates-dropped-periodicity

Syntax	<code>g-duplicates-dropped-periodicity seconds;</code>
Hierarchy Level	[edit services dynamic-flow-capture]
Release Information	Statement introduced in Junos OS Release 9.2.
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.
Related Documentation	<ul style="list-style-type: none">• duplicates-dropped-periodicity on page 246• <i>Limiting the Number of Duplicates of a Packet</i>

g-max-duplicates

Syntax	<code>g-max-duplicates <i>number</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture]
Release Information	Statement introduced in Junos OS Release 9.2.
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	<i>number</i> —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.
Related Documentation	<ul style="list-style-type: none">• max-duplicates on page 311• <i>Limiting the Number of Duplicates of a Packet</i>

generate-snmp-traps

Syntax	generate-snmp-traps;
Hierarchy Level	[edit services] [edit services video-monitoring]
Release Information	Statement introduced in Junos OS Release 15.1.
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, mdiDFAlarm trap be triggered.
<div>  NOTE: SNMP traps are not generated if SNMP trap generation is not enabled. </div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding Inline Video Monitoring on MX Series Routers</i> • alarms on page 205

hard-limit

Syntax	hard-limit <i>bandwidth</i> ;
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
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	<i>bandwidth</i> —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.
Related Documentation	<ul style="list-style-type: none"> • hard-limit-target on page 284 • <i>Configuring the Content Destination</i>

hard-limit-target

Syntax	<code>hard-limit-target <i>bandwidth</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify a bandwidth threshold at which the dynamic flow capture application stops deleting criteria.
Options	<i>bandwidth</i> —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.
Related Documentation	<ul style="list-style-type: none">• hard-limit on page 283• <i>Configuring the Content Destination</i>

hardware-timestamp

Syntax	<code>hardware-timestamp;</code>
Hierarchy Level	[edit services rpm probe <i>owner</i> test <i>test-name</i>]
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 10.3 for EX Series switches.
Description	Enable timestamping of RPM probe messages in the Packet Forwarding Engine host processor. This feature is supported only with <code>icmp-ping</code> , <code>icmp-ping-timestamp</code> , <code>udp-ping</code> , and <code>udp-ping-timestamp</code> probe types.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

history-size

Syntax	<code>history-size size;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test test-name], [edit services rpm twamp client control-connection <i>control-client-name</i>]
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 at the [edit services rpm twamp client control-connection <i>control-client-name</i>] 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.
Description	Specify the number of stored history entries.
Options	size —Value from 0 to 255. Default: 50
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i> <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>

host-outbound media-interface

Syntax	host-outbound media-interface;
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in Junos OS Release 13.2 on MX Series 3D Universal Edge routers.
Description	<p>Enable Layer 2 port mirroring of host-generated outbound packets only on MPCs on MX Series 3D Universal Edge routers.</p> <p>This statement enables all Routing Engine-generated Layer 2 injections to execute egress logical interface filters.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Examples: Layer 2 Port Mirroring at Multiple Levels of the Chassis</i>• <i>Configuring Port Mirroring</i>• <i>Understanding Layer 2 Port Mirroring</i>


in-service (RFC2544 Benchmarking)

Syntax	in-service;
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers.
Description	<p>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.</p> <p>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.</p>
Default	The default service mode for the reflecting egress interface for an E-LAN service is out-of-service mode.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • rfc2544-benchmarking on page 373 • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

inactivity-timeout (Services RPM)

Syntax	<code>inactivity-timeout <i>seconds</i>;</code>
Hierarchy Level	<code>[edit services rpm twamp server]</code>
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Inactivity timeout period, in seconds.
Options	<i>seconds</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

inline-jflow

Syntax	<pre>inline-jflow { source-address address; flow-export-rate rate; }</pre>
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family inet output]
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 14.2 for T4000 Series routers with Type 5 FPC.</p> <p>Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.</p> <p>Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.</p>
Description	<p>Specify inline flow monitoring for traffic from the designated address.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
<div>  <p>NOTE: If you configure inline flow monitoring with <code>inline-jflow</code> then you have to disable it before performing ISSU. For more information, see <i>Before You Begin a Unified ISSU</i>.</p> </div>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115

input (Port Mirroring)

Syntax	<pre>input { maximum-packet-length bytes rate number; run-length number; }</pre>
Hierarchy Level	[edit forwarding-options port-mirroring], [edit forwarding-options port-mirroring instance <i>instance-name</i>], [edit forwarding-options port-mirroring family (inet inet6)]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure port mirroring 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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i>

input (Sampling)

Syntax	<pre>input { max-packets-per-second number; rate number; run-length number; maximum-packet-length bytes; }</pre>
Hierarchy Level	[edit forwarding-options sampling], [edit forwarding-options sampling instance <i>instance-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

input-interface-index

Syntax	<code>input-interface-index <i>number</i>;</code>
Hierarchy Level	[edit forwarding-options monitoring name output interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a value for the input interface index that overrides the default supplied by SNMP.
Options	<i>number</i> —Input interface index value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

input-packet-rate-threshold

Syntax	<code>input-packet-rate-threshold <i>rate</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Specify a packet rate threshold value that triggers a system log warning message.
Options	<i>rate</i> —Threshold value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Thresholds

instance (Sampling)

```
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]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Sampling Instance on MX, M and T Series Routers or QFX Series Switches</i>

interface (Accounting or Sampling)

Syntax	<pre>interface <i>interface-name</i> { engine-id <i>number</i>; engine-type <i>number</i>; source-address <i>address</i>; }</pre>
Hierarchy Level	[edit forwarding-options accounting name output], [edit forwarding-options sampling family (inet inet6 mpls) output], [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the output interface for monitored traffic.
Options	<i>interface-name</i> —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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Discard Accounting on M and T Series Routers</i> • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

interfaces

Syntax	<code>interfaces { ... }</code>
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i>

interface (Services Flow Tap)

Syntax	<code>interface sp-fpc/pic/port.logical-unit-number;</code>
Hierarchy Level	[edit services flow-tap]
Release Information	Statement introduced in Junos OS Release 8.1.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Flow-Tap Interface</i>

interface-map

Syntax	<pre>interface-map { collector <i>interface-name</i>; file-specification <i>variant-number</i>; <i>interface-name</i> { collector <i>interface-name</i>; file-specification <i>variant-number</i>; } }</pre>
Hierarchy Level	[edit services flow-collector]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Match an input interface with a flow collector interface and apply the preset file specifications to the input interface.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Interface Mappings</i>

interfaces (Services Dynamic Flow Capture)

Syntax	interfaces <i>interface-name</i> ;
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Specify the DFC interface used with the control source configured in the same capture group.
Options	<i>interface-name</i> —Name of the DFC interface.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the DFC PIC Interface</i>

interfaces (Video Monitoring)

```

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](#)]

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.

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.

Related Documentation

- *Configuring Inline Video Monitoring on MX Series Routers*

inet6-options (Services)

Syntax

```
inet6-options {
  source-address address;
}
```

Hierarchy Level [edit **services** rpm **probe** owner **test** test-name]

Release Information Statement introduced in Junos OS Release 14.1R4.

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 interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches*

ip-swap (RFC 2544 Benchmarking)

Syntax	ip-swap;
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 14.2 for MX Series routers.
Description	Swaps source and destination IPv4 addresses. This statement is applicable only for family bridge .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>• <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>• rfc2544-benchmarking on page 373

ipv4-flow-table-size

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.

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:



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.

```
[edit]
user@router# set chassis fpc slot-number inline-services flow-table-size
ipv4-flow-table-size 15
```

Options **units**—Number of 256K flow entries available for the IPv4 flow table.
Range: 1 through 245
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
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**
- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)

ipv4-template

Syntax	ipv4-template;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2. Statement introduced in Junos OS Release 15.1F4 for PTX Series routers with third-generation FPCs installed. Supported at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level.
Description	Specify the version 9 or IPFIX template properties for one of the following: <ul style="list-style-type: none"> • 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 <i>template-name</i>] hierarchy level.
	<div>  <p>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 <i>template-name</i>] hierarchy level.</p> </div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates • Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers • Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250 • Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers

ipv6-flow-table-size

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: 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 interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)

ipv6-extended-attrib

Syntax	ipv6-extended-attrib;
Hierarchy Level	[edit chassis fpc <i>slot-number</i> 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
	<div> NOTE: Collection of IPv4 fragmentation IDs occurs automatically without having to configure this setting explicitly.</div>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115

ipv6-template

Syntax	ipv6-template;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 9.4. Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2. Statement introduced in Junos OS Release 15.1F4 for PTX Series routers with third-generation FPCs installed. Supported at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level.
Description	Specify that the flow aggregation version 9 or IPFIX template is used only for IPv6 records.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates• Configuring Flow Aggregation to Use IPFIX Flow Templates on PTX Series Routers

jflow-log (Interfaces)

Syntax	<pre>jflow-log { message-rate-limit <i>messages-per-second</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>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).</p> <p>The remaining statement is described separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i> • <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i> • <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i> • <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

jflow-log (Services)

Syntax	<pre>jflow-log { collector <i>collector-name</i> { source-ip <i>address</i>; destination-address <i>address</i>; destination-port <i>port-number</i>; } collector-group <i>collector-group-name</i> { [<i>collector-name1 collector-name2</i>]; } template-profile <i>template-profile-name</i> { collector <i>collector-name</i> ; collector-group <i>collector-group-name</i> ; template-type nat; version (ipfix v9); refresh-rate <i>packets packets seconds seconds</i>; message-rate-limit <i>messages-per-second</i></pre>
Hierarchy Level	[edit services]
Release Information	Statement introduced in Junos OS Release 14.2R2.
Description	<p>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.</p> <p>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.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

label-position

Syntax	label-position [<i>positions</i>];
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i> mpls-ipv4-template], [edit services flow-monitoring version9 template <i>template-name</i> mpls-template]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Specify positions for up to three labels in the active flow monitoring version 9 template.
Default	[1 2 3]
Options	<i>positions</i> —Numbered positions for the labels.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i>

license-server

Syntax	<pre>license-server { ip-address <i>address</i>; log-interval <i>seconds</i>; services (jflow cgnat firewall); }</pre>
Hierarchy Level	[edit]
Release Information	Statement introduced in Junos OS Release 15.1 for MX Series routers.
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	<p>ip-address <i>address</i>—Use the specified IP address of the license log server.</p> <p>log-interval <i>seconds</i>—Use the specified frequency at which throughput data must be sent from the router to the log collector. Range: 60 through 86,400 seconds</p> <p>services—Specify the services for which throughput data must be exported.</p> <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.
Related Documentation	<ul style="list-style-type: none">• <i>License Server Management for Throughput Data Export on MX Series Routers for NAT, Firewall, and Inline Flow Monitoring Services</i>

local-dump

Syntax	(local-dump no-local-dump);
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server <i>hostname</i>], [edit forwarding-options sampling family (inet inet6 mpls) output flow-server <i>hostname</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Flow Aggregation on T and M Series Routers</i>

logical-system

Syntax	logical-system <i>logical-system-name</i> { [routing-instances <i>instance-name</i>]; }
Hierarchy Level	[edit services rpm bgp]
Release Information	Statement introduced in Junos OS Release 7.6.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i>

match

Syntax	<code>match expression;</code>
Hierarchy Level	[edit forwarding-options port-mirroring traceoptions file], [edit forwarding-options sampling traceoptions file]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i>• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>



max-connection-duration

Syntax	<code>max-connection-duration hours;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 11.1.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

max-duplicates

Syntax	<code>max-duplicates <i>number</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
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	<i>number</i> —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.
Related Documentation	<ul style="list-style-type: none">• g-max-duplicates on page 282• <i>Limiting the Number of Duplicates of a Packet</i>


max-packets-per-second

Syntax	<code>max-packets-per-second <i>number</i>;</code>
Hierarchy Level	[edit forwarding-options sampling input], [edit forwarding-options sampling instance <i>instance-name</i> input]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
	<div>NOTE: The <code>max-packets-per-second</code> statement is not supported when you configure inline flow monitoring (by including the <code>inline-jflow</code> statement at the [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6) output] hierarchy level).</div>
	<div>NOTE: When you configure active monitoring and specify a Monitoring Services, Adaptive Services, or Multiservices PIC in the output statement, the <code>max-packets-per-second</code> value is ignored.</div>
Options	<i>number</i> —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.
Related Documentation	<ul style="list-style-type: none"><i>Configuring Traffic Sampling on MX, M and T Series Routers</i>


maximum-age

Syntax	maximum-age <i>minutes</i> ;
Hierarchy Level	[edit services flow-collector transfer-log-archive]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Maximum age of transfer log file.
Options	<i>minutes</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Transfer Logs</i>

maximum-connections

Syntax	<code>maximum-connections count;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Configure the maximum number of allowed connections between the server and all control client hosts.
	<div> 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.</div>
Options	<p>count—Maximum number of connections.</p> <p>Range: 1 through 1000</p> <p>Default: 64</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

maximum-connections-per-client

Syntax	<code>maximum-connections-per-client count;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Configure the maximum number of allowed connections between the server and a single control client host.
	<div>  <p>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.</p> </div>
Options	<p>count—Maximum number of connections.</p> <p>Range: 1 through 500</p> <p>Default: 64</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

maximum-packet-length

Syntax	<code>maximum-packet-length <i>bytes</i>;</code>
Hierarchy Level	[edit forwarding-options analyzer analyzer-name input], [edit forwarding-options port-mirroring input], [edit forwarding-options port-mirroring instance <i>instance-name</i> input], [edit forwarding-options sampling input], [edit forwarding-options sampling instance <i>instance-name</i> input]
Release Information	Statement introduced in Junos OS Release 9.6. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport routers. Support at the [edit forwarding-options analyzer analyzer-name input] hierarchy level introduced in Junos OS Release 14.1 for MX Series routers. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.
Description	Set the maximum length of the packet used for port mirroring or traffic sampling. Packets with lengths greater than the specified maximum are truncated.



NOTE: The `maximum-packet-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).



NOTE: The `maximum-packet-length` statement is not supported on MX80 Series routers.



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.

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: For PTX Series routers with third-generation FPCs installed, the `maximum-packet-length` statement at the `[edit forwarding-options sampling input]` and `[edit forwarding-options sampling instance instance-name input]` hierarchy levels is not supported.

Options `bytes`—Maximum length (in bytes) of the mirrored packet or the sampled packet.



BEST PRACTICE: Juniper Networks recommends that you configure the packet length equal to or greater than the IP header. 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

Default: 0

For MX Series routers with Modular Port Concentrators (MPCs) and EX9200 switches, port-mirrored or sampled packets can be truncated (or clipped) to any length in the range of 1 through 255 bytes.


For other devices, the range is from 0 through 9216. A maximum-packet-length value of zero represents that truncation is disabled, and the entire packet is mirrored or sampled.

Required Privilege Level `interface`—To view this statement in the configuration.
`interface-control`—To add this statement to the configuration.

Related Documentation

- *Configuring Port Mirroring*
- *Configuring Traffic Sampling on MX, M and T Series Routers*

maximum-sessions

Syntax	<code>maximum-sessions <i>count</i>;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	Configure the maximum number of allowed test sessions the server can have running at one time.
	<div> 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.</div>
Options	<i>count</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

maximum-sessions-per-connection

Syntax	<code>maximum-sessions-per-connection count;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 9.3.
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.
Related Documentation	<ul style="list-style-type: none"> Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches

media-loss-rate

Syntax `media-loss-rate {
 no-syslog-generation;
 generate-snmp-traps;
 storm-control {
 count number;
 interval number;
 }
 alarm-mode {
 immediate;
 }
 }`

Hierarchy Level [edit services]

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- *Understanding Inline Video Monitoring on MX Series Routers*
- [alarms on page 205](#)

media-rate-variation

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]

Release Information Statement introduced in Junos OS Release 15.1.

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.

