

Next Gen Services Interfaces User Guide for Routing Devices

Next Gen Services Interfaces User Guide for Routing Devices

Published
2020-09-24

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

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

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

Next Gen Services Interfaces User Guide for Routing Devices
Copyright © 2020 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

About the Documentation | xxviii

Documentation and Release Notes | xxviii

Using the Examples in This Manual | xxviii

 Merging a Full Example | xxix

 Merging a Snippet | xxix

Documentation Conventions | xxx

Documentation Feedback | xxxiii

Requesting Technical Support | xxxiii

 Self-Help Online Tools and Resources | xxxiv

 Creating a Service Request with JTAC | xxxiv

1

Overview

Next Gen Services Overview | 2

Next Gen Services Overview | 2

 MX Series 5G Universal Router Services Overview | 2

 Adaptive Services Overview | 3

 Inline Services | 4

 Next Gen Services | 4

 Summary of Services Supported on MX Series 5G Universal Routers | 4

 Next Gen Services Documentation | 6

 Enabling Next Gen Services | 7

 Compatibility with Other Services Cards | 7

 Configuring the MX-SPC3 Services Card | 8

 Methods for Applying Services to Traffic | 8

Configuration Overview | 10

Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10

 Overview | 10

 Interfaces | 12

 Service Set | 15

Stateful Firewall | 18**Rules and Policies | 18****Address Lists and Ranges | 20****Applications | 21****Traceoptions and Counters | 22****Carrier Grade Network Address Translation (CGNAT) | 23****Intrusion Detection System (IDS) | 49****Migrate from the MS Card to the MX-SPC3 | 57****Next Gen Services Feature Configuration Overview | 59****Service Rules and Rule Sets | 59****Service Sets | 59****Interface-Style Service Set | 59****Next-Hop-Style Service Set | 59****Services Interfaces | 60****How to Configure Services Interfaces for Next Gen Services | 60****How to Configure Interface-Style Service Sets for Next Gen Services | 62****How to Configure Next-Hop Style Service Sets for Next Gen Services | 63****How to Configure Service Set Limits for Next Gen Services | 65****Enabling and Disabling Next Gen Services | 67****Loading the Software Images on RE-S-X6-64G-UB | 67****Enabling Next Gen Services on an MX Series Router | 68****Disabling Next Gen Services on an MX Series Router | 69****Determining Whether Next Gen Services is Enabled on an MX Series Router | 70****Global System Logging Overview and Configuration | 71****Understanding Next Gen Services CGNAT Global System Logging | 71****Next Gen Services CGNAT Global System Logging | 71****Modes of Operation for Next Gen Services System Logging | 72****Understanding Stream Mode | 72****System Logging Configuration Overview | 72****Disabling Session Open Information in Syslogs | 73****Enabling Global System Logging for Next Gen Services | 73****Configuring Local System Logging for Next Gen Services | 74**

Configuring System Logging to One or More Remote Servers for Next Gen Services | 76