Related Documentation

- *Understanding Inline Video Monitoring on MX Series Routers*
- [alarms on page 205](#)

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

Syntax	<code>message-rate-limit <i>messages-per-second</i></code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options jflow-log]
Release Information	Statement introduced in Junos OS Release 15.1.
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).
	<div> NOTE: The <code>message-rate-limit</code> option can be configured only for multiservices interfaces (<code>ms-x/x/x</code>) and not with other interface types.</div>
Options	<i>messages-per-second</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

minimum-priority

Syntax	minimum-priority <i>value</i> ;
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Control Source</i>

mode (RFC 2544 Benchmarking)

Syntax	mode reflect;
Hierarchy Level	[edit services rpm rfc2544-benchmarkingtests <i>test-name</i> <i>test-name</i>]
Release Information	Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.
Description	Specify the test mode for the packets that are sent during the benchmarking test.
Options	reflect —Reflect the test frames on the chosen service (IPv4 or Ethernet).
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • rfc2544-benchmarking on page 373

monitoring

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]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation • [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

moving-average-size

Syntax	<code>moving-average-size <i>number</i>;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test <i>test-name</i>], [edit services rpm twamp client control-connection <i>control-client-name</i>]
Release Information	Statement introduced in Junos OS Release 8.5. 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 at the [edit services rpm twamp client control-connection <i>control-client-name</i>] 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.
Description	Enable statistical calculation operations to be performed across a configurable number of the most recent samples.
Options	<i>number</i> —Number of samples to be used in calculations. Range: 0 through 255
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>


mpls-flow-table-size

Syntax	<code>mpls-flow-table-size <i>units</i>;</code>
Hierarchy Level	[edit chassis fpc <i>slot-number</i> inline-services flow-table-size]
Release Information	Statement introduced in Junos OS Release 16.1 for MX Series routers.
Description	Configure the size of the MPLS flow table in units of 256,000 entries. <div><div></div><div><p>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.</p></div></div>
Options	<i>units</i> —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.
Related Documentation	<ul style="list-style-type: none">• Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115


mpls-ipv4-template

Syntax	<code>mpls-ipv4-template { label-position [<i>positions</i>]; }</code>
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 8.3. Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level introduced in Junos OS Release 16.1.
Description	Specify the flow aggregation version 9 or IPFIX properties for templates that combine IPv4 and MPLS records. The remaining statement is explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i> • <i>Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers</i> • <i>Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers</i>

mpls-ipvx-template

Syntax	<code>mpls-ipvx-template;</code>
Hierarchy Level	<code>[edit services flow-monitoring version9 template <i>template-name</i>],</code> <code>[edit services flow-monitoring version-ipfix template <i>template-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 18.1R1 on PTX Series routers.
Description	<p>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 <code>[edit services flow-monitoring (version 9 version-ipfix) template <i>template-name</i>]</code> hierarchy level.</p> <p>Starting in Junos OS Release 18.2R1, use mpls-template instead of <code>mpls-ipvx-template</code>.</p> <div>NOTE: For an MPLS-over-UDP flow that is carried between IPv4 endpoints, configure <code>ipv4-template</code> at the <code>[edit services flow-monitoring (version9 version-ipfix) template <i>template-name</i>]</code> hierarchy level.</div>
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers</i>

mpls-template

Syntax	<pre>mpls-template { label-position [<i>positions</i>]; }</pre>
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>]
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level introduced in Junos OS Release 16.1.</p> <p>Statement introduced in Junos OS Release 18.2R1 on PTX Series routers.</p>
Description	<p>Specify the flow aggregation IPFIX or version 9 properties for templates used only for MPLS records.</p> <p>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 <i>template-name</i>] hierarchy level.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> 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 <i>template-name</i>] hierarchy level.</p> </div> <p>The remaining statement is explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i>

multiservice-options

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]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

name-format

Syntax	<code>name-format "format";</code>
Hierarchy Level	<code>[edit services flow-collector file-specification variant <i>variant-number</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
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	<p>format—Specify the filename format, within quotation marks. The name format can include the following macros, typed in braces:</p> <ul style="list-style-type: none"> • {%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_abbrev}**—From **00** through **99**

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • *Configuring File Formats*

next-hop (Forwarding Options)

Syntax next-hop *address*;

Hierarchy Level [edit forwarding-options **port-mirroring family** (inet | inet6) **output interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation • *Configuring Port Mirroring on M, T MX, and PTX Series Routers*

next-hop-group (Forwarding Options)

Syntax `next-hop-group group-name {
 interface interface-name {
 next-hop address;
 }
 }`

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

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 interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation • *Configuring Port Mirroring on M, T MX, and PTX Series Routers*

next-hop-group (Port Mirroring)

Syntax	<code>next-hop-group <i>group-name</i>;</code>
Hierarchy Level	[edit forwarding-options port-mirroring family (inet vpls) output], [edit forwarding-options port-mirroring instance <i>instance-name</i> family (inet vpls) output]
Release Information	Statement introduced in Junos OS Release 9.6.
Description	<p>Specify the next-hop address for sending copies of packets to an analyzer. This configuration enables multipacket port mirroring on MX Series routers and EX Series switches without the use of a Tunnel PIC.</p> <p>The commit operation fails when a next-hop group has only one interface configured. It is implicitly assumed that a subgroup is up only if more than one interface in the subgroup is up.</p>
Options	<i>group-name</i> —Name of next-hop group.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Port Mirroring with Next-Hop Groups</i>

nexthop-learning

Syntax	<code>nexthop-learning (disable enable);</code>
Hierarchy Level	[edit <code>services flow-monitoring(version-ipfix version9) template template-name</code>]
Release Information	<p>Statement introduced in Junos OS Release 15.1F2.</p> <p>Statement introduced in Junos OS Release 17.4R1 for PTX Series routers.</p> <p>nexthop-learning is supported for bridge-template in Junos OS Release 18.2R1 for MX Series routers.</p>
Description	<p>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 flow monitoring. When learning next-hop addresses is disabled, data is reported as follows:</p> <ul style="list-style-type: none"> • If the destination address of the sampled IPv4 flow is reachable through multiple paths, the <code>ipNextHopIPv4Address</code> (Element ID 15) and <code>egressInterface</code> (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 <code>ipNextHopIPv6Address</code> (Element ID 62) and <code>egressInterface</code> (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, <code>destinationIPv4PrefixLength</code> (Element ID 13), <code>bgpDestinationAsNumber</code> (Element ID 17), <code>ipNextHopIPv4Address</code> (Element ID 15), and <code>egressInterface</code> (Element ID 14) are set to 0 in IPv4 flow records and <code>destinationIPv6PrefixLength</code> (Element ID 30), <code>bgpDestinationAsNumber</code> (Element ID 17), <code>ipNextHopIPv6Address</code> (Element ID 62), and <code>egressInterface</code> (Element ID 14) are set to 0 in IPv6 flow records. <p>When learning of next-hop addresses is enabled, output SNMP, destination IP address, and destination IP mask values in the flow records are reported correctly. In addition, when enabled, <code>mplsTopLabelIPv4Address</code> (Element ID 47) in IPv4 flow records reports correctly when MPLS ingress sampling is enabled.</p> <p>To enable next-hop learning, include the nexthop-learning enable statement at the [edit <code>services flow-monitoring version-ipfix template template-name</code>] hierarchy level.</p>
Options	<p>disable—Disable the learning of next hop information required for inline jflow.</p> <p>enable—Enable the learning of next hop information required for inline jflow.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

- Related Documentation**
- *Configuring Next-Hop Address Learning on MX Series Routers for Destinations Accessible Over Multiple Paths*

no-filter-check

Syntax	no-filter-check;
Hierarchy Level	[edit forwarding-options port-mirroring family (inet inet6) output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Disable filter checking on the port-mirroring interface.</p> <p>This statement is required when you send port-mirrored traffic to a Tunnel PIC that has a filter applied to it.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i>


no-remote-trace (Trace Options)

Syntax	no-remote-trace;
Hierarchy Level	[edit forwarding-options port-mirroring traceoptions], [edit forwarding-options sampling traceoptions]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Disable remote tracing.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Tracing Traffic Sampling Operations</i>

no-syslog

Syntax	no-syslog;
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none"> • Configuring System Logging

no-syslog-generation

Syntax	no-syslog-generation;
Hierarchy Level	[edit services]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Disable system log generation.
<div style="display: flex; align-items: center;">  <div> <p>NOTE: If this statement is not configured, <code>edit services</code> generates a system log with respective severity level for values not within the configured range.</p> </div> </div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Understanding Inline Video Monitoring on MX Series Routers • alarms on page 205

notification-targets

Syntax	<code>notification-targets address port port-number;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Control Source</i>

observation-domain-id

Syntax	<code>observation-domain-id <i>domain-id</i>;</code>
Hierarchy Level	[edit services flow-monitoring version-ipfix template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.
Description	<p>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 inique 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.</p> <p>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.</p>
Options	<p><i>domain-id</i>—Unique identifier for the observation domain for IPFIX flows.</p> <p>Range: 0 through 255</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows</i> • <i>Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows</i>

one-way-hardware-timestamp

Syntax	one-way-hardware-timestamp;
Hierarchy Level	[edit services rpm probe owner test test-name]
Release Information	Statement introduced in Junos OS Release 8.5. Statement introduced in Junos OS Release 9.3 for EX Series switches.
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.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Timestamping on MX, M and T Series Routers and EX Series Switches</i>• destination-interface on page 235• hardware-timestamp on page 284

option-refresh-rate

Syntax	<code>option-refresh-rate packets <i>packets</i> seconds <i>seconds</i>;</code>
Hierarchy Level	[<code>edit services flow-monitoring version9 template <i>template-name</i></code>], [<code>edit services flow-monitoring version-ipfix template <i>template-name</i></code>],
Release Information	Statement introduced in Junos OS Release 8.3. Support at the [<code>edit services flow-monitoring version-ipfix template <i>template-name</i></code>] hierarchy level added in Junos OS Release 10.2. Statement introduced in Junos OS Release 15.1F4 for PTX Series routers with third-generation FPCs installed. Supported at the [<code>edit services flow-monitoring version-ipfix template <i>template-name</i></code>] hierarchy level. Support at the [<code>edit services flow-monitoring version9 template <i>template-name</i></code>] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.
Description	Specify the refresh rate, in either packets or seconds.
Options	<i>packets</i> —Refresh rate, in number of packets. Range: 1 through 480,000 Default: 4800 <i>seconds</i> —Refresh rate, in number of seconds. 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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i> • <i>Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250</i>

options-template-id

Syntax	<code>options-template-id id;</code>
Hierarchy Level	[edit services flow-monitoring version9 template template-name], [edit services flow-monitoring version-ipfix template template-name]
Release Information	Statement introduced in Junos OS Release 14.1. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.
Description	Define a unique options template ID to be used for flow aggregation of version 9 and IPFIX flows. 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.
Options	id —Unique identifier for the options template to be used for version 9 or IPFIX flows. Range: 1024 through 65,535
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows</i>• <i>Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows</i>

output (Accounting)

```
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](#)]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- *Configuring Discard Accounting on M and T Series Routers*

output (Monitoring)

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]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege interface—To view this statement in the configuration.
 Level interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

output (Port Mirroring)

Syntax	<pre>output { interface <i>interface-name</i> { next-hop <i>address</i>; } no-filter-check; }</pre>
Hierarchy Level	[edit forwarding-options port-mirroring family (inet inet6)]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure output interfaces and flow properties.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i>

output (Sampling)

```
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)]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.
Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.

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.

Related Documentation

- [Configuring Traffic Sampling on MX, M and T Series Routers](#)

output-interface-index

Syntax output-interface-index *number*;

Hierarchy Level [edit forwarding-options [monitoring name](#) [output interface](#) *interface-name*]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

packet-size (RFC 2544 Benchmarking)

Syntax	<code>packet-size bytes;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking profiles test-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Description	Define up to 10 packet sizes that are used sequentially for the test.
Options	<p>bytes—Size of the test packet. If you enter multiple packet sizes, you must separate each number with a space.</p> <p>Range: 64 through 9136 bytes</p> <div><div></div><div><p>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 <code>test services rpm rfc2544-benchmarking test test-name start</code> command, an error message is displayed if you configured an invalid packet size in the test profile associated with the test name.</p></div></div>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>RFC 2544-Based Benchmarking Tests Overview</i>• <i>Configuring RFC 2544-Based Benchmarking Tests</i>• rfc2544-benchmarking on page 373

passive-monitor-mode

Syntax	<code>passive-monitor-mode;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers</i> • multiservice-options on page 330

password (Flow Collector File Servers)

Syntax	<code>password "password";</code>
Hierarchy Level	[edit services flow-collector destination ftp: <i>url</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the primary and secondary destination FTP server password.
Options	<i>password</i> —FTP server password.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Destination FTP Servers for Flow Records</i>

password (Transfer Log File Servers)

Syntax	<code>password "password";</code>
Hierarchy Level	[edit services flow-collector transfer-log-archive archive-sites]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the primary and secondary destination FTP server password.
Options	<i>password</i> —FTP server password.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Transfer Logs</i>

peer-as-billing-template

Syntax	<code>peer-as-billing-template;</code>
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i>

pic-memory-threshold

Syntax	<code>pic-memory-threshold percentage <i>percentage</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture <code>capture-group</code> <i>client-name</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Specify a PIC memory usage percentage that triggers a system log warning message.
Options	<i>percentage</i> —PIC memory threshold value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Thresholds</i>

pop-all-labels

Syntax	<pre>pop-all-labels { required-depth number; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options mpls], [edit interfaces <i>interface-name</i> fastether-options mpls], [edit interfaces <i>interface-name</i> gigheter-options mpls], [edit interfaces <i>interface-name</i> sonet-options mpls]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>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.</p> <p>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.</p> <p>The remaining statement is explained separately. See CLI Explorer.</p>
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.
Related Documentation	<ul style="list-style-type: none">• <i>Passive Flow Monitoring for MPLS Encapsulated Packets</i>• <i>Junos OS Network Interfaces Library for Routing Devices</i>

port (Flow Monitoring)

Syntax	<code>port port-number;</code>
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]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the User Datagram Protocol (UDP) port number on the cflowd host system or flow server.
Options	<i>port-number</i> —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.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Flow Aggregation on T and M Series Routers</i>

port (RPM)

Syntax	<code>port number;</code>
Hierarchy Level	[edit services rpm probe-server (tcp udp)]
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 Release 18.1 for QFX Series switches.
Description	Specify the port number for the probe server.
Options	<i>number</i> —Port number for the probe server. The value can be 7 or 49,160 through 65,535.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RPM Receiver Servers on M, MX, T and PTX Series Routers and EX Series Switches</i>

port (TWAMP)

Syntax	<code>port <i>number</i>;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 9.3.
Description	TWAMP server listening port.
Options	<i>number</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

port-mirroring

```
Syntax  port-mirroring {
        input {
            maximum-packet-length bytes
            rate rate;
            run-length number;
        }
        family any {
            output {
                (next-hop-group group-name | interface interface-name);
            }
        }
        family inet {
            output {
                interface interface-name {
                    next-hop address;
                }
                no-filter-check;
            }
        }
        instance instance-name {
            input {
                rate rate;
                maximum-packet-length number;
            }
            family any {
                output {
                    (next-hop-group group-name | interface interface-name);
                }
            }
            family inet {
                output {
                    next-hop-group group-name;
                }
            }
        }
        traceoptions {
            file filename <files number> <size bytes> <world-readable | no-world-readable>;
        }
    }
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the input, output, and traceoptions properties for sending copies of packets to an analyzer.



NOTE: Option `run-length` is not supported on MX Series routers with MPCs.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege	interface—To view this statement in the configuration.
Level	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i>

post-cli-implicit-firewall

Syntax	post-cli-implicit-firewall;
Hierarchy Level	[edit services rpm twamp]
Release Information	Statement introduced in Junos OS Release 15.1.
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.

On issuing the command `restart remote-operations` all TWAMP sessions (both client and server) are aborted. 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	interface—To view this statement in the configuration.
Level	interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Two-Way Active Measurement Protocol on Routers</i>

pre-rewrite-tos

Syntax	pre-rewrite-tos;
Hierarchy Level	[edit forwarding-options sampling]
Release Information	Statement introduced in Junos OS Release 14.1.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

probe

```
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;
            hardware-timestamp;
            history-size size;
            inet6-options source-address ipv6-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 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]
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 18.1 for QFX Series switches.
Description	Specify an owner name. The owner name combined with the test name represent a single RPM configuration instance.
Options	<p>owner—Owner name up to 32 characters in length.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

probe-count

Syntax	<code>probe-count count;</code>
Hierarchy Level	<code>[edit services rpm bgp],</code> <code>[edit services rpm probe owner test test-name],</code> <code>[edit services rpm twamp client control-connection control-client-name test-session session-name]</code>
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 <code>[edit services rpm twamp client control-connection control-client-name]</code> 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.
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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i><i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>

probe-interval

Syntax	<code>probe-interval <i>interval</i>;</code>
Hierarchy Level	<code>[edit services rpm bgp],</code> <code>[edit <i>services</i> rpm <i>probe</i> owner <i>test</i> <i>test-name</i>],</code> <code>[edit services rpm twamp client control-connection <i>control-client-name</i> test-session <i>session-name</i>]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.</p> <p>Support at the <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 18.1 for QFX Series switches.</p>
Description	Specify the time to wait between sending packets, in seconds.
Options	<i>interval</i> —Number of seconds, from 1 through 255.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i> • <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i> • <i>Understanding Two-Way Active Measurement Protocol on Routers</i>

probe-limit

Syntax	<code>probe-limit <i>limit</i>;</code>
Hierarchy Level	[edit services rpm]
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 Release 18.1 for QFX Series switches.
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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"><i>Limiting the Number of Concurrent RPM Probes on M, MX, T and PTX Routers and EX Series Switches</i>

probe-server

```
Syntax  probe-server {
        tcp {
            destination-interface interface-name;
            port number;
        }
        udp {
            destination-interface interface-name;
            port number;
        }
    }
```

Hierarchy Level [edit [services](#) rpm]

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 Release 18.1 for QFX Series switches.

Description Specify the server to act as a receiver for the probes.

The remaining statements are explained separately. See [CLI Explorer](#).



NOTE: The `destination-interface` statement is not supported on PTX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring RPM Receiver Servers on M, MX, T and PTX Series Routers and EX Series Switches*

probe-type

Syntax	<code>probe-type type;</code>
Hierarchy Level	<code>[edit services rpm bgp],</code> <code>[edit services rpm probe owner test test-name]</code>
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 Release 18.1 for QFX Series switches.
Description	Specify the packet and protocol contents of a real-time performance monitoring (RPM) probe.
Options	<p>type—One of the following probe type values:</p> <ul style="list-style-type: none">• http-get—(Not available at the <code>[edit services rpm bgp]</code> hierarchy level.) Sends a Hypertext Transfer Protocol (HTTP) get request to a target URL.• http-metadata-get—(Not available at the <code>[edit services rpm bgp]</code> 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.• icmp6-ping—Sends ICMP IPv6 echo requests to a target address. This option is supported only when you configure the generation of RPM probes on an MS-MPC or MS-MIC.• 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. <p>If you are configuring the generation of RPM probes on an MS-MPC or MS-MIC, the <i>type</i> 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.</p>
Required Privilege Level	interface —To view this statement in the configuration. interface-control —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i>

profiles (RFC 2544 Benchmarking)

Syntax	<pre> profiles { test-profile profile-name { test-type (throughput latency frame-loss back-back-frames); packet-size bytes; step-percent percent; bandwidth-kbps kpbs; } } </pre>
Hierarchy Level	[edit services rpm rfc2544-benchmarking]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
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	<p>profiles—Define the test profile for the RFC 2544-based benchmarking test to examine and analyze the performance characteristics of a network interconnecting device.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>RFC 2544-Based Benchmarking Tests Overview</i> • <i>Configuring RFC 2544-Based Benchmarking Tests</i> • rfc2544-benchmarking on page 373

rate (Forwarding Options)

Syntax	<code>rate number;</code>
Hierarchy Level	<code>[edit forwarding-options analyzer <i>analyzer-name</i> input]</code> , <code>[edit forwarding-options port-mirroring family (inet inet6) input]</code> , <code>[edit forwarding-options port-mirroring input]</code> , <code>[edit forwarding-options sampling input]</code> , <code>[edit forwarding-options sampling instance <i>instance-name</i> input]</code>
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport routers. Support at the <code>[edit forwarding-options analyzer <i>analyzer-name</i> input]</code> hierarchy level for MX Series routers introduced in Junos OS Release 14.1. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.
Description	<p>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.</p> <p>Native analyzer sessions (that is, the <code>[edit forwarding-options analyzer <i>analyzer-name</i> input]</code> 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.</p>
Options	number —Denominator of the ratio. Range: 1 through 16000000(16M)
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Port Mirroring</i>• <i>Configuring Traffic Sampling</i>

receive-options-packets

Syntax	receive-options-packets;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers</i>

receive-ttl-exceeded

Syntax	receive-ttl-exceeded;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Passive Flow Monitoring on M Series, MX Series or T Series Routers</i>

refresh-rate (Flow Monitoring Logs for NAT)

Syntax	refresh-rate packets <i>packets</i> seconds <i>seconds</i> ;
Hierarchy Level	[edit services jflow-log template-profile <i>template-profile-name</i>]
Release Information	Statement introduced in Junos OS Release 15.1.
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	<p>packets— Number of packets after which templates are sent to the collector. Range: 1 through 480,000 Default: 4800</p> <p>seconds—Number of seconds after which templates are sent to the collector Range: 10 through 600 Default: 600</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

reflect-mode (RFC2544 Benchmarking)

Syntax	reflect-mode (mac-rewrite mac-swap no-mac-swap);
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers.
Description	Specify the reflection mode for the benchmarking test.
Options	<p>mac-rewrite—(ACX Series routers only) Enable rewriting of the MAC address on the reflected frames. The MAC addresses specified in the source-mac-address and destination-mac-address options are used.</p> <p>mac-swap—Swap the source and destination MAC addresses in the test frame. This is the default behavior.</p>



NOTE: In bridge families, when the service type is ELAN, MAC addresses are swapped by default, on the reflected frames. And, 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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding RFC2544-Based Benchmarking Tests for E-LAN and E-Line Services on MX Series Routers</i> • rfc2544-benchmarking on page 373 • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

reflect-etype (RFC 2544 Benchmarking)

Syntax	reflect-etype <i>ethertype-value</i> ;
Hierarchy Level	[edit services rpm rfc2544-benchmarkingtests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 15.1 for MX104 3D Universal Edge routers.
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	<i>ethertype-value</i> —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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>• <i>Supported RFC2544-Based Benchmarking Statements on MX Series Routers</i>• <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

required-depth

Syntax	<code>required-depth <i>number</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> atm-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> fastether-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> gigether-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> sonet-options mpls pop-all-labels]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>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 MX Series routers with MPCs, all labels are popped by default and the required-depth statement is ignored.</p> <p>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.</p>
Options	<p><i>number</i>—Number of MPLS labels on incoming IP packets.</p> <p>Range: 1 through 2 labels.</p> <p>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.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Passive Flow Monitoring for MPLS Encapsulated Packets</i> <i>Junos OS Network Interfaces Library for Routing Devices</i>

retry (Services Flow Collector)

Syntax	<code>retry number;</code>
Hierarchy Level	[edit services flow-collector]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the maximum number of attempts the flow collector interface make to transfer log files to the FTP server.
Options	<i>number</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Retry Attempts</i>

retry-delay

Syntax	<code>retry-delay seconds;</code>
Hierarchy Level	[edit services flow-collector]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the amount of time the flow collector interface waits between retry attempts.
Options	<i>seconds</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Retry Attempts</i>

rfc2544-benchmarking

```
Syntax  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 {
                test-interface interface-name;
                mode reflect;
                family (bridge| inet | ccc);
                destination-ipv4-address address;
                destination-udp-port port-number;
                source-ipv4-address address;
                source-udp-port port-number;
                direction (egress | ingress);
            }
        }
    }
```

Hierarchy Level [edit [services](#) rpm]

Release Information Statement introduced in Junos OS Release 12.3X52 for ACX Series routers.
Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.

Description Configure the parameters for the RFC 2544-based benchmarking test. You must configure a test profile, 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, for the test. You can associate the same test profile with multiple test names.

Define the attributes 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.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.


Related Documentation

- *Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers*
- *Understanding RFC2544-Based Benchmarking Tests on MX Series Routers*
- *show services rpm rfc2544-benchmarking*
- *show services rpm rfc2544-benchmarking test-id*

rfc6514-compliant-safi129 (Protocols BGP)

Syntax	rfc6514-compliant-safi129
Hierarchy Level	<pre>[edit logical-systems <i>logical-system-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit protocols bgp], [edit protocol bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>neighbor-name</i>]</pre>
Release Information	Statement introduced in Junos OS Release 16.1 for MX Series routers.
Description	Parse and send BGP VPN multicast traffic according to Subsequent Address Family Identifier (SAFI) 129, as defined in RFC 6514 (that is, <i>length, prefix</i>). The Network Layer Reachability Information (NLRI) format used for BGP VPN multicast in previous releases of Junos OS was SAFI 128, which was <i>length, label, prefix</i> .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i>

routing-instance

Syntax	<code>routing-instance <i>instance-name</i>;</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code> <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code>
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 <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> 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.
Description	Specify the routing instance used by the probes. The routing instance is also applicable for control connection.
<div>  <p>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.</p> </div>	
Options	<i>instance-name</i> —Routing instance configured at the <code>[edit routing-instance]</code> hierarchy level. Default: Internet (IPv4) routing table <code>inet.0</code> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i> <i>Understanding Two-Way Active Measurement Protocol on Routers</i>

routing-instance (cflowd)

Syntax	<code>routing-instance <i>instance-name</i>;</code>
Hierarchy Level	[edit forwarding-options sampling family (inet inet6 mpls) output flow-server <i>hostname</i>]
Release Information	Statement introduced in Junos OS Release 13.3.
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 <i>instance-name</i>] hierarchy level.
Options	<i>instance-name</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Directing Traffic Sampling Output to a Server Running the cflowd Application</i>


routing-instance-list (TWAMP)

Syntax	<pre> routing-instance-list { instance-name { port number; } } </pre>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	<p>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.</p>
Options	<p>instance-name—Name of the routing instance, a maximum of 31 characters.</p> <p>number—Port number.</p> <p>Range: 1 through 65,535</p>
Required Privilege Level	<p>system—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches

routing-instances

Syntax	<code>routing-instances <i>instance-name</i>;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm bgp logical-system <i>logical-system-name</i>]
Release Information	Statement introduced in Junos OS Release 7.6. 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 Release 18.1 for QFX Series switches.
Description	Specify the routing instance used by the probes.
Options	<i>instance-name</i> —A routing instance configured at the [edit routing-instances] hierarchy level. Default: Internet routing table <code>inet.0</code> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i>

rpm (Interfaces)

Syntax	<code>rpm (client <i>client</i> server <i>server</i> twamp-client <i>twamp-client</i> twamp-server <i>twamp-server</i>);</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</code>
Release Information	<p>Statement introduced in Junos OS Release 8.1.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 15.1 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 18.1 for QFX Series switches.</p>
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.
<div>  NOTE: The TWAMP client is applicable only for si- interfaces. </div>	
Options	<p><i>client</i>—Identifier for RPM client router or switch.</p> <p><i>server</i>—Identifier for RPM server.</p> <p><i>twamp-client</i>—Identifier for RPM TWAMP client router.</p> <p><i>twamp-server</i>—Identifier for RPM TWAMP server.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RPM Timestamping on MX, M and T Series Routers and EX Series Switches</i>

rpm (Services)

```
Syntax 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;
    }
    probe owner {
        test test-name {
            data-fill data;
            data-size size;
            destination-interface interface-name;
            destination-port port;
            dscp-code-point (Services) dscp-bits;
            hardware-timestamp;
            history-size size;
            moving-average-size number;
            one-way-hardware-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;
        }
    }
    probe-server {
        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 {
        test-interface interface-name;
        mode reflect;
        family (bridge| inet | ccc);
        destination-ipv4-address address;
        destination-udp-port port-number;
        source-ipv4-address address;
        source-udp-port port-number;
        direction (egress | ingress);
    }
}
}
traceoptions {
    file filename <files number> <match regular-expression > <size maximum-file-size>
    <world-readable | no-world-readable>;
    flag flag;
}
twamp {
    server {
        authentication-mode (authenticated | encrypted | none);
        authentication-key-chain identifier {
            key-id identifier {
                secret password-string;
            }
        }
        client-list list-name {
            [ address address ];
        }
        inactivity-timeout seconds;
        maximum-connections-duration hours;
        maximum-connections count;
        maximum-connections-per-client count;
        maximum-sessions count;
        maximum-sessions-per-connection count;
        port number;
        server-inactivity-timeout minutes;
    }
}
rfc2544-benchmarking {
    tests{
        test-name test-name {
            test-interface interface-name;
            mode reflect;
            family (inet | ccc);
            destination-ipv4-address address;
            destination-udp-port port-number;
            source-ipv4-address address;
            source-udp-port port-number;
            direction (egress | ingress);
        }
    }
}

```

```
}  
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure BGP neighbor discovery through RPM.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM*

rpm-scale

```
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*]

Release Information Statement introduced in Junos OS Release 17.4R1 on MX Series routers.

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 *interface-name.logical-unit-number*—The services interface that is generating RPM probes and the logical unit number that is used for the first test that is generated.

ipv4-address-base—The IPv4 source or target address that is incremented to generate the addresses used in the RPM tests.

ipv4-count—The maximum number of IPv4 source or target addresses to use for the generated RPM tests.

ipv4-step—The amount to increment the IPv4 source or target address for each generated RPM test.

ipv6-address-base—The IPv6 source or target address that is incremented to generate the addresses used in the RPM tests.

ipv6-count—The maximum number of IPv6 source or target addresses to use for the generated RPM tests.

ipv6-step—The amount to increment the IPv6 source or target address for each generated RPM test.

subunit-cnt—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 *interface-name.logical-unit-number* 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.


tests-count—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.

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.

Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>
------------------------------	--

run-length

Syntax	<code>run-length <i>number</i>;</code>
Hierarchy Level	<p>[edit forwarding-options port-mirroring input], [edit forwarding-options port-mirroring instance <i>port-mirroring-instance-name</i> input], [edit forwarding-options port-mirroring family (inet inet6) input], [edit forwarding-options sampling input], [edit forwarding-options sampling instance instance-name input]</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 12.1X48 for PTX Series Packet Transport routers. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.</p>
Description	Set the number of samples following the initial trigger event. The configuration enables you to sample packets following those already being sampled.
<div>  <p>NOTE: The <code>run-length</code> statement is not supported when you configure inline flow monitoring (by including the <code>inline-jflow</code> statement at the [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6) output] hierarchy level).</p> </div>	
Options	<p><i>number</i>—Number of samples. Range: 0 through 20 Default: 0</p>
Required Privilege Level	<p>interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Applying Forwarding Table Filters</i> • <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i> • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

sample-once

Syntax	sample-once;
Hierarchy Level	[edit forwarding-options sampling]
Release Information	Statement introduced in Junos OS Release 9.6.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

sampling (Forwarding Options)

```
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]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 12.3R2 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 16.1X65 for PTX1000 routers.</p> <p>Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.</p> <p>Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.</p>
Description	<p>Configure traffic sampling.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i> • <i>Applying Forwarding Table Filters</i> • <i>Collecting Traffic Sampling Output in the Cisco Systems NetFlow Services Export Version 9 Format</i> • <i>Directing Traffic Sampling Output to a Server Running the cflowd Application</i> • <i>Configuring Port Mirroring</i> • <i>Tracing Traffic-Sampling Operations</i> • Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115

sampling (Interfaces)

Syntax	<code>sampling <i>direction</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4. Support for sampling on both input and output for bridge family introduced in Junos OS Release 18.2R1. Support for sampling on both input and output for vpls family introduced in Junos OS Release 18.2R1.
Description	Configure the direction of traffic to be sampled.
Options	<i>direction</i> 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 expect egress point.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Services Interfaces Library for Routing Devices</i>• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

server

Syntax server {
 client-list *list-name* {
 [address *address*];
 }
 inactivity-timeout *seconds*;
 maximum-connections *count*;
 maximum-connections-per-client *count*;
 maximum-sessions *count*;
 maximum-sessions-per-connection *count*;
 port *number*;
 }

Hierarchy Level [edit services rpm [twamp](#)]

Release Information Statement introduced in Junos OS Release 9.3.

Description TWAMP server configuration settings.

Options The remaining statements are described separately.

Required Privilege Level system—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation • *Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches*


server-inactivity-timeout

Syntax	<code>server-inactivity-timeout <i>minutes</i>;</code>
Hierarchy Level	[edit services rpm twamp server]
Release Information	Statement introduced in Junos OS Release 11.1.
Description	The maximum time the Two-Way Active Measurement Protocol (TWAMP) server has to finish the TWAMP control protocol negotiation.
Options	<i>minutes</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i>

service-port

Syntax	<code>service-port <i>port-number</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	Identify the User Datagram Protocol (UDP) port number for control protocol requests.
Options	<i>port-number</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Control Source</i>

service-type (RFC2544 Benchmarking)

Syntax	<code>service-type (elan eline) ;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers.
Description	Mention the service under test. Possible values are elan and eline . This statement is applicable only for the bridge 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 is not supported in 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.
	<div>  <p>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.</p> </div>
Options	<p>elan—Specify elan service type.</p> <p>eline—Specify eline service type.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • rfc2544-benchmarking on page 373 • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>

services

Syntax	services { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure router services.</p> <p>The underlying statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i>

services

Syntax	<code>services rpm { ... }</code>
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4. Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.
Description	Define the service rules to be applied to traffic.
Options	rpm —Use the RPM set of rules statements.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i> • <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i> • <i>Configuring RPM Receiver Servers on M, MX, T and PTX Series Routers and EX Series Switches</i> • <i>Limiting the Number of Concurrent RPM Probes on M, MX, T and PTX Routers and EX Series Switches</i> • <i>Configuring RPM Timestamping on MX, M and T Series Routers and EX Series Switches</i> • <i>Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches</i> • <i>Enabling RPM on MX, M and T Series Routers and SRX Firewalls for the Services SDK</i>

services-options

Syntax