System Log Error Messages for Next Gen Services | 79

Session Open Logs | 79

MS-MPC Services Card | 79

MX-SPC3 Services Card | 79

Session Open Logs With NAT | 80

Session Open Logs Without NAT | 80

Session Close Logs | 81

MS-MPC Services Card | 81

MX-SPC3 Services Card | 81

NAT Out of Address Logs | 82

MS-MPC Services Card | 82

MX-SPC3 Services Card: | 82

NAT Out of Ports Logs | 82

MS-MPC Services Card | 82

MX-SPC3 Services Card | 82

NAT Rule Match Logs | 82

MS-MPC Services Card | 82

MX-SPC3 Services Card | 82

NAT Pool Release Logs | 82

MS-MPC Services Card | 83

MX-SPC3 Services Card | 83

NAT Port Block Allocation Logs | 83

MS-MPC Services Card-Example 1 | 83

MX-SPC3 Services Card-Example 1 | 83

MS-MPC Services Card-Example 2 | 83

MX-SPC3 Services Card-Example 2 | 83

NAT Port Block Allocation Interim Logs | 83

MS-MPC Services Card | 83

MX-SPC3 Services Card | 83

NAT Port Block Release Logs | 83

MS-MPC Services Card | 84

MX-SPC3 Services Card | 84

Deterministic NAT Logs | 84

MS-MPC Services Card | 84

Stateful Firewall Rule Accept Logs | 84

MS-MPC Services Card | 84

MX-SPC3 Services Card | 84

Stateful Firewall Rule Reject Logs | 85

MS-MPC Services Card | 85

MX-SPC3 Services Card | 85

Stateful Firewall Rule Discard Logs | 85

MS-MPC Services Card | 85

MX-SPC3 Services Card | 85

Stateful Firewall Rule No Rule Drop Logs | 85

MS-MPC Services Card | 86

MX-SPC3 Services Card | 86

Stateful Firewall No Policy Drop Logs | 86

MS-MPC Services Card | 86

MX-SPC3 Services Card | 86

Configuring Syslog Events for NAT Rule Conditions with Next Gen Services | 86

Next Gen Services SNMP MIBs and Traps | 88

Next Gen Services SNMP MIBs and Traps | 88

Service-Set Related SNMP MIBs | 88

Summary Mapping of MX-SPC3 CLI Services Operational Commands to SNMP MIBs | 94

NAT SNMP MIBs | 97

SNMP Traps | 100

Configuring SNMP Trap Generation | 101

SNMP Trace Logs for Traps | 104

2

Carrier Grade NAT (CGNAT)

Deterministic NAT Overview and Configuration | 109

Deterministic NAT Overview for Next Gen Services | 109

Benefits of Deterministic NAT | 109

Understanding Deterministic NAT Algorithms | 109

- Deterministic NAT Restrictions | 113

- Configuring Deterministic NAT for Next Gen Services | 114

- Configuring the NAT Pool for Deterministic NAT for Next Gen Services | 114

- Configuring the NAT Rule for Deterministic NAT44 for Next Gen Services | 116

- Configuring the NAT Rule for Deterministic NAT64 for Next Gen Services | 117

- Configuring the Service Set for Deterministic NAT for Next Gen Services | 118

- Clearing the Don't Fragment Bit | 119

Dynamic Address-Only Source NAT Overview and Configuration | 120

- Dynamic Address-Only Source Translation Overview | 120

- Benefits of Dynamic Address-Only Source Translation | 120

- Configuring Dynamic Address-Only Source NAT for Next Gen Services | 121

- Configuring the Source Pool for Dynamic Address-Only Source NAT | 121

- Configuring the NAT Source Rule for Dynamic Address-Only Source NAT | 122

- Configuring the Service Set for Dynamic Address-Only Source NAT | 124

Network Address Port Translation Overview and Configuration | 125

- Network Address Port Translation (NAPT) Overview | 125

- Benefits of NAPT | 126

- Configuring Network Address Port Translation for Next Gen Services | 126

- Configuring the Source Pool for NAPT | 126

- Configuring the NAT Source Rule for NAPT | 130

- Configuring the Service Set for NAPT | 132

- Configuring Syslog Events for NAT Rule Conditions with Next Gen Services | 133

NAT46 | 134

- NAT46 Next Gen Services Configuration Examples | 134

- NAT46 Support Summary | 135

- NAT46 Sample Configuration | 136

Stateful NAT64 Overview and Configuration | 138

Stateful NAT64 Overview | 138

Benefits of Stateful NAT64 | 138

IPv4 Addresses Embedded in IPv6 Addresses | 139

Configuring Next Gen Services Stateful NAT64 | 140

Configuring the Source Pool for Stateful NAT64 | 140

Configuring the NAT Rules for Stateful NAT64 | 144

Configuring the Service Set for Stateful NAT64 | 146

Clearing the Don't Fragment Bit | 147

IPv4 Connectivity Across IPv6-Only Network Using 464XLAT Overview and Configuration | 148

464XLAT Overview | 148

Benefits of 464XLAT | 150

IPv4 Addresses Embedded in IPv6 Addresses | 150

Configuring 464XLAT Provider-Side Translator for IPv4 Connectivity Across IPv6-Only Network for Next Gen Services | 151

Configuring the Source Pool for 464XLAT | 151

Configuring the NAT Rules for 464XLAT | 153

Configuring the Service Set for 464XLAT | 156

Clearing the Don't Fragment Bit | 157

IPv6 NAT Protocol Translation (NAT PT) | 158

IPv6 NAT PT Overview | 158

IPv6 NAT-PT Communication Overview | 159

Stateless Source Network Prefix Translation for IPv6 Overview and Configuration | 161