```

services-options {
  cgn-pic;
  close-timeout
  fragment-limit
  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
  session-limit {
    maximum number;
    rate 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*]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- *Configuring Default Timeout Settings for Services Interfaces.*
- *Configuring System Logging for Services Interfaces*


shared-key

Syntax	<code>shared-key value;</code>
Hierarchy Level	[edit services dynamic-flow-capture <code>capture-group client-name control-source identifier</code>]
Release Information	Statement introduced in Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Control Source</i>

size

Syntax	<code>size bytes;</code>
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]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>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.</p> <p>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.</p>
Options	<p>bytes—Maximum size of each traffic sampling file or trace log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).</p> <p>Syntax: <i>xk</i> to specify KB, <i>xm</i> to specify MB, or <i>xg</i> to specify GB</p> <p>Range: 10 KB through the maximum file size supported on your router</p> <p>Default: 1 MB for sampling data; 128 KB for log information</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Port Mirroring on M, T MX, and PTX Series Routers</i>• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

slamon-services

Syntax	slamon-services rfc2544;
Hierarchy Level	[edit chassis fpc slot-number]
Release Information	Statement introduced in Junos OS Release 16.1R1.
Description	(MX240, MX480, MX960, MX2010, and MX2020 routers only) Enable support for RFC2544-based benchmarking tests on MX Series routers with MPC1 (MX-MPC1-3D), MPC2 (MX-MPC2-3D), and the 16-port 10-Gigabit Ethernet MPC (MPC-3D-16XGE-SFP) that are hosting test interfaces. For aggregated interfaces, enable support for RFC2544-based benchmarking tests on all MPCs hosting child links. A system log is generated when you enable support for RFC2544-based benchmarking tests on unsupported MPCs.
	<div>  <p>NOTE: On MX104 and MX80 Series routers that have a single fixed FPC, this configuration is not required.</p> </div>
Options	rfc2544—RFC2544-based benchmarking tests.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Enabling Support for RFC2544-Based Benchmarking Tests on MX Series Routers</i>

soft-limit

Syntax	<code>soft-limit <i>bandwidth</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
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	<i>bandwidth</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Content Destination</i>

soft-limit-clear

Syntax	<code>soft-limit-clear <i>bandwidth</i>;</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> content-destination <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Specify a bandwidth threshold at which the latch set by the soft-limit threshold is cleared.
Options	<i>bandwidth</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring the Content Destination</i>• soft-limit on page 400

source-address (Forwarding Options)

Syntax	<code>source-address <i>address</i>;</code>
Hierarchy Level	<p>[edit forwarding-options accounting name output interface interface-name],</p> <p>[edit forwarding-options monitoring name family family inet output interface interface-name],</p> <p>[edit forwarding-options sampling instance instance-name family (inet inet6 mpls vpls) output interface interface-name],</p> <p>[edit forwarding-options sampling family (inet inet6 mpls) output interface interface-name],</p> <p>[edit forwarding-options sampling instance instance-name family inet output inline-jflow]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the source address for monitored packets.
Options	<i>address</i> —Interface source address.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Discard Accounting on M and T Series Routers</i> • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8 • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

source-address (Services)

Syntax	<code>source-address <i>address</i>;</code>
Hierarchy Level	[edit services rpm probe owner test <i>test-name</i>]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.</p> <p>Statement introduced in Junos OS Release 18.1 for QFX Series switches.</p>
Description	<p>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 use the outgoing interface's address as its source.</p> <p>The following addresses cannot be used for the source IP address used 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)
Options	<i>address</i> —Valid IP address.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>

source-addresses

Syntax	<code>source-addresses [<i>addresses</i>];</code>
Hierarchy Level	[edit services dynamic-flow-capture capture-group <i>client-name</i> control-source <i>identifier</i>]
Release Information	Statement introduced in Junos OS Release 7.4.
Description	List the IP addresses from which the control source can send control protocol requests to the Juniper Networks router.
Options	<i>address</i> —Allowed IP source address.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring the Control Source</i>

source-id

Syntax	<code>source-id <i>source-id</i>;</code>
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1.
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.
Options	<i>source-id</i> —Unique identifier for the source for version 9 flows. Range: 0 through 255
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows</i> • <i>Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows</i>

source-ip (Flow Monitoring Logs for NAT)

Syntax	<code>source-ip address;</code>
Hierarchy Level	<code>[edit services jflow-log collector <i>collector-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 15.1.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

source-ipv4-address (RFC 2544 Benchmarking)

Syntax	<code>source-ipv4-address <i>address</i>;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.
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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • rfc2544-benchmarking on page 373

source-mac-address (RFC2544 Benchmarking)

Syntax	<code>source-mac-address <i>mac-address</i>;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests <i>test-name</i> <i>test-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers.
Description	Specify the source MAC address used in generated test frames. This parameter is applicable for a bridge family.
Options	<i>mac-address</i> —Source MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> ; for example, 0000:5e00:5355 or 00:00:5e:00:53:55 .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• rfc2544-benchmarking on page 373• <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>• <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

source-udp-port (RFC 2544 Benchmarking)

Syntax	<code>source-udp-port <i>port-number</i>;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X52 for ACX Series routers. Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.
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	<i>port-number</i> —Source 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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • rfc2544-benchmarking on page 373

stamp

Syntax	<code>(stamp no-stamp);</code>
Hierarchy Level	[edit forwarding-options sampling family (inet inet6 mpls) output file]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Include a timestamp with each line in the output file.
Options	<i>no-stamp</i> —Do not include timestamps. This is the default. <i>stamp</i> —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.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>



storm-control

Syntax	<pre>storm-control { count <i>number</i>; interval <i>number</i>; }</pre>
Hierarchy Level	[edit services]
Release Information	Statement introduced in Junos OS Release 15.1.
Description	Configure the count and the interval to control the flooding of SNMP traps per flow.
Options	<p>count <i>number</i>—Use the specified maximum number of SNMP traps generated in the configured interval.</p> <p>interval <i>number</i>—Use the specified minimum time period, in seconds, between the generation of successive traps.</p> <p>Default: The default count value is 1. The default interval is 1 second.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• <i>Understanding Inline Video Monitoring on MX Series Routers</i>• alarms on page 205

syslog

Syntax	(syslog no-syslog);
Hierarchy Level	[edit interfaces mo-fpc/pic/port multiservice-options]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the <code>/var/log</code> directory.</p> <ul style="list-style-type: none">• syslog—Enable PIC system logging.• no-syslog—Disable PIC system logging.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

target (Services RPM)

Syntax	<code>target (url <i>url</i> address <i>address</i>);</code>
Hierarchy Level	<code>[edit services rpm probe owner test <i>test-name</i>]</code> <code>[edit services rpm twamp client control-connection <i>control-client-name</i> test-session <i>session-name</i>]</code>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Packet Transport routers.</p> <p>Support at the <code>[edit services rpm twamp client control-connection <i>control-client-name</i>]</code> hierarchy level introduced in Junos OS Release 15.1 for MX Series routers.</p> <p>Statement introduced in Junos OS Release 18.1 for QFX Series switches.</p>
Description	Specify the destination address or URL used for the probes.
Options	<p>url <i>url</i>—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.</p>
<p> NOTE: The <i>url</i> is for RPM only.</p>	
<p>address <i>address</i>—For all probe types other than the HTTP probes, use the specified IPv4 or IPv6 address for the target host.</p>	
<p> 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.</p>	
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i> • <i>Understanding Two-Way Active Measurement Protocol on Routers</i> • <i>Configuring the Interface for RPM Timestamping for Client/Server on a Switch (CLI Procedure)</i>

tcp

Syntax	<pre>tcp { destination-interface <i>interface-name</i>; port <i>port</i>; }</pre>
Hierarchy Level	[edit services rpm probe-server]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Statement introduced in Junos OS Release 9.3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 13.2 for PTX Series Packet Transport routers.</p> <p>Statement introduced in Junos OS Release 18.1 for QFX Series switches.</p>
Description	<p>Specify the port information for the TCP server.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring RPM Receiver Servers on M, MX, T and PTX Series Routers and EX Series Switches</i>

template (Flow Monitoring IPFIX Version)

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-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](#)]

Release Information Statement introduced in Junos OS Release 10.2.
Statement introduced in Junos OS Release 12.R3 for EX Series switches.
Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.
Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.

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 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 on page 115](#)
- [Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250](#)

template (Flow Monitoring Version 9)

Syntax	<pre> template <i>template-name</i> { flow-active-timeout <i>seconds</i>; flow-inactive-timeout <i>seconds</i>; flow-key { flow-direction; vlan-id; output-interface; } (bridge-template ipv4-template ipv6-template mpls-template vpls-templatelabel-position [<i>positions</i>] mpls-ipv4-template label-position [<i>positions</i>] mpls-ipvx-template); option-refresh-rate <i>packets packets seconds seconds</i>; options-template-id peer-as-billing-template; source-id template-id template-refresh-rate <i>packets packets seconds seconds</i>; tunnel-observation [ipv4 ipv6 mpls-over-udp]; } </pre>
Hierarchy Level	[edit services flow-monitoring version9]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Specify the version 9 output template properties to support flow monitoring.
Options	<p><i>template-name</i>—Name of the version 9 template.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

template (Forwarding Options)

Syntax	template <i>template-name</i> ;
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server <i>hostname</i> version9], [edit forwarding-options sampling family (inet inet6 mpls) output flow-server <i>hostname</i> version9]
Release Information	Statement introduced in Junos OS Release 8.3.
Description	Specify flow monitoring version 9 template to be used for output of sampling records.
Options	<i>template-name</i> —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.
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates

template (Forwarding Options Version IPFIX)

Syntax	template;
Hierarchy Level	[edit forwarding-options sampling instance family (inet inet6 mpls vpls) output flow-server <i>hostname</i> version-ipfix]
Release Information	Statement introduced in Junos OS Release 10.2. Statement introduced in Junos OS Release 12.R3 for EX Series switches.
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.
Related Documentation	<ul style="list-style-type: none">• Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115• Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250

template-id

Syntax	<code>template-id <i>id</i>;</code>
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1. Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches. Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.
Description	Define a template ID to be used for flow aggregation of version 9 and IPFIX flows. 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.
Options	<i>id</i> —Unique identifier for the template to be used for version 9 or IPFIX flows. Range: 1024 through 65,535
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Observation Domain ID and Source ID for Version 9 and IPFIX Flows</i> <i>Configuring Template ID and Options Template ID for Version 9 and IPFIX Flows</i>

template-profile (Flow Monitoring Logs for NAT)

Syntax	template-profile <i>template-profile-name</i> ;
Hierarchy Level	[edit services jflow-log], [edit services service-set <i>service-set-name</i> jflow-log]
Release Information	Statement introduced in Junos OS Release 15.1.
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 <i>service-set-name</i> jflow-log] hierarchy level.
Options	<i>template-profile-name</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

template-refresh-rate

Syntax	template-refresh-rate packets <i>packets</i> seconds <i>seconds</i> ;
Hierarchy Level	[edit services flow-monitoring version9 template <i>template-name</i>], [edit services flow-monitoring version-ipfix template <i>template-name</i>],
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Support at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level added in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 15.1F4 for PTX Series routers with third-generation FPCs installed. Supported at the [edit services flow-monitoring version-ipfix template <i>template-name</i>] hierarchy level.</p> <p>Support at the [edit services flow-monitoring version9 template <i>template-name</i>] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.</p> <p>Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.</p> <p>Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.</p>
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	<p>packets—Refresh rate, in number of packets.</p> <p>Range: 1 through 480,000</p> <p>Default: 4800</p> <p>seconds—Refresh rate, in number of seconds.</p> <p>Range: 10 through 600</p> <p>Default: 600</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i> • <i>Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250</i>

template-type (Flow Monitoring Logs for NAT)

Syntax	template-type nat;
Hierarchy Level	[edit services jflow-log template-profile <i>template-profile-name</i>]
Release Information	Statement introduced in Junos OS Release 15.1.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

templates

```
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](#)]

Release Information Statement introduced in Junos OS Release 14.1.

Description Configure the media delivery index template containing the measurement parameters for video monitoring.

Options **delay-factor**—Define delay factor syslog threshold levels.

delay-factor-threshold—Delay factor threshold in milliseconds. When the threshold is exceeded, a syslog message is generated.

Default: 0—Do not generate syslogs.

Range: 0 though 65,535 milliseconds

disable—Disable logging for the threshold.

inactive-timeout—Number of seconds of flow inactivity after which time media delivery index statistics collection for a flow is terminated.

Range: 30 through 300 seconds

info | warning | critical—Level of syslog message generated when a threshold is exceeded.

interval-duration—Number of seconds after which time media delivery index flow monitoring statistics for the interval are reported.

Range: 1 through 50

layer3-packets-per-second—Layer 3 packet rate in IP packets per second.

Range: 0 though 4,294,967,295 pps

media-bits-per-second—Media bit rate for the stream in bits per second.

media-loss-rate—Define media loss rate syslog threshold levels.

media-packets-count-in-layer-3—Number of media packets in an IP packet.

Range: 1 through 32

media-packet-size—Size of media packet in bits.

Default: 188

Range: 1 through 2048

media-rate-variation—Define delay factor syslog threshold levels.

mlr-packet-count—Media loss rate threshold expressed as the number of packets dropped. When the threshold is exceeded, a syslog message is generated.

mlr-percentage—Media loss rate threshold expressed as the percentage of total packets dropped. When the threshold is exceeded, a syslog message is generated.

Range: 0 through 100

mrv-variation—Media rate variation threshold. The variation is the ratio of actual media rate to the configured media rate, expressed as a percentage.

template-name—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.
---------------------------------	---

Related Documentation	<ul style="list-style-type: none">• Configuring Inline Video Monitoring on MX Series Routers
------------------------------	--

test

```

Syntax  test test-name {
        data-fill data;
        data-size size;
        destination-interface interface-name;
        destination-port port;
        dscp-code-point (Services) dscp-bits;
        hardware-timestamp;
        history-size size;
        moving-average-size number;
        inet6-options;
        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
    }

```

Hierarchy Level [edit [services](#) rpm [probe](#) owner]

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. inet6-options option added in Junos OS Release 14.1R4 for MX Series routers. Statement introduced in Junos OS Release 18.1 for QFX Series switches.
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 —Test name. The name can be up to 32 characters in length. The remaining statements are explained separately. See CLI Explorer .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i>

tests (RFC 2544 Benchmarking)

Syntax