Stateless Source Network Prefix Translation for IPv6 | 161

Stateless Source Network Prefix Translation for IPv6 for IPv6 | 161

Benefits of Stateless Source Network Prefix Translation | 162

Configuring NPTv6 for Next Gen Services | 162

Configuring the Source Pool | 162

Configuring the NAT Rule | 162

Configuring the Service Set | 164

Transitioning to IPv6 Using Softwires | 165

6rd Softwires in Next Gen Services | 165

6rd Softwires in Next Gen Services Overview | 165

Benefits | 166

Configuring Inline 6rd for Next Gen Services | 166

Configuring a 6rd Softwire Concentrator | 166

Configuring a 6rd Softwire Rule | 167

Configuring Inline Services and an Inline Services Interface | 168

Configuring the IPv4-Facing and IPv6-Facing Interfaces for 6rd | 169

Configuring the Service Set | 169

Transitioning to IPv6 Using DS-Lite Softwires | 171

DS-Lite Softwires—IPv4 over IPv6 for Next Gen Services | 171

DS-Lite Softwires—IPv4 over IPv6 | 172

DS-Lite and NAT in Next Gen Services | 172

DS-Lite and AMS | 173

Configuring Next Gen Services DS-Lite Softwires | 174

Configuring Next Gen Services Softwire Rules | 175

Configuring Service Sets for Next Gen Services Softwires | 176

Configuring the DS-Lite Softwire | 178

DS-Lite Subnet Limitation | 180

DS-Lite Per Subnet Limitation Overview | 180

Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks | 182

Protecting CGN Devices Against Denial of Service (DOS) Attacks | 184

Mapping Refresh Behavior | 184

EIF Inbound Flow Limit | 184

Reducing Traffic and Bandwidth Requirements Using Port Control Protocol | 185

Port Control Protocol Overview | 185

Benefits of Port Control Protocol | 186

Port Control Protocol Version 2 | 187

Configuring Port Control Protocol | 188

Configuring PCP Server Options | 189

Configuring a PCP Rule | 191

- Configuring a NAT Rule | 192
- Configuring a Service Set to Apply PCP | 193
- SYSLOG Message Configuration | 193

Transitioning to IPv6 Using Mapping of Address and Port with Encapsulation (MAP-E) | 195

Mapping of Address and Port with Encapsulation (MAP-E) for Next Gen Services | 195

- Understanding Mapping of Address and Port with Encapsulation (MAP-E) | 195
 - Benefits of Mapping of Address and Port with Encapsulation (MAP-E) | 196
 - Mapping of Address and Port with Encapsulation (MAP-E) Terminology | 196
 - Mapping of Address and Port with Encapsulation (MAP-E) Functionality | 196
 - Mapping of Address and Port with Encapsulation (MAP-E) Supported and Unsupported Features | 197
- Configuring Mapping of Address and Port with Encapsulation (MAP-E) for Next Gen Services | 198

Equal Cost Multiple Path (ECMP) support for Mapping of Address and Port with Encapsulation (MAP-E) | 202

- Equal Cost Multiple Path (ECMP) support for Mapping of Address and Port with Encapsulation (MAP-E) | 202
 - Benefits | 202
- Disabling auto-routes to support ECMP with Mapping of Address and Port with Encapsulation (MAP-E) | 203

Monitoring and Troubleshooting Softwires | 207

- Ping and Traceroute for DS-Lite | 207
- Monitoring Softwire Statistics | 207
- Monitoring CGN, Stateful Firewall, and Softwire Flows | 209

Port Forwarding Overview and Configuration | 211

Port Forwarding for Next Gen Services | 211

- Port Forwarding Overview | 212
 - Benefits | 212
- Configuring Port Forwarding with Static Destination Address Translation for Next Gen Services | 212
 - Configuring the Destination Pool for Destination Address Translation | 212
 - Configuring the Mappings for Port Forwarding | 213
 - Configuring the NAT Rule for Port Forwarding with Destination Address Translation | 213

Configuring the Service Set for Port Forwarding with Destination Address Translation	215
Configuring Port Forwarding without Static Destination Address Translation for Next Gen Services	215
Configuring the Mappings for Port Forwarding	216
Configuring the NAT Rule for Port Forwarding without Destination Address Translation	216
Configuring the Service Set for Port Forwarding without Destination Address Translation	217

Port Translation Features Overview and Configuration | 219

Address Pooling and Endpoint Independent Mapping for Port Translation | 219

Address Pooling	219
Endpoint Independent Mapping and Endpoint Independent Filtering	220

Round-Robin Port Allocation | 221

Secured Port Block Allocation for Port Translation | 222

Static Source NAT Overview and Configuration | 223

Static Source NAT Overview | 223

Benefits	224
Configuring Static Source NAT44 or NAT66 for Next Gen Services	224
Configuring the Source Pool for Static Source NAT44 or NAT66	224
Configuring the NAT Rule for Static Source NAT44 or NAT66	225
Configuring the Service Set for Static Source NAT44 or NAT66	226

Static Destination NAT Overview and Configuration | 228

Static Destination NAT Overview | 228

Benefits of Static Destination NAT	228
------------------------------------	-----

Configuring Static Destination NAT for Next Gen Services | 229

Configuring the Destination Pool for Static Destination NAT	229
Configuring the NAT Rule for Static Destination NAT	229
Configuring the Service Set for Static Destination NAT	231

Twice NAPT Overview and Configuration | 233

Twice NAPT Overview | 233

- Benefits | 233**

Configuring Twice NAPT for Next Gen Services | 234

- Configuring the Source and Destination Pools for Twice NAPT | 234**

- Configuring the NAT Rules for Twice NAPT | 238**

- Configuring the Service Set for Twice NAPT | 241**

Twice NAT Overview and Configuration | 243

Twice Static NAT Overview | 243

- Benefits | 243**

Configuring Twice Static NAT44 for Next Gen Services | 244

- Configuring the Source and Destination Pools for Twice Static NAT44 | 244**

- Configuring the NAT Rules for Twice Static NAT44 | 245**

- Configuring the Service Set for Twice Static NAT44 | 247**

Twice Dynamic NAT Overview | 248

- Benefits | 248**

Configuring Twice Dynamic NAT for Next Gen Services | 249

- Configuring the Source and Destination Pools for Twice Dynamic NAT | 249**

- Configuring the NAT Rules for Twice Dynamic NAT | 250**

- Configuring the Service Set for Twice Dynamic NAT | 253**

Class of Service Overview and Configuration | 255

Class of Service for Services PICs (Next Gen Services) | 255

- Class of Service Overview for Services PICs (Next Gen Services) | 255**

- Benefits | 256**

Configuring CoS for Traffic Processed by a Services PIC (Next Gen Services) | 256

- Configuring CoS Rules | 256**

- Configuring Application Profiles for CoS Rules | 259**

- Configuring CoS Rule Sets | 260**

- Configuring the Service Set for CoS | 261**

3

Stateful Firewall Services

Stateful Firewall Services Overview and Configuration | 264

Stateful Firewall Overview for Next Gen Services | 264

- Benefits | 264

- Flows and Conversations | 264

- Stateful Firewall Rules | 265

- Stateful Firewall Anomaly Checking | 265

Configuring Stateful Firewalls for Next Gen Services | 267

- Configuring Stateful Firewall Rules for Next Gen Services | 267

- Configuring Stateful Firewall Rule Sets for Next Gen Services | 269

- Configuring the Service Set for Stateful Firewalls for Next Gen Services | 269

4

Intrusion Detection Services

IDS Screens for Network Attack Protection Overview and Configuration | 272

Understanding IDS Screens for Network Attack Protection | 272

- Intrusion Detection Services | 272

- Benefits | 273

- Session Limits | 273

- Suspicious Packet Patterns | 274

Configuring Network Attack Protection With IDS Screens for Next Gen Services | 276