```
tests {
  test-name test-name {
    test-interface interface-name;
    mode reflect;
    family (inet | ccc);
    destination-ipv4-address address;
    destination-udp-port port-number;
    source-ipv4-address address;
    source-udp-port port-number;
    direction (egress | ingress);
  }
}
```

Hierarchy Level [edit [services](#) rpm [rfc2544-benchmarking](#)]

Release Information Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.

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 interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers*
- *Understanding RFC2544-Based Benchmarking Tests on MX Series Routers*
- [rfc2544-benchmarking on page 373](#)

test-interface (RFC 2544 Benchmarking)

Syntax	<code>test-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name <i>test-name</i>]
Release Information	Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.
Description	<p>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 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.</p>
Options	<i>interface-name</i> —Name of the logical interface on which the test needs to be run.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>• <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i>• rfc2544-benchmarking on page 373

test-interval

Syntax	<code>test-interval seconds;</code>
Hierarchy Level	[edit services rpm bgp], [edit services rpm probe owner test test-name], [edit services rpm twamp client control-connection control-client-name]
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 Release 18.1 for QFX Series switches.
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] and [edit services rpm probe owner test 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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring BGP Neighbor Discovery on MX, M, T and PTX Series Routers Through RPM</i> • <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i> • <i>Understanding Two-Way Active Measurement Protocol on Routers</i>

test-name (RFC 2544 Benchmarking)

Syntax `test-name test-name {
 test-interface interface-name;
 mode reflect;
 family (inet | ccc);
 destination-ipv4-address address;
 destination-udp-port port-number;
 source-ipv4-address address;
 source-udp-port port-number;
 direction (egress | ingress);
 }`

Hierarchy Level [edit [services](#) rpm [rfc2544-benchmarking tests](#)]

Release Information Statement introduced in Junos OS Release 13.3 for MX104 3D Universal Edge routers.

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 interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers*
- *Understanding RFC2544-Based Benchmarking Tests on MX Series Routers*
- [rfc2544-benchmarking on page 373](#)

test-profile (RFC 2544 Benchmarking)

Syntax	<code>test-profile <i>profile-name</i>;</code>
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests <i>test-name</i> <i>test-name</i>] [edit services rpm rfc2544-benchmarking profiles]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
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	<i>profile-name</i> —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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>RFC 2544-Based Benchmarking Tests Overview</i> • <i>Configuring RFC 2544-Based Benchmarking Tests</i> • rfc2544-benchmarking on page 373

test-session

Syntax `test-session session-name {
 data-fill-with zeros data;
 data-size size;
 dscp-code-point (Services) dscp-bits;
 probe-count count;
 probe-interval seconds;
 target (url url | address address);
}`

Hierarchy Level [edit services rpm twamp client control connection *session-name*]

Release Information Statement introduced in Junos OS Release 15.1.

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 interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Understanding Two-Way Active Measurement Protocol on Routers*

test-type (RFC 2544 Benchmarking)

Syntax	test-type (throughput latency frame-loss back-back-frames);
Hierarchy Level	[edit services rpm rfc2544-benchmarking profiles test-profile <i>profile-name</i>]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers.
Description	<p>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.</p> <p>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.</p>
Options	<p>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.</p> <p>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.</p> <p>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.</p> <p>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.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>RFC 2544-Based Benchmarking Tests Overview</i> • <i>Configuring RFC 2544-Based Benchmarking Tests</i> • rfc2544-benchmarking on page 373

thresholds

Syntax	<code>thresholds thresholds;</code>
Hierarchy Level	[edit <code>services rpm probe owner test test-name</code>], [edit <code>services rpm twamp client control-connection control-client-name</code>]
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 Series Transport routers. Support at the [edit <code>services rpm twamp client control-connection control-client-name</code>] 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.
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	<p><i>thresholds</i>—Specify one or more threshold measurements. The following options are supported:</p> <ul style="list-style-type: none"> • 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	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches</i> • <i>Understanding Two-Way Active Measurement Protocol on Routers</i>

traceoptions (Dynamic Flow Capture)

Syntax	<pre>traceoptions { file <i>filename</i> <files <i>number</i>> <size <i>size</i>> <world-readable non-world-readable>; }</pre>
Hierarchy Level	[edit services dynamic-flow-capture]
Release Information	Statement introduced in Junos OS Release 9.2.
Description	Enable and define tracing options for dynamic flow capture events.
Options	<p>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.</p> <p>files <i>number</i>—(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> <p>Range: 2 through 1000 files.</p> <p>Default: 10 files.</p> <p>no-world-readable—(Optional) Restrict access to the file.</p> <p>world-readable—(Optional) Enable free access to the file.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Junos Capture Vision on M and T Series Routers</i>

traceoptions (Forwarding Options)

Syntax	<pre>traceoptions { no-remote-trace; file filename <files number> <size bytes> <match expression> <world-readable no-world-readable>; }</pre>
Hierarchy Level	[edit forwarding-options port-mirroring], [edit forwarding-options sampling]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Configure traffic sampling tracing operations.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Tracing Traffic Sampling Operations</i>

traceoptions (RPM)

Syntax	<pre> traceoptions { file <i>filename</i> <files <i>number</i>> <match <i>regular-expression</i> > <size <i>maximum-file-size</i>> <world-readable no-world-readable>; flag <i>flag</i>; } </pre>
Hierarchy Level	[edit services rpm]
Release Information	Statement introduced in Junos OS Release 13.2.
Description	Define tracing operations for RPM processes.
Options	<p>file <i>filename</i>—Use the specified file to receive the output of the tracing operation. All files are placed in the directory /var/log. Default: rmopd</p> <p>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 size option. Range: 2 through 1000 Default: 3 files</p> <p>match <i>regular-expression</i>—(Optional) Use the specified regular expression to refine the output to include lines that contain the regular expression.</p> <p>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 files option. Range: 10 KB through 1 GB Default: 128 KB</p> <p>world-readable—(Optional) Enable unrestricted file access.</p> <p>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.</p> <p>flag <i>flag</i>—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"> • all—Trace all operations. • configuration—Trace configuration events. • error—Trace events related to catastrophic errors in daemon. • ipc—Trace IPC events.

- **ppm**—Trace ppm events.
- **statistics**—Trace statistics.

Required Privilege Level trace—To view this statement in the configuration.
 trace-control—To add this statement to the configuration.

Related Documentation • *Tracing RPM Operations on MX, M, T and ACX Series Routers*

transfer

Syntax transfer {
 record-level *number*;
 timeout *seconds*;
 }

Hierarchy Level [edit services flow-collector file-specification variant *variant-number*]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation • *Configuring File Formats*

transfer-log-archive

Syntax transfer-log-archive {
 archive-sites {
 ftp:url {
 password "password";
 username username;
 }
 }
 filename-prefix *prefix*;
 maximum-age *minutes*;
 }

Hierarchy Level [edit services flow-collector]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation • *Configuring Transfer Logs*

traps

Syntax	<code>traps traps;</code>
Hierarchy Level	[edit services rpm probe owner test test-name], [edit services rpm twamp client control-connection control-client-name]
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 Release 18.1 for QFX Series switches.
Description	Set the trap bit to generate traps for probes. Traps are sent if the configured threshold is met or exceeded.
Options	traps —Specify one or more traps. The following options are supported: <ul style="list-style-type: none">• control-connection-closed—Generate traps when the control connection is closed.• egress-jitter-exceeded—Generate traps when the jitter in egress time threshold is met or exceeded.• egress-std-dev-exceeded—Generate traps when the egress time standard deviation threshold is met or exceeded.• egress-time-exceeded—Generate traps when the maximum egress time threshold is met or exceeded.• ingress-jitter-exceeded—Generate traps when the jitter in ingress time threshold is met or exceeded.• ingress-std-dev-exceeded—Generate traps when the ingress time standard deviation threshold is met or exceeded.• ingress-time-exceeded—Generate traps when the maximum ingress time threshold is met or exceeded.• jitter-exceeded—Generate traps when the jitter in round-trip time threshold is met or exceeded.• probe-failure—Generate traps when successive probe loss thresholds are crossed.• rtt-exceeded—Generate traps when the maximum round-trip time threshold is met or exceeded.• std-dev-exceeded—Generate traps when the round-trip time standard deviation threshold is met or exceeded.• test-completion—Generate traps when a test is completed.• test-failure—Generate traps when the total probe loss threshold is met or exceeded.

- **test-iteration-done**—Generate traps when all test sessions under control connections complete one test iteration.



NOTE: For RPM traps to be generated, 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.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring RPM Probes on M, MX and T Series Routers and EX Series Switches*
- *Understanding Two-Way Active Measurement Protocol on Routers*

tth

Syntax `tth hops;`

Hierarchy Level [edit services dynamic-flow-capture **capture-group** *client-name* **content-destination** *identifier*]

Release Information Statement introduced in Junos OS Release 7.4.

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.

Related Documentation

- *Configuring the Content Destination*

twamp

Syntax

```
twamp {
  server {
    authentication-mode mode;
    authentication-key-chain identifier {
      key-id identifier {
        secret password-string;
      }
    }
    client-list list-name {
      [ address address ];
    }
    inactivity-timeout seconds;
    max-connection-duration hours;
    maximum-connections count;
    maximum-connections-per-client count;
    maximum-sessions count;
    maximum-sessions-per-connection count;
    port number;
    routing-instance-list {
      instance-name {
        port number;
      }
    }
    server-inactivity-timeout minutes;
  }
}
```

Hierarchy Level [edit services rpm]

Release Information Statement introduced in Junos OS Release 9.3.

Description Configure the Two-Way Active Measurement Protocol (TWAMP) responder or sever settings on all M Series and T Series routers that support Multiservices PICs (running in either Layer 2 or Layer 3 mode), and on MX Series routers.

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 described separately.

Required Privilege Level system—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches*

twamp-server

Syntax twamp-server;

Hierarchy Level [edit interfaces *sp-fpc/pic/port* unit *logical-unit-number*]

Release Information Statement introduced in Junos OS Release 9.3.


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.

Related Documentation

- *Configuring TWAMP on MX, M and T Series Routers and EX4300 Series Switches*

trio-flow-offload

Syntax	trio-flow-offload minimum-bytes <i>minimum-bytes</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Enable any plug-in or daemon on a PIC to generate a request to off-load flows to the Packet Forwarding Engine. This command is available on MX Series routers with Modular Port Concentrators (MPCs) and Modular Interface Cards (MICs).
	<div> NOTE: This feature is not supported for Broadband Edge subscribers (given that service PIC off load is not available with aggregate Ethernet (AE)).</div>
Options	<i>minimum-bytes</i> —Minimum number of bytes that trigger offloading. When this option is omitted, offloading is triggered when both the forward and reverse flows of the session have begun, meaning that at least one packet has flowed in each direction.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Flow Offloading on MX Series Routers</i>

tunnel-observation

Syntax	<code>tunnel-observation [ipv4 ipv6 mpls-over-udp];</code>
Hierarchy Level	<code>[edit services flow-monitoring version9 template <i>template-name</i>]</code> <code>[edit services flow-monitoring version-ipfix template <i>template-name</i>]</code>
Release Information	Statement introduced in Junos OS Release 18.1R1 on PTX Series routers. <code>ipv4</code> and <code>ipv6</code> options added in Junos OS Release 18.2R1.
Description	Specify the types of MPLS flows on which to enable inline flow monitoring.
Options	<p><code>ipv4</code>—Enable flow monitoring for MPLS-IPv4 traffic. You must also configure <code>mpls-template</code> at the <code>[edit services flow-monitoring (version9 version-ipfix) template <i>template-name</i>]</code> hierarchy level.</p> <p><code>ipv6</code>—Enable flow monitoring for MPLS-IPv6 traffic. You must also configure <code>mpls-template</code> at the <code>[edit services flow-monitoring (version9 version-ipfix) template <i>template-name</i>]</code> hierarchy level.</p> <p><code>mpls-over-udp</code>—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 <code>ipv4-template</code> at the <code>[edit services flow-monitoring (version9 version-ipfix) template <i>template-name</i>]</code> hierarchy level. For an MPLS-over-UDP flow that is encapsulated in an RSVP-TE LSP, you must also configure <code>mpls-ipvx-template</code> in Junos OS Release 18.1 or <code>mpls-template</code> starting in Junos OS 18.2R1 at the <code>[edit services flow-monitoring (version 9 version-ipfix) template <i>template-name</i>]</code> hierarchy level.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Inline Active Flow Monitoring on PTX Series Routers</i> <i>Inline Active Flow Monitoring of MPLS-over-UDP Flows on PTX Series Routers</i>

udp

Syntax `udp {
 destination-interface interface-name;
 port port;
 }`

Hierarchy Level [edit `services rpm probe-server`]

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 Release 18.1 for QFX Series switches.

Description Specify the port information for the UDP server.

The remaining statements are explained separately. See [CLI Explorer](#).



NOTE: The `destination-interface` statement is not supported on PTX Series routers.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring RPM Receiver Servers on M, MX, T and PTX Series Routers and EX Series Switches*

udp-tcp-port-swap (RFC 2544 Benchmarking)

Syntax	udp-tcp-port-swap;
Hierarchy Level	[edit services rpm rfc2544-benchmarking tests test-name test-name]
Release Information	Statement introduced in Junos OS Release 12.3X53 for ACX Series routers. Statement introduced in Junos OS Release 14.2 for MX104 3D Universal Edge routers.
Description	Swap source and destination UDP ports in the test packets. Only UDP port swap and UDP over IPv4 traffic is supported.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • rfc2544-benchmarking on page 373 • <i>Understanding RFC2544-Based Benchmarking Tests on MX Series Routers</i> • <i>Configuring RFC 2544-Based Benchmarking Tests on MX Series Routers</i>

unit

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](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options *logical-unit-number*—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.

Related Documentation

- *Junos OS Network Interfaces Library for Routing Devices* for other statements that do not affect services interfaces.
- *Junos OS Network Interfaces Library for Routing Devices*

use-extended-flow-memory

Syntax	use-extended-flow-memory;
Hierarchy Level	[edit chassis fpc slot-number inline-services]
Release Information	Statement introduced in Junos OS Release 16.1 for MX Series routers.
Description	<p>Configure the service to extended flow memory. This service provides more scale in flows for inline services sampling.</p> <p>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 Lookup (LU) platform. It is not applicable for XL line card because XL has dedicated DMEM memory to hold 64M flow entries.</p>
	<div>  <p>NOTE: This configuration is supported only in Lookup (LU) platform.</p> </div>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115 • Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers

username (Services)

Syntax	<code>username <i>user-name</i>;</code>
Hierarchy Level	[edit services flow-collector transfer-log-archive archive-sites]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the username for the transfer log server.
Options	<i>user-name</i> —FTP server username.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Transfer Logs</i>

variant

Syntax	<pre>variant <i>variant-number</i> { data-format <i>format</i>; name-format <i>format</i>; transfer { record-level <i>number</i>; timeout <i>seconds</i>; } }</pre>
Hierarchy Level	[edit services flow-collector file-specification]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring File Formats</i>

version

Syntax	<code>version <i>format</i>;</code>
Hierarchy Level	[edit forwarding-options accounting name output flow-server hostname], [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output flow-server hostname], [edit forwarding-options sampling family (inet inet6 mpls) output flow-server hostname]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the version format of the aggregated flows exported to a cflowd server.
Options	<i>format</i> —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.
Related Documentation	<ul style="list-style-type: none"> • export-format on page 250 • <i>Enabling Flow Aggregation on T and M Series Routers</i>

version (Flow Monitoring Logs for NAT)

Syntax	version (ipfix v9);
Hierarchy Level	[edit services jflow-log template-profile <i>template-profile-name</i>]
Release Information	Statement introduced in Junos OS Release 15.1.
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.
Related Documentation	<ul style="list-style-type: none">• <i>Understanding NAT Event Logging in Flow Monitoring Format on an MX Series Router or NFX250</i>• <i>Configuring Log Generation of NAT Events in Flow Monitoring Record Format on an MX Series Router or NFX250</i>• <i>Monitoring NAT Events on MX Series Routers by Logging NAT Operations in Flow Template Formats</i>• <i>Example: Configuring Logs in Flow Monitoring Format for NAT Events on MX Series Routers for Troubleshooting</i>

version9 (Forwarding Options)

Syntax	<pre>version9 { template template-name; }</pre>
Hierarchy Level	<p>[edit forwarding-options sampling instance <i>instance-name</i> family (bridge inet inet6 mpls vpls) output flow-server <i>hostname</i>],</p> <p>[edit forwarding-options sampling family (inet inet6 mpls vpls bridge) output flow-server <i>hostname</i>]</p>
Release Information	<p>Statement introduced in Junos OS Release 8.3.</p> <p>Statement introduced in Junos OS Release 17.2R1 for PTX Series routers with third-generation FPCs installed.</p> <p>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].</p>
Description	<p>Specify flow monitoring version 9 properties to apply to output sampling records.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates</i>

version9 (Flow Monitoring)

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]`

Release Information Statement introduced in Junos OS Release 8.3.
Statement introduced in Junos OS Release 17.2R1 for PTX Series routers with third-generation FPCs installed.

Description Specify the version 9 output template properties to support flow monitoring.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- *Configuring Flow Aggregation on MX, M, vMX and T Series Routers and NFX250 to Use Version 9 Flow Templates*
- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)

version-ipfix (Forwarding Options)

Syntax	<pre>version-ipfix { template <i>template-name</i>; }</pre>
Hierarchy Level	[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls vpls) output flow-server <i>hostname</i>]
Release Information	<p>Statement introduced in Junos OS Release 10.2.</p> <p>Statement introduced in Junos OS Release 12.R3 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 15.1F4 for PTX Series routers with third-generation FPCs installed.</p> <p>Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.</p> <p>Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.</p>
Description	<p>Specify flow monitoring version IPFIX properties to apply to output sampling records.</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115 • Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250

version-ipfix (Services)

```
Syntax  version-ipfix {
        template template-name {
            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](#)]

Release Information Statement introduced in Junos OS Release 10.2.
 Statement introduced in Junos OS Release 12.R3 for EX Series switches.
 Statement introduced in Junos OS Release 15.1F4 for PTX Series routers with third-generation FPCs installed.
 Statement introduced in Junos OS Release 17.2R1 for QFX10002 switches.
 Statement introduced in Junos OS Release 17.4R1 for QFX10008 and QFX10016 switches.

Description Specify the IPFIX output template properties to support flow monitoring.
 The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)
- [Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250](#)
- [Configuring Inline Active Flow Monitoring on PTX Series Routers](#)

video-monitoring

```
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) mrv-variation;
            }
        }
        media-packets-count-in-layer3 media-packets-count-in-layer3;
        media-packet-size media-packet-size;
    }
}

```

Hierarchy Level	[edit services]
Release Information	Statement introduced in Junos OS Release 14.1.
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.

Related Documentation • [Configuring Inline Video Monitoring on MX Series Routers](#)

vpls-flow-table-size

Syntax vpls-flow-table-size *units*;

Hierarchy Level [edit chassis fpc *slot-number* inline-services flow-table-size]

Release Information Statement introduced in Junos OS Release 13.2.

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.

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.

Related Documentation • [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)

vpls-template

Syntax	vpls-template;
Hierarchy Level	[edit services flow-monitoring version-ipfix version9 template template-name]
Release Information	Statement introduced in Junos OS Release 13 .2. Support at the [edit services flow-monitoring version9 template template-name] hierarchy level introduced in Junos OS Release 18.2R1.
Description	Specify that the IPFIX and version 9 template is used only for VPLS records.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115

world-readable

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]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• Configuring Port Mirroring on M, T MX, and PTX Series Routers• Configuring Traffic Sampling on MX, M and T Series Routers

CHAPTER 10

Basic Flow and Active Flow EX9200 Monitoring Configuration Statements

- [address \(Interfaces\) on page 458](#)
- [cflowd \(Discard Accounting\) on page 459](#)
- [core-dump on page 460](#)
- [destination \(Interfaces\) on page 461](#)
- [engine-id \(Forwarding Options\) on page 462](#)
- [engine-type on page 463](#)
- [export-format on page 464](#)
- [family \(Monitoring\) on page 465](#)
- [filter on page 466](#)
- [flow-active-timeout on page 467](#)
- [flow-export-destination on page 468](#)
- [flow-inactive-timeout on page 469](#)
- [flow-table-size on page 470](#)
- [input-interface-index on page 471](#)
- [interface \(Accounting or Sampling\) on page 471](#)
- [ipv4-flow-table-size on page 472](#)
- [ipv6-flow-table-size on page 473](#)
- [monitoring on page 474](#)
- [multiservice-options on page 475](#)
- [output-interface-index on page 475](#)
- [output \(Monitoring\) on page 476](#)
- [port \(Flow Monitoring\) on page 477](#)
- [sampling \(Interfaces\) on page 478](#)
- [source-address \(Forwarding Options\) on page 479](#)
- [syslog on page 480](#)
- [unit on page 481](#)


address (Interfaces)

Syntax	<code>address address { destination address; }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure the interface address.
Options	<i>address</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Network Interfaces Library for Routing Devices</i> for other options not associated with flow monitoring.• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8• <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

cflowd (Discard Accounting)

Syntax	<pre> cflowd <i>hostname</i> { aggregation { autonomous-system; destination-prefix; protocol-port; source-destination-prefix { caida-compliant; } source-prefix; } autonomous-system-type (origin peer); label-position { template <i>template-name</i>; } (local-dump no-local-dump); port <i>port-number</i>; source-address <i>address</i>; version <i>format</i>; } </pre>
Hierarchy Level	[edit forwarding-options accounting name output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>Collect an aggregate of sampled flows and send the aggregate to a specified host system that runs the collection utility <code>cfcollect</code>.</p> <p>You can configure up to one version 5 and one version 8 flow format at the [edit forwarding-options accounting name output] hierarchy level.</p>
Options	<p>hostname—IP address or identifier of the host system (the workstation running the <code>cflowd</code> utility).</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Enabling Flow Aggregation on T and M Series Routers</i>

core-dump

Syntax	(core-dump no-core-dump);
Hierarchy Level	[edit interfaces mo-fpc/pic/port multiservice-options]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>A useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory <code>/var/tmp</code> 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):</p> <div> 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.</div> <ul style="list-style-type: none">• 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.
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8


destination (Interfaces)

Syntax	<code>destination address;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>For CoS on ATM interfaces, specify the remote address of the connection.</p> <p>For point-to-point interfaces only, specify the address of the interface at the remote end of the connection.</p> <p>For tunnel and encryption interfaces, specify the remote address of the tunnel.</p>
Options	address —Address of the remote side of the connection.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Linear RED Profiles on ATM Interfaces</i> • <i>Multilink and Link Services Logical Interface Configuration Overview</i> • <i>Configuring Encryption Interfaces</i> • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i> • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8 • <i>Configuring Unicast Tunnels</i>

engine-id (Forwarding Options)

Syntax	<code>engine-id <i>number</i>;</code>
Hierarchy Level	<code>[edit forwarding-options accounting name output interface <i>interface-name</i>],</code> <code>[edit forwarding-options monitoring name output interface <i>interface-name</i>],</code> <code>[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output</code> <code>interface <i>interface-name</i>],</code> <code>[edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the engine ID number for flow monitoring and accounting services.
Options	<i>number</i> —Identity of accounting interface.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Traffic Sampling on MX, M and T Series Routers• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8• Configuring Discard Accounting on M and T Series Routers

engine-type

Syntax	<code>engine-type <i>number</i>;</code>
Hierarchy Level	<p>[edit forwarding-options accounting name output interface <i>interface-name</i>],</p> <p>[edit forwarding-options monitoring name output interface <i>interface-name</i>],</p> <p>[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output interface <i>interface-name</i>],</p> <p>[edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>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.</p>
<div>  <p>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.</p> </div>	
Options	<i>number</i> —Platform-specific accounting interface type.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Traffic Sampling on MX, M and T Series Routers • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8 • Configuring Discard Accounting on M and T Series Routers

export-format

Syntax	<code>export-format <i>format</i>;</code>
Hierarchy Level	[edit forwarding-options monitoring name output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Flow monitoring export format.
Options	<i>format</i> —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.
Related Documentation	<ul style="list-style-type: none">• version on page 447• Exporting Flows on page 11

family (Monitoring)

```
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](#)]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

filter

Syntax	<pre>filter { input <i>filter-name</i>; output <i>filter-name</i>; group <i>filter-group-number</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Apply a firewall filter to an interface. You can also use filters for encrypted traffic.
Options	<p>group <i>filter-group-number</i>—Use the specified interface to be part of a filter group. The default filter group number is 0.</p> <p>input <i>filter-name</i>—Use the specified filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Use the specified filter to evaluate when packets are transmitted on the interface.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide</i> or the <i>Junos OS Administration Library</i>• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

flow-active-timeout

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*]

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.

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.

- Related Documentation**
- [Configuring Time Periods When Flow Monitoring Is Active and Inactive on page 12](#)
 - [Configuring the Version 9 Template Properties](#)
 - [Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250](#)

flow-export-destination

Syntax	<pre>flow-export-destination { (cflowd-collector collector-pic); }</pre>
Hierarchy Level	[edit forwarding-options monitoring <i>group-name</i> family inet output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure flow collection.
Options	<p>cflowd-collector—Use the cflowd collector.</p> <p>collector-pic—Use the collector PIC.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Exporting Flows on page 11

flow-inactive-timeout

Syntax	<code>flow-inactive-timeout <i>seconds</i>;</code>
Hierarchy Level	<p>[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],</p>
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>Support at the [edit services flow-monitoring version-ipfix template template-name] hierarchy level added in Junos OS Release 10.2.</p> <p>Support at the [edit services flow-monitoring version9 template template-name] hierarchy level added in Junos OS Release 16.1 for MPLS traffic flows.</p>
Description	Set the interval of inactivity that marks a flow inactive.
<div>  <p>NOTE: The router must include an Adaptive Services, Multiservices, or Monitoring Services PIC for this statement to take effect.</p> </div>	
Options	<p><i>seconds</i>—Duration of the timeout period.</p> <p>Range: 15 through 1800 seconds (for forwarding-options configurations); 10 through 600 seconds (for services configurations)</p> <p>Default: 60 seconds (for forwarding-options configurations); 60 seconds (for services configurations)</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Time Periods When Flow Monitoring Is Active and Inactive on page 12 • Configuring the Version 9 Template Properties • Configuring Flow Aggregation to Use IPFIX Flow Templates on MX, vMX and T Series Routers, EX Series Switches and NFX250

flow-table-size

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;
 bridge-flow-table-size units;
 }`

Hierarchy Level [edit chassis fpc slot-number inline-services]

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.

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
```



NOTE: If you are using a Junos release prior to Junos OS Release 15.1F2, this command initiates an automatic reboot of the FPC, and we recommend you run this command during a maintenance window.

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.

Related Documentation

- [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)
- [Including Fragmentation Identifier and IPv6 Extension Header Elements in IPFIX Templates on MX Series Routers](#)

input-interface-index

Syntax	<code>input-interface-index <i>number</i>;</code>
Hierarchy Level	[edit forwarding-options monitoring name output interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a value for the input interface index that overrides the default supplied by SNMP.
Options	<i>number</i> —Input interface index value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

interface (Accounting or Sampling)

Syntax	<pre>interface <i>interface-name</i> { engine-id <i>number</i>; engine-type <i>number</i>; source-address <i>address</i>; }</pre>
Hierarchy Level	[edit forwarding-options accounting name output], [edit forwarding-options sampling family (inet inet6 mpls) output], [edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the output interface for monitored traffic.
Options	<i>interface-name</i> —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.
Related Documentation	<ul style="list-style-type: none">• Configuring Discard Accounting on M and T Series Routers• Configuring Traffic Sampling on MX, M and T Series Routers

ipv4-flow-table-size

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.

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:



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.

```
[edit]
user@router# set chassis fpc slot-number inline-services flow-table-size
ipv4-flow-table-size 15
```

Options *units*—Number of 256K flow entries available for the IPv4 flow table.

Range: 1 through 245

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
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 • [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)

ipv6-flow-table-size

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: 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 interface—To view this statement in the configuration.

Level interface-control—To add this statement to the configuration.

Related Documentation • [Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250 on page 115](#)

monitoring

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]

Release Information Statement introduced before Junos OS Release 7.4.

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.

Related Documentation

- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

multiservice-options

Syntax	<pre>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; } }</pre>
Hierarchy Level	[edit interfaces <i>mo-fpc/pic/port</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
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.
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

output-interface-index

Syntax	<pre>output-interface-index <i>number</i>;</pre>
Hierarchy Level	[edit forwarding-options monitoring <i>name</i> output interface <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify a value for the output interface index that overrides the default supplied by SNMP.
Options	<i>number</i> —Output interface index value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

output (Monitoring)

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]

Release Information Statement introduced before Junos OS Release 7.4.

Description Configure cflowd, output interfaces, and flow properties.

The remaining statements are explained separately. See [CLI Explorer](#).

Required Privilege interface—To view this statement in the configuration.
Level interface-control—To add this statement to the configuration.

Related Documentation

- [Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8](#)

port (Flow Monitoring)

Syntax	<code>port <i>port-number</i>;</code>
Hierarchy Level	<code>[edit forwarding-options accounting name output cflowd hostname],</code> <code>[edit forwarding-options monitoring name family inet output cflowd hostname],</code> <code>[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls) output</code> <code>flow-server hostname],</code> <code>[edit forwarding-options sampling family (inet inet6 mpls) output flow-server hostname]</code>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the User Datagram Protocol (UDP) port number on the cflowd host system or flow server.
Options	<i>port-number</i> —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.
Related Documentation	<ul style="list-style-type: none">• <i>Enabling Flow Aggregation on T and M Series Routers</i>

sampling (Interfaces)

Syntax	<code>sampling <i>direction</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before Junos OS Release 7.4. Support for sampling on both input and output for bridge family introduced in Junos OS Release 18.2R1. Support for sampling on both input and output for vpls family introduced in Junos OS Release 18.2R1.
Description	Configure the direction of traffic to be sampled.
Options	<i>direction</i> 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 expect egress point.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Junos OS Services Interfaces Library for Routing Devices</i>• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

source-address (Forwarding Options)

Syntax	<code>source-address <i>address</i>;</code>
Hierarchy Level	<p>[edit forwarding-options accounting name output interface <i>interface-name</i>],</p> <p>[edit forwarding-options monitoring name family <i>family</i> inet output interface <i>interface-name</i>],</p> <p>[edit forwarding-options sampling instance <i>instance-name</i> family (inet inet6 mpls vpls) output interface <i>interface-name</i>],</p> <p>[edit forwarding-options sampling family (inet inet6 mpls) output interface <i>interface-name</i>],</p> <p>[edit forwarding-options sampling instance <i>instance-name</i> family inet output inline-jflow]</p>
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the source address for monitored packets.
Options	<i>address</i> —Interface source address.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Discard Accounting on M and T Series Routers</i> • Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8 • <i>Configuring Traffic Sampling on MX, M and T Series Routers</i>

syslog

Syntax	(syslog no-syslog);
Hierarchy Level	[edit interfaces mo- <i>fpc/pic/port</i> multiservice-options]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	<p>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.</p> <ul style="list-style-type: none">• syslog—Enable PIC system logging.• no-syslog—Disable PIC system logging.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring Flow Monitoring on T Series and M Series Routers and EX9200 Switches on page 8

unit

Syntax	<pre> unit <i>logical-unit-number</i> { family inet { address <i>address</i> { destination <i>destination-address</i>; } filter { group <i>filter-group-number</i>; input <i>filter-name</i>; output <i>filter-name</i>; } sampling <i>direction</i>; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<p><i>logical-unit-number</i>—Number of the logical unit.</p> <p>Range: 0 through 16,384</p> <p>The remaining statements are explained separately. See CLI Explorer.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Junos OS Network Interfaces Library for Routing Devices</i> for other statements that do not affect services interfaces. • <i>Junos OS Network Interfaces Library for Routing Devices</i>

CHAPTER 11

Basic Flow and Active Flow EX9200 Monitoring Operational Commands

- `show services accounting aggregation`
- `show services accounting aggregation template`
- `show services accounting errors`
- `show services accounting flow`
- `show services accounting flow-detail`
- `show services accounting memory`
- `show services accounting packet-size-distribution`
- `show services accounting status`
- `show services accounting usage`

show services accounting aggregation

Syntax `show services accounting aggregation aggregation-type <aggregation-value>`
 <detail | extensive | terse>
 <limit *limit-value*>
 < name *service-name*>
 <order (bytes | packets)>

Release Information Command introduced before Junos OS Release 7.4.

Description Display information about the aggregated active flows being processed by the accounting service.

Options *aggregation-type* <*aggregation-value*>—Display information for the specified aggregation type and optional value:

- *as* <*source-as-value* | *destination-as-value* | *input-snmp-interface-index-value* | *output-snmp-interface-index-value*>—Aggregate by autonomous system (AS).
- *destination-prefix* <*destination-prefix-value* | *destination-as-value* | *output-snmp-interface-index-value*>—Aggregate by destination prefix.
- *protocol-port* <*protocol-value* | *source-port-value* | *destination-port-value*>—Aggregate by protocol and port.
- *source-destination-prefix* <*source-prefix-value* | *destination-prefix-value* | *destination-as-value* | *source-as-value* | *input-snmp-interface-index-value* | *output-snmp-interface-index-value*>—Aggregate by source and destination prefix.
- *source-prefix* <*source-prefix-value* | *source-as-value* | *input-snmp-interface-index-value*>—Aggregate by source prefix.

detail | extensive | terse—(Optional) Display the specified level of output.

limit *limit-value*—(Optional) Limit the display output to the specified number of flows. The default is no limit.

name *service-name*—(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

List of Sample Output [show services accounting aggregation protocol-port detail on page 486](#)
[show services accounting aggregation source-destination-prefix on page 486](#)
[show services accounting aggregation source-destination- prefix order packet detail on page 486](#)
[show services accounting aggregation source-destination- prefix extensive limit on page 487](#)
[show services accounting aggregation source-destination-prefix name terse on page 487](#)

Output Fields [Table 34 on page 485](#) lists the output fields for the **show services accounting aggregation** command. Output fields are listed in the approximate order in which they appear.