- Configuring the IDS Screen Name, Direction, and Alarm Option | 276

- Configuring Session Limits in the IDS Screen | 277

- Configuring Suspicious Packet Pattern Detection in the IDS Screen | 281

- Configuring the Service Set for IDS | 284

5

Traffic Load Balancing

Traffic Load Balancing Overview and Configuration | 287

Traffic Load Balancer Overview | 287

- Traffic Load Balancing Support Summary | 287

- Traffic Load Balancer Application Description | 288

- Traffic Load Balancer Modes of Operation | 289

- Transparent Mode Layer 2 Direct Server Return | 290

- Translated Mode | 291

Transparent Mode Layer 3 Direct Server Return	292
Traffic Load Balancer Functions	292
Traffic Load Balancer Application Components	293
Servers and Server Groups	293
Server Health Monitoring – Single Health Check and Dual Health Check	293
Virtual Services	294
Traffic Load Balancer Configuration Limits	295
Configuring TLB	296

Loading the TLB Service Package	296
Configuring a TLB Instance Name	297
Configuring Interface and Routing Information	297
Configuring Servers	300
Configuring Network Monitoring Profiles	301
Configuring Server Groups	302
Configuring Virtual Services	304
Configuring Tracing for the Health Check Monitoring Function	307

6

DNS Request Filtering

DNS Request Filtering Overview and Configuration | 311

DNS Request Filtering for Disallowed Website Domains | 311

Overview of DNS Request Filtering	311
Benefits	312
Disallowed Domain Filter Database File	313
DNS Filter Profile	313
How to Configure DNS Request Filtering	313
How to Configure a Domain Filter Database	313
How to Configure a DNS Filter Profile	314
How to Configure a Service Set for DNS Filtering	319

DNS Request Filtering System Logging Error Messages | 320

System Logging for DNS Request Filtering Overview	320
DNS Match-Event Syslog Format	321
Reason Mask Values & Interpretations for DNS Filtering	324
Per-Term Statistics Syslog Format	325
DNS Filtering Disallow-List File Add/Change Syslog Format	326

7

DNS Filtering Summary Report Statistics Syslog Format | 328

DNS Filtering Per-Client-IP Statistics Syslog Format | 328

URL Filtering

URL Filtering | 332

URL Filtering Overview | 332

URL Filter Database File | 333

URL Filter Profile Caveats | 334

Configuring URL Filtering | 335

8

Integration of Juniper Sky ATP and Web filtering on MX Routers

Integration of Juniper Sky ATP and Web filtering on MX Routers | 342

Integration of Juniper ATP Cloud and Web filtering on MX Routers | 342

Overview | 342

Benefits | 343

Understanding Policy Enforcer and Juniper ATP Cloud | 343

Security Intelligence (SecIntel) - Overview | 344

Web Filtering (URL-Filterd) - Overview | 345

Configuring the Web Filter Profile for Sampling | 347

Associate a Sampling Instance with the FPC | 347

Configure a Sampling Instance and Associate the Template With the Sampling Instance. | 348

Configure the sample instance and associate the flow-server IP address and other parameters. | 349

Example: Configuring Web-filter Profile to Define Different Threat-Levels | 350

9

Aggregated Multiservices Interfaces

Enabling Load Balancing and High Availability Using Multiservices Interfaces | 353

Understanding Aggregated Multiservices Interfaces for Next Gen Services | 353

Aggregated Multiservices Interface | 353

IPv6 Traffic on AMS Interfaces Overview | 356

Member Failure Options and High Availability Settings | 357

Warm Standby Redundancy | 358

Configuring Aggregated Multiservices Interfaces | 359

Configuring Load Balancing on AMS Infrastructure | 361

Configuring AMS Infrastructure | 362

Configuring High Availability | 363

Load Balancing Network Address Translation Flows | 364

Configuring Warm Standby for Services Interfaces | 365

Inter-Chassis Services PIC High Availability

Inter-Chassis Services PIC High Availability Overview and Configuration | 368

Next Gen Services Inter-chassis High Availability Overview for NAT, Stateful Firewall, and IDS Flows | 368

Inter-chassis High Availability Overview for NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 368

Benefits | 369

Example: Next Gen Services Inter-Chassis Stateful High Availability for NAT and Stateful Firewall (MX-SPC3) | 369

Inter-Chassis Stateful Synchronization for Long Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 381

Inter-Chassis Stateful Synchronization Overview | 381

Benefits | 382

Configuring Inter-Chassis Stateful Synchronization for Long- Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 382

Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with non-AMS Interface | 383

Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with AMS Interface | 385

Inter-Chassis Services Redundancy Overview for Next Gen Services | 388

Introduction to Inter-Chassis Services Redundancy | 389

Benefits | 389

Services Redundancy Components | 389

Services Redundancy Operation | 390

Configuring Inter-Chassis Services Redundancy for Next Gen Services | 391

Configuring Non-Stop Services Redundancy for Next Gen Services Service Set | 392

Configuring One-Way Services Redundancy for Next Gen Services Service Set | 398

Application Layer Gateways

Enabling Traffic to Pass Securely Using Application Layer Gateways | 411

Next Gen Services Application Layer Gateways | 411

RTSP | 411

SIP | 412

Configuring SIP | 412

SIP ALG Interaction with Network Address Translation | 414

Junos OS SIP ALG Limitations | 420

Configuring Application Sets | 421

Configuring Application Properties for Next Gen Services | 421

Configuring an Application Protocol | 422

Configuring the Network Protocol | 424

Configuring the ICMP Code and Type | 425

Configuring Source and Destination Ports | 427

Configuring the Inactivity Timeout Period | 428

Configuring SIP | 428

SIP ALG Interaction with Network Address Translation | 430

Junos OS SIP ALG Limitations | 436

Configuring an SNMP Command for Packet Matching | 437

Examples: Configuring Application Protocols | 437

Verifying the Output of ALG Sessions | 438

FTP Example | 438

Sample Output | 439

FTP System Log Messages | 441

Analysis | 442

Troubleshooting Questions | 444

RTSP ALG Example | 444

Sample Output for MS-MPCs | 444

Sample Output for MX-SPC3 Services Card | 445

Analysis | 446

Troubleshooting Questions | 446

System Log Messages | 448

System Log Configuration | 448

System Log Output | 449

NAT, Stateful Firewall, and IDS Flows

Inline NAT Services Overview and Configuration | 451

Inline Static Source NAT Overview | 451

Benefits | 451

Configuring Inline Static Source NAT44 for Next Gen Services | 452

Configuring the Source Pool for Inline Static Source NAT44 | 452

Configuring the NAT Rule for Inline Static Source NAT44 | 453

Configuring the Service Set for Inline Static Source NAT44 | 454

Configuring Inline Services and an Inline Services Interface | 454

Inline Static Destination NAT Overview | 455

Benefits | 455

Configuring Inline Static Destination NAT for Next Gen Services | 455

Configuring the Destination Pool for Inline Static Destination NAT | 456

Configuring the NAT Rule for Inline Static Destination NAT | 456

Configuring the Service Set for Inline Static Destination NAT | 458

Configuring Inline Services and an Inline Services Interface | 458

Inline Twice Static NAT Overview | 459

Benefits | 459

Configuring Inline Twice Static NAT44 for Next Gen Services | 460

Configuring the Source and Destination Pools for Inline Twice Static NAT44 | 460

Configuring the NAT Rules for Inline Twice Static NAT44 | 461

Configuring the Service Set for Inline Twice Static NAT44 | 463

Configuring Inline Services and an Inline Services Interface | 464

Configuration Statements

Configuration Statements | 466

address (Address Book Next Gen Services) | 473

address (NAT Pool Next Gen Services) | 474

address-pooling (Source NAT Next Gen Services) | 475

aggregations (IDS Screen Next Gen Services) | 476

alarm-without-drop (IDS Screen Next Gen Services) | 477

white-list | 478

allow-overlapping-pools (NAT Next Gen Services) | 479

application (NAT Next Gen Services) | 480

application-profile (Services CoS Next Gen Services) | 481

application-protocol | 483

application-set | 485

applications (Services ALGs) | 486

automatic (Source NAT Next Gen Services) | 487

bad-option (IDS Screen Next Gen Services) | 488

block-allocation (Source NAT Next Gen Services) | 489

block-frag (IDS Screen Next Gen Services) | 490

by-destination (IDS Screen Next Gen Services) | 491

bypass-traffic-on-exceeding-flow-limits | 493

by-protocol (IDS Screen Next Gen Services) | 494