Table 34: 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 [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.
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.
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.

Table 34: show services accounting aggregation Output Fields (continued)

Field Name	Field Description
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 Prefix	Destination Prefix	Input SNMP Index	Output SNMP Index	SNMP Count	Flow Count	Packet Count	Byte 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	

show services accounting aggregation template

Syntax	show services accounting aggregation template <template-name <i>template-name</i>>
Release Information	Command introduced in Junos OS Release 8.3.
Description	Display information for flow aggregation version 9 templates.
Options	none —Display information for all flow aggregation version version 9 templates. template-name <i>template-name</i> —(Optional) Display information for the specified template only.
Required Privilege Level	view
List of Sample Output	show services accounting aggregation template template-name on page 488
Output Fields	Table 35 on page 488 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 35: 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

show services accounting errors

Syntax	<code>show services accounting errors</code> <code><inline-jflow name (* all <i>service-name</i>)></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display active flow error statistics.
Options	<p>none—Display error statistics for all services accounting instances.</p> <p>inline-jflow fpc-slot <i>slot-number</i>—(Optional) Display error statistics for inline jflow.</p> <p>name (* all <i>service-name</i>)—(Optional) Display active flow error statistics. Use a wildcard character, specify all services, or provide a specific service name.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show services accounting flow on page 494
List of Sample Output	<p>show services accounting errors (Monitoring PIC interface) on page 491</p> <p>show services accounting errors (Service PIC interface) on page 492</p> <p>show services accounting errors inline-jflow fpc-slot (When Only IPv6 Is Configured) on page 492</p> <p>show services accounting errors inline-jflow fpc-slot (When IPv4, IPv6, VPLS, and Bridge Are Configured) on page 492</p> <p>show services accounting errors inline-jflow (MX80 Router When Both IPv4 and IPv6 Are Configured) on page 493</p> <p>show services accounting errors inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured) on page 493</p>
Output Fields	Table 36 on page 490 lists the output fields for the show services accounting errors command. Output fields are listed in the approximate order in which they appear.

Table 36: 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.
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.)

Table 36: show services accounting errors Output Fields (continued)

Field	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.
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 .
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 flow

Syntax	<code>show services accounting flow</code> <code><inline-jflow fpc-slot <i>slot-number</i> logical-system (all <i>logical-system</i>) name (* all <i>service-name</i>)></code>
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.
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>) —(Optional) Display inline flow statistics for the specified FPC. 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
Related Documentation	<ul style="list-style-type: none">• show services accounting status on page 508
List of Sample Output	show services accounting flow (Flow Aggregation v5/v8 Configuration) on page 495 show services accounting flow (Flow Aggregation v9 Configuration) on page 496 show services accounting flow name on page 496 show services accounting flow name all on page 496 show services accounting flow (Multiple Sampling Instances) on page 497 show services accounting flow inline-jflow fpc-slot (for IPv4 Flow) on page 497 show services accounting flow inline-jflow fpc-slot (with IPv4, IPv6, VPLS, and Bridge Configuration) on page 497 show services accounting flow inline-jflow (MX80 Router with IPv4 and IPv6 Configuration) on page 498 show services accounting flow inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured) on page 498
Output Fields	Table 37 on page 495 lists the output fields for the show services accounting flow command. Output fields are listed in the approximate order in which they appear.

Table 37: *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.
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 slot-number 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-detail

Syntax show services accounting flow-detail
<detail | extensive | terse>
<filters>
<limit *limit-value*>
<name (* | all | *service-name*)>
<order (bytes | packets)>

Release Information Command introduced before Junos OS Release 7.4.

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.

filters—(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:

- **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 *limit-value*—(Optional) Limit the display output to the specified number of flows. The default is no limit.

name (* | all | *service-name*)—(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

List of Sample Output [show services accounting flow-detail on page 501](#)
[show services accounting flow-detail limit on page 502](#)
[show services accounting flow-detail name extensive on page 502](#)
[show services accounting flow-detail limit order bytes on page 502](#)
[show services accounting flow-detail name detail source-port on page 503](#)

Output Fields [Table 38 on page 500](#) 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 38: 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
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

Table 38: show services accounting flow-detail Output Fields (continued)

Field Name	Field Description	Output Level
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
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 interface	Source address	Source port	Output 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 address	Destination port	Packet count	Byte count	Time since last 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      Source      Source  Output
         interface address      port    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      Destination      Packet      Byte      Time since last
```


address	port	count	count	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	Byte count for
last active timeout	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

```

show services accounting memory

Syntax	show services accounting memory
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display memory and flow record statistics.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show services accounting memory (Monitoring PIC Interface) on page 505 show services accounting memory (Service PIC Interface) on page 505
Output Fields	Table 39 on page 504 lists the output fields for the show services accounting memory command. Output fields are listed in the approximate order in which they appear.

Table 39: 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).
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
```

show services accounting packet-size-distribution

Syntax	show services accounting packet-size-distribution <name (* all <i>service-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display a packet size distribution histogram.
Options	<p>none—Display a packet size distribution histogram of all accounting services.</p> <p>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.</p>
Required Privilege Level	view
List of Sample Output	show services accounting packet-size-distribution name on page 506
Output Fields	Table 40 on page 506 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 40: 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 [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.
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
```

show services accounting status

Syntax	<pre>show services accounting status <inline-jflow fpc-slot slot-number name (* all service-name)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 13.2R2 for EX Series switches.</p>
Description	Display available Physical Interface Cards (PICs) for accounting services.
Options	<p>none—Display available PICs for all accounting services.</p> <p>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 master 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.</p> <p>name (* all service-name)—(Optional) Display available PICs. Use a wildcard character, specify all services, or provide a specific services name.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show services accounting flow on page 494 Inline Flow Monitoring for Virtual Chassis Overview
List of Sample Output	<p>show services accounting status name (Monitoring PIC Interface) on page 509</p> <p>show services accounting status name (Service PIC Interface) on page 510</p> <p>show services accounting status inline-jflow fpc-slot (When IPv4, IPv6 and Bridge Family Are Configured) on page 510</p> <p>show services accounting status inline-jflow (MX80 Router When Both IPv4 and IPv6) on page 510</p> <p>show services accounting status inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured) on page 510</p>
Output Fields	<p>Table 41 on page 508 lists the output fields for the show services accounting status command. Output fields are listed in the approximate order in which they appear.</p>

Table 41: show services accounting status Output Fields

Field	Field Description
Service Accounting interface	Name of the service accounting interface.

Table 41: show services accounting status Output Fields (continued)

Field	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.
FPC Slot	Slot number of the FPC for which the flow information is displayed.
Local interface index	Index counter of the local interface.
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
```

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 usage

Syntax	<code>show services accounting usage</code> <code><name service-name></code>
Release Information	Command introduced before Junos OS Release 7.4.
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
List of Sample Output	show services accounting usage (Monitoring PIC Interface) on page 512 show services accounting usage (Service PIC Interface) on page 512
Output Fields	Table 42 on page 511 lists the output fields for the show services accounting usage command. Output fields are listed in the approximate order in which they appear.

Table 42: show services accounting usage Output Fields

Output Field	Output 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, (default sampling) , 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.
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%
```

CHAPTER 12

Active Flow Monitoring Commands

- `show forwarding-options next-hop-group`
- `show forwarding-options port-mirroring`
- `show services accounting aggregation`
- `show services accounting aggregation template`
- `show services accounting errors`
- `show services accounting flow`
- `show services accounting flow-detail`
- `show services accounting memory`
- `show services accounting packet-size-distribution`
- `show services accounting status`
- `show services accounting usage`

show forwarding-options next-hop-group

Syntax	show forwarding-options next-hop-group <terse brief detail> <group-name>
Release Information	Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 12.3R2 for EX Series switches. Support for IPv6 introduced in Junos OS Release 14.2 for the MX Series routers.
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
Related Documentation	<ul style="list-style-type: none"> • show forwarding-options port-mirroring on page 517
List of Sample Output	show forwarding-options next-hop-group terse on page 515 show forwarding-options next-hop-group brief on page 515 show forwarding-options next-hop-group detail on page 515
Output Fields	Table 43 on page 514 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 43: 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 43: `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      State
                  next-hop 203.0.113.10     up
                  next-hop 198.51.100.10.10 up
                  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    State
                  next-hop 2001:db8::20:10   up
                  up
Member Subgroup: nhsg6
                  up
  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    State
                    next-hop 2001:db8::4:1    up
                    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
```

show forwarding-options port-mirroring

Syntax	show forwarding-options port-mirroring <terse detail> <instance-name>
Release Information	Command introduced in Junos OS Release 9.6. Command introduced in Junos OS Release 12.3R2 for EX Series switches.
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
Related Documentation	
List of Sample Output	show forwarding-options port-mirroring terse on page 518 show forwarding-options port-mirroring detail on page 518
Output Fields	Table 44 on page 517 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 44: 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
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

Table 44: show forwarding-options port-mirroring Output Fields (continued)

Field Name	Field Description	Level of Output
Destination	Destination (next-hop group name).	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: &global_instance
Instance Id: 1      State: up
  Input parameters:
    Rate:          10
    Run-length:     4
    Maximum-packet-length: 0
  Output parameters:
    Family: inet    State: up Destination: inet_nhg
    Family: vpls/eth-switch State: up Destination: vpls_nhg

Instance Name: inst1
Instance Id: 2      State: up
  Input parameters:
    Rate:          1
    Run-length:     0
    Maximum-packet-length: 200
  Output parameters:
    Family: inet    State: up Destination: inet_nhg
    Family: vpls/eth-switch State: down Destination: vpls_nhg_2
```


show services accounting aggregation

Syntax	<pre>show services accounting aggregation <i>aggregation-type</i> <<i>aggregation-value</i>> <detail extensive terse> <limit <i>limit-value</i>> < name <i>service-name</i>> <order (bytes packets)></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display information about the aggregated active flows being processed by the accounting service.
Options	<p><i>aggregation-type</i> <<i>aggregation-value</i>>—Display information for the specified aggregation type and optional value:</p> <ul style="list-style-type: none"> as <<i>source-as-value</i> <i>destination-as-value</i> <i>input-snmp-interface-index-value</i> <i>output-snmp-interface-index-value</i>>—Aggregate by autonomous system (AS). destination-prefix <<i>destination-prefix-value</i> <i>destination-as-value</i> <i>output-snmp-interface-index-value</i>>—Aggregate by destination prefix. protocol-port <<i>protocol-value</i> <i>source-port-value</i> <i>destination-port-value</i>>—Aggregate by protocol and port. source-destination-prefix <<i>source-prefix-value</i> <i>destination-prefix-value</i> <i>destination-as-value</i> <i>source-as-value</i> <i>input-snmp-interface-index-value</i> <i>output-snmp-interface-index-value</i>>—Aggregate by source and destination prefix. source-prefix <<i>source-prefix-value</i> <i>source-as-value</i> <i>input-snmp-interface-index-value</i>>—Aggregate by source prefix. <p>detail extensive terse—(Optional) Display the specified level of output.</p> <p>limit <i>limit-value</i>—(Optional) Limit the display output to the specified number of flows. The default is no limit.</p> <p>name <i>service-name</i>—(Optional) Display information about the aggregated flows for a specified service name.</p> <p>order (bytes packets)—(Optional) Display the flow with the ordering of the highest number, either by byte count or by packet count.</p>
Additional Information	For information about aggregation configuration options, see the <i>Junos OS Services Interfaces Library for Routing Devices</i> .
Required Privilege Level	view

List of Sample Output [show services accounting aggregation protocol-port detail on page 521](#)
[show services accounting aggregation source-destination-prefix on page 521](#)
[show services accounting aggregation source-destination- prefix order packet detail on page 521](#)
[show services accounting aggregation source-destination- prefix extensive limit on page 522](#)
[show services accounting aggregation source-destination-prefix name terse on page 522](#)

Output Fields [Table 34 on page 485](#) lists the output fields for the **show services accounting aggregation** command. Output fields are listed in the approximate order in which they appear.

Table 45: 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 [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.
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.
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.

Table 45: show services accounting aggregation Output Fields (continued)

Field Name	Field Description
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      Destination      Input      Output      Flow  Packet  Byte
prefix      prefix           interface  interface  count count   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 Prefix	Destination Prefix	Input SNMP Index	Output SNMP Index	SNMP Count	Flow Count	Packet Count	Byte 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	

show services accounting aggregation template

Syntax	show services accounting aggregation template <template-name <i>template-name</i>>
Release Information	Command introduced in Junos OS Release 8.3.
Description	Display information for flow aggregation version 9 templates.
Options	none —Display information for all flow aggregation version version 9 templates. template-name <i>template-name</i> —(Optional) Display information for the specified template only.
Required Privilege Level	view
List of Sample Output	show services accounting aggregation template template-name on page 523
Output Fields	Table 35 on page 488 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 46: 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

show services accounting errors

Syntax	<code>show services accounting errors</code> <code><inline-jflow name (* all <i>service-name</i>)></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display active flow error statistics.
Options	<p>none—Display error statistics for all services accounting instances.</p> <p>inline-jflow fpc-slot <i>slot-number</i>—(Optional) Display error statistics for inline jflow.</p> <p>name (* all <i>service-name</i>)—(Optional) Display active flow error statistics. Use a wildcard character, specify all services, or provide a specific service name.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show services accounting flow on page 494
List of Sample Output	<p>show services accounting errors (Monitoring PIC interface) on page 526</p> <p>show services accounting errors (Service PIC interface) on page 527</p> <p>show services accounting errors inline-jflow fpc-slot (When Only IPv6 Is Configured) on page 527</p> <p>show services accounting errors inline-jflow fpc-slot (When IPv4, IPv6, VPLS, and Bridge Are Configured) on page 527</p> <p>show services accounting errors inline-jflow (MX80 Router When Both IPv4 and IPv6 Are Configured) on page 528</p> <p>show services accounting errors inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured) on page 528</p>
Output Fields	Table 36 on page 490 lists the output fields for the show services accounting errors command. Output fields are listed in the approximate order in which they appear.

Table 47: 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.
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.)

Table 47: show services accounting errors Output Fields (continued)

Field	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.
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 .
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
  TFE 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 flow

Syntax	<pre>show services accounting flow <inline-jflow fpc-slot <i>slot-number</i> logical-system (all <i>logical-system</i>) name (* all <i>service-name</i>)></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Junos OS Release 10.0 added the capability to display output from multiple sampling instances.</p>
Description	Display active flow statistics.
Options	<p>none—Display active flow statistics for all service instances.</p> <p>logical-system (all <i>logical-system</i>)—(Optional) Display active flow statistics for the specified logical system or all logical systems on the device.</p> <p>inline-jflow (fpc-slot <i>slot-number</i>)—(Optional) Display inline flow statistics for the specified FPC.</p> <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.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show services accounting status on page 508
List of Sample Output	<p>show services accounting flow (Flow Aggregation v5/v8 Configuration) on page 530</p> <p>show services accounting flow (Flow Aggregation v9 Configuration) on page 531</p> <p>show services accounting flow name on page 531</p> <p>show services accounting flow name all on page 531</p> <p>show services accounting flow (Multiple Sampling Instances) on page 532</p> <p>show services accounting flow inline-jflow fpc-slot (for IPv4 Flow) on page 532</p> <p>show services accounting flow inline-jflow fpc-slot (with IPv4, IPv6, VPLS, and Bridge Configuration) on page 532</p> <p>show services accounting flow inline-jflow (MX80 Router with IPv4 and IPv6 Configuration) on page 533</p> <p>show services accounting flow inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured) on page 533</p>
Output Fields	<p>Table 37 on page 495 lists the output fields for the show services accounting flow command. Output fields are listed in the approximate order in which they appear.</p>

Table 48: 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.
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 slot-number 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-detail

Syntax show services accounting flow-detail
 <detail | extensive | terse>
 <filters>
 <limit *limit-value*>
 <name (* | all | *service-name*)>
 <order (bytes | packets)>

Release Information Command introduced before Junos OS Release 7.4.

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.

filters—(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:

- **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 *limit-value*—(Optional) Limit the display output to the specified number of flows. The default is no limit.

name (* | all | *service-name*)—(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

List of Sample Output [show services accounting flow-detail on page 536](#)
[show services accounting flow-detail limit on page 537](#)
[show services accounting flow-detail name extensive on page 537](#)
[show services accounting flow-detail limit order bytes on page 537](#)
[show services accounting flow-detail name detail source-port on page 538](#)

Output Fields [Table 38 on page 500](#) 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 49: 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
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

Table 49: show services accounting flow-detail Output Fields (continued)

Field Name	Field Description	Output Level
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
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 interface	Source address	Source port	Output 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 address	Destination port	Packet count	Byte count	Time since last 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      Source      Source  Output
         interface address      port    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      Destination      Packet      Byte      Time since last

```

address	port	count	count	active timeout...
192.168.128.2		0	16	12148
192.168.144.2		0	16	15229
192.168.192.2		0	16	13296
192.168.16.2		0	16	13924
192.168.48.2		0	16	13428

Packet count for	Byte count for
last active timeout	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
```

show services accounting memory

Syntax	show services accounting memory
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display memory and flow record statistics.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show services accounting memory (Monitoring PIC Interface) on page 540 show services accounting memory (Service PIC Interface) on page 540
Output Fields	Table 39 on page 504 lists the output fields for the show services accounting memory command. Output fields are listed in the approximate order in which they appear.

Table 50: 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).
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
```

show services accounting packet-size-distribution

Syntax	show services accounting packet-size-distribution <name (* all <i>service-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display a packet size distribution histogram.
Options	<p>none—Display a packet size distribution histogram of all accounting services.</p> <p>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.</p>
Required Privilege Level	view
List of Sample Output	show services accounting packet-size-distribution name on page 541
Output Fields	Table 40 on page 506 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 51: 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 [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.
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

show services accounting status

Syntax	show services accounting status <inline-jflow fpc-slot <i>slot-number</i> name (* all <i>service-name</i>)>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 13.2R2 for EX Series switches.
Description	Display available Physical Interface Cards (PICs) for accounting services.
Options	<p>none—Display available PICs for all accounting services.</p> <p>inline-jflow fpc-slot <i>slot-number</i>—(Optional) Display inline flow accounting status for the specified FPC. For a two-member MX Series Virtual Chassis or EX9200 Virtual Chassis, the master 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.</p> <p>name (* all <i>service-name</i>)—(Optional) Display available PICs. Use a wildcard character, specify all services, or provide a specific services name.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> show services accounting flow on page 494 Inline Flow Monitoring for Virtual Chassis Overview
List of Sample Output	<p>show services accounting status name (Monitoring PIC Interface) on page 544</p> <p>show services accounting status name (Service PIC Interface) on page 545</p> <p>show services accounting status inline-jflow fpc-slot (When IPv4, IPv6 and Bridge Family Are Configured) on page 545</p> <p>show services accounting status inline-jflow (MX80 Router When Both IPv4 and IPv6) on page 545</p> <p>show services accounting status inline-jflow fpc-slot (PTX1000 Router When Both IPv4 and IPv6 Are Configured) on page 545</p>
Output Fields	Table 41 on page 508 lists the output fields for the show services accounting status command. Output fields are listed in the approximate order in which they appear.

Table 52: show services accounting status Output Fields

Field	Field Description
Service Accounting interface	Name of the service accounting interface.

Table 52: show services accounting status Output Fields (continued)

Field	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.
FPC Slot	Slot number of the FPC for which the flow information is displayed.
Local interface index	Index counter of the local interface.
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
```

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 usage

Syntax	<code>show services accounting usage</code> <code><name service-name></code>
Release Information	Command introduced before Junos OS Release 7.4.
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
List of Sample Output	show services accounting usage (Monitoring PIC Interface) on page 547 show services accounting usage (Service PIC Interface) on page 547
Output Fields	Table 42 on page 511 lists the output fields for the show services accounting usage command. Output fields are listed in the approximate order in which they appear.

Table 53: show services accounting usage Output Fields

Output Field	Output 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, (default sampling) , 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.
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%
```


CHAPTER 13

Dynamic Flow Capture Commands

- `clear services dynamic-flow-capture`
- `show services dynamic-flow-capture content-destination`
- `show services dynamic-flow-capture control-source`
- `show services dynamic-flow-capture statistics`

clear services dynamic-flow-capture

Syntax	<code>clear services dynamic-flow-capture capture-group <i>group-name</i></code> <code><criteria-identifier <i>identifier</i>></code> <code><destination-identifier <i>identifier</i>></code> <code><force></code> <code><static></code>
Release Information	Command introduced in Junos OS Release 7.4.
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
List of Sample Output	clear services dynamic-flow-capture on page 550
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
```


show services dynamic-flow-capture content-destination

Syntax	<code>show services dynamic-flow-capture content-destination capture-group <i>group-name</i> destination-identifier <i>identifier</i> <terse></code>
Release Information	Command introduced in Junos OS Release 7.4.
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	<p>capture-group <i>group-name</i>—Display information for the specified capture-group identifier.</p> <p>destination-identifier <i>identifier</i>—Display information for the specified content destination identifier.</p> <p>terse—(Optional) Display summary information.</p>
Required Privilege Level	view
List of Sample Output	show services dynamic-flow-capture content-destination capture-group on page 552
Output Fields	Table 54 on page 551 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 54: 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.
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
```

show services dynamic-flow-capture control-source

Syntax	<code>show services dynamic-flow-capture control-source capture-group <i>group-name</i> control-source source-identifier <i>identifier</i> <detail terse></code>
Release Information	Command introduced in Junos OS Release 7.4.
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	<p>capture-group <i>group-name</i>—Capture group identifier.</p> <p>source-identifier <i>identifier</i>—Control source identifier.</p> <p>detail terse—(Optional) Display the specified level of output.</p>
Required Privilege Level	view
List of Sample Output	<p>show services dynamic-flow-capture control-source source-identifier capture-group on page 554</p> <p>show services dynamic-flow-capture control-source source-identifier capture-group detail on page 554</p>
Output Fields	Table 55 on page 553 lists the output fields for the show services dynamic-flow-capture control-source command. Output fields are listed in the approximate order in which they appear.

Table 55: 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.

Table 55: show services dynamic-flow-capture control-source Output Fields (continued)

Output Field	Output Field Description
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 source-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

show services dynamic-flow-capture statistics

Syntax	<code>show services dynamic-flow-capture statistics capture-group <i>group-name</i></code>
Release Information	Command introduced in Junos OS Release 7.4.
Description	(M320 Series routers and T Series routers only) Display statistics information about the capture group specified for dynamic flow capture.
Options	capture-group <i>group-name</i> —Display information for the specified capture group identifier.
Required Privilege Level	view
List of Sample Output	show services dynamic-flow-capture statistics capture-group on page 557
Output Fields	Table 56 on page 556 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 56: 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 56: *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.
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

CHAPTER 14

Flow Collection Commands

- clear services flow-collector statistics
- request services flow-collector change-destination primary interface
- request services flow-collector change-destination secondary interface
- request services flow-collector test-file-transfer
- show services flow-collector file interface
- show services flow-collector input interface
- show services flow-collector interface

clear services flow-collector statistics

Syntax	clear services flow-collector statistics (all interface <i>interface-name</i>)
Release Information	Command introduced before Junos OS Release 7.4.
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 <i>interface-name</i> —Clear statistics for the specified flow collector interface (<i>cp-fpc/pic/port</i>).
Required Privilege Level	network
List of Sample Output	clear services flow-collector statistics on page 560
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
```

request services flow-collector change-destination primary interface

Syntax	request services flow-collector change-destination primary interface <i>cp-fpc/pic/port</i> <clear-files> <clear-logs> <immediately gracefully>
Release Information	Command introduced before Junos OS Release 7.4.
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	<p>none—Switch to the primary FTP server.</p> <p>cp-fpc/pic/port—Use the specified flow collector interface name for the primary destination.</p> <p>clear-files—(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.</p> <p>clear-logs—(Optional) Request clearing of existing logs when the switch takes place.</p> <p>immediately gracefully—(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.</p>
Required Privilege Level	maintenance
List of Sample Output	request services flow-collector change-destination primary interface on page 561
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
```

request services flow-collector change-destination secondary interface

Syntax	<code>request services flow-collector change-destination secondary interface <i>cp-fpc/pic/port</i></code> <code><clear-files></code> <code><clear-logs></code> <code><immediately gracefully></code>
Release Information	Command introduced before Junos OS Release 7.4.
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	<p>none—Switch to the secondary FTP server.</p> <p><i>cp-fpc/pic/port</i>—Use the specified flow collector interface name (<i>cp-fpc/pic/port</i>) for the secondary destination.</p> <p>clear-files—(Optional) Request clearing of existing data files in the FTP wait queue when the switch takes place.</p> <p>clear-logs—(Optional) Request clearing of existing logs when the switch takes place.</p> <p>immediately gracefully—(Optional) Specify whether you want the switch to take place immediately, or to affect only newly created files.</p>
Required Privilege Level	maintenance
List of Sample Output	request services flow-collector change-destination secondary interface on page 562
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
```

request services flow-collector test-file-transfer

Syntax	<code>request services flow-collector test-file-transfer <i>filename</i> interface (all <i>cp-fpc/pic/port</i>) (channel-zero channel-one) (primary secondary)</code>
Release Information	Command introduced before Junos OS Release 7.4.
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	<p><i>filename</i>—Name of the test file to transfer.</p> <p>interface (all <i>cp-fpc/pic/port</i>)—Transfer a test file of flows from all configured flow collector interfaces or from only the specified interface.</p> <p>channel-zero channel-one—Transfer a file from export channel 0 (unit 0) or channel 1 (unit 1) of the PIC.</p> <p>primary secondary—Transfer a file to the primary or secondary server configured as a flow collector.</p>
Required Privilege Level	network
List of Sample Output	request services flow-collector test-file-transfer interface channel-one primary on page 563
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
```

show services flow-collector file interface

Syntax	show services flow-collector file interface (all cp-fpc/pic/port) <detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M40e, M160, and M320 Series routers and T Series routers only) Display information about flow collector files.
Options	<p>none—Display file information for all configured flow collector interfaces.</p> <p>all cp-fpc/pic/port—Display file information for all configured flow collector interfaces or for the specified interface.</p> <p>detail extensive terse—(Optional) Display the specified level of output.</p>
Additional Information	No entries are displayed for files that have been successfully transferred.
Required Privilege Level	view
List of Sample Output	show services flow-collector file interface extensive on page 565
Output Fields	Table 57 on page 564 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 57: 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 57: *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

```

show services flow-collector input interface

Syntax	<code>show services flow-collector input interface (all cp-fpc/pic/port)</code> <detail extensive terse>
Release Information	Command introduced before Junos OS Release 7.4.
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	<p>none—Display packets received by all configured flow collector interfaces.</p> <p>all cp-fpc/pic/port—Display packets received by all configured flow collector interfaces or by the specified interface.</p> <p>detail extensive terse—(Optional) Display the specified level of output.</p>
Required Privilege Level	view
List of Sample Output	show services flow-collector input interface on page 566 show services flow-collector input interface all on page 567
Output Fields	Table 58 on page 566 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 58: 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
```

show services flow-collector interface

Syntax	<code>show services flow-collector interface (all <i>cp-fpc/pic/port</i>)</code> <code><detail extensive terse></code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M40e, M160, and M320 Series routers and T Series routers only) Display overall statistics for the flow collector application.
Options	<p>none—Display statistics for flow collector applications on all interfaces.</p> <p>all <i>cp-fpc/pic/port</i>—Display statistics for flow collector applications on all interfaces or for the specified interface.</p> <p>detail extensive terse—(Optional) Display the specified level of output.</p>
Required Privilege Level	view
List of Sample Output	show services flow-collector interface all detail on page 571 show services flow-collector interface all extensive on page 571 show services flow-collector interface all terse on page 573 show services flow-collector interface extensive on page 573
Output Fields	Table 59 on page 568 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 59: 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
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

Table 59: show services flow-collector interface Output Fields (continued)

Output Field	Output Field Description	Level of Output
Memory	Bytes used on the PIC and bytes free.	detail extensive
Input	Incoming flow collector packet statistics: <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
Allocation	Data block statistics: <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	File statistics, incremented since the PIC last booted: <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
Throughput	Throughput statistics: <ul style="list-style-type: none"> • Uncompressed bytes—Total uncompressed data size for all files created on this PIC. <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 compressed data size for all files created on this PIC. <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. 	detail extensive

Table 59: show services flow-collector interface Output Fields (continued)

Output Field	Output Field Description	Level of Output
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
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
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      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      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
```


CHAPTER 15

Passive Flow Monitoring Commands

- `clear passive-monitoring statistics`
- `show passive-monitoring error`
- `show passive-monitoring flow`
- `show passive-monitoring memory`
- `show passive-monitoring status`
- `show passive-monitoring usage`

clear passive-monitoring statistics

Syntax	clear passive-monitoring statistics (all interface <i>interface-name</i>)
Release Information	Command introduced in Junos OS Release 7.6.
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 (<i>mo-fpc/pic/port</i>).
Required Privilege Level	network
List of Sample Output	clear passive-monitoring statistics on page 576
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
```

show passive-monitoring error

Syntax	<code>show passive-monitoring error (* all mo-fpc/pic/port)</code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring error statistics.
Options	<code>* all mo-fpc/pic/port</code> —Display error statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.
Required Privilege Level	view
List of Sample Output	show passive-monitoring error all on page 578
Output Fields	Table 60 on page 577 lists the output fields for the show passive-monitoring error command. Output fields are listed in the approximate order in which they appear.

Table 60: 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	State of the passive monitoring interface: <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.
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.

Table 60: show passive-monitoring error Output Fields (continued)

Field Name	Field Description
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

```

show passive-monitoring flow

Syntax	show passive-monitoring flow (* all mo- <i>fpc/pic/port</i>)
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M40e, M160, and M320 Series routers and T Series routers only) Display passive flow statistics.
Options	* all mo- <i>fpc/pic/port</i> —Display passive flow statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.
Required Privilege Level	view
List of Sample Output	show passive-monitoring flow all on page 580
Output Fields	Table 61 on page 579 lists the output fields for the show passive-monitoring flow command. Output fields are listed in the approximate order in which they appear.

Table 61: 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.
Interface state	State of the passive monitoring interface: <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.

Table 61: show passive-monitoring flow Output Fields (continued)

Field Name	Field Description
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 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

```

show passive-monitoring memory

Syntax	<code>show passive-monitoring memory (* all mo-fpc/pic/port)</code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring memory and flow record statistics
Options	<code>* all mo-fpc/pic/port</code> —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
List of Sample Output	show passive-monitoring memory all on page 582
Output Fields	Table 62 on page 581 lists the output fields for the <code>show passive-monitoring memory</code> command. Output fields are listed in the approximate order in which they appear.

Table 62: 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
```


show passive-monitoring status

Syntax	<code>show passive-monitoring status (* all mo-fpc/pic/port)</code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring status.
Options	<code>* all mo-fpc/pic/port</code> —Display status for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.
Required Privilege Level	view
List of Sample Output	show passive-monitoring status all on page 584
Output Fields	Table 63 on page 583 lists the output fields for the show passive-monitoring status command. Output fields are listed in the approximate order in which they appear.

Table 63: 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.

Table 63: show passive-monitoring status Output Fields (continued)

Output Field	Output Field Description
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
```

show passive-monitoring usage

Syntax	<code>show passive-monitoring usage (* all mo-fpc/pic/port)</code>
Release Information	Command introduced before Junos OS Release 7.4.
Description	(M40e, M160, and M320 Series routers and T Series routers only) Display passive monitoring usage statistics.
Options	<code>* all mo-fpc/pic/port</code> —Display usage statistics for monitoring interfaces. Use a wildcard character, specify all interfaces, or provide a specific interface name.
Required Privilege Level	view
List of Sample Output	show passive-monitoring usage all on page 585
Output Fields	Table 64 on page 585 lists the output fields for the <code>show passive-monitoring usage</code> command. Output fields are listed in the approximate order in which they appear.

Table 64: 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.
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%
```