by-source (IDS Screen Next Gen Services) | 496

category (System Logging) | 498

child-inactivity-timeout | 500

clat-prefix (Source NAT Next Gen Services) | 501

clear-dont-fragment-bit (NAT Next Gen Services) | 502

close-timeout | 503

cos-rule-sets (Service Set Next Gen Services) | 504

cos-rules (Service Set Next Gen Services) | 505

cpu-load-threshold | 506

cpu-throttle (Next Gen Services) | 507

data (FTP) | 508

description (Security Policies Next Gen Services) | 509

destination-address (NAT Next Gen Services) | 510

destination-address-name (NAT Next Gen Services) | 511

destination-prefix (Destination NAT Next Gen Services) | 512

deterministic (Source NAT Next Gen Services) | 513

deterministic-nat-configuration-log-interval (Source NAT Next Gen Services) | 514

disable-global-timeout-override | 515

dns-filter | 516

dns-filter-template | **518**

drop-member-traffic (Aggregated Multiservices) | **521**

dscp (Services CoS) | **522**

ds-lite | **523**

ei-mapping-timeout (Source NAT Next Gen Services) | **525**

enable-asymmetric-traffic-processing (Service Set Next Gen Services) | **526**

enable-rejoin (Aggregated Multiservices) | **527**

enable-subscriber-analysis (Services Options VMS Interfaces) | **528**

event-rate (Next Gen Services Service-Set Local System Logging) | **529**

file (Next Gen Services Global System Logging) | **530**

files (Next Gen Services Global System Logging) | **531**

filename (Next Gen Services Global System Logging) | **532**

filtering-type (Source NAT Next Gen Services) | **533**

fin-no-ack (IDS Screen Next Gen Services) | **534**

flag (Next Gen Services Global System Logging) | **535**

format (Next Gen Services Service-Set Remote System Logging) | **536**

forwarding-class (Services PIC Classifiers) | **537**

forwarding-class (Services PIC Classifiers) | **538**

forwarding-class (Services PIC Classifiers) | **539**

fragment (IDS Screen Next Gen Services) | **540**

fragment-limit | **541**

ftp (Services CoS Next Gen Services) | **542**

gate-timeout | **543**

global-dns-stats-log-timer | **544**

group (Traffic Load Balancer) | **545**

hash-keys (Interfaces) | **547**

header-integrity-check (Next Gen Services) | **549**

high-availability-options (Aggregated Multiservices) | **551**

host (Next Gen Services Service-Set Remote System Logging) | **552**

host-address-base (Source NAT Next Gen Services) | **553**

inactivity-timeout | **554**

inactivity-asymm-tcp-timeout (Service Set Next Gen Services) | **555**

icmp (IDS Screen Next Gen Services) | **556**

icmp-type | **557**

icmpv6-malformed (IDS Screen Next Gen Services) | 558

ip (IDS Screen Next Gen Services) | 559

ipv6-extension-header (IDS Screen Next Gen Services) | 561

limit-session (IDS Screen Next Gen Services) | 563

inline-services (PIC level) | 565

ipv6-extension-header (IDS Screen Next Gen Services) | 566

instance (Traffic Load Balancer) | 568

interface-service (Services Interfaces) | 570

land (IDS Screen Next Gen Services) | 571

large (IDS Screen Next Gen Services) | 572

limit-session (IDS Screen Next Gen Services) | 573

load-balancing-options (Aggregated Multiservices) | 575

local-category (Next Gen Services Service-Set Local System Logging) | 577

local-log-tag (Next Gen Services Service-Set System Logging) | 579

loose-source-route-option (IDS Screen Next Gen Services) | 580

many-to-one (Aggregated Multiservices) | 581

map-e | 583

mapping-timeout (Source NAT Next Gen Services) | 585

mapping-type (Source NAT Next Gen Services) | 586

match (Next Gen Services Global System Logging) | 587

match (Services CoS Next Gen Services) | 588

match (Stateful Firewall Rule Next Gen Services) | 590

match-direction (NAT Next Gen Services) | 591

match-rules-on-reverse-flow (Next Gen Services) | 592

max-session-setup-rate (Service Set) | 593

max-sessions-per-subscriber (Service Set Next Gen Services) | 594

maximum | 595

member-failure-options (Aggregated Multiservices) | 596

member-interface (Aggregated Multiservices) | 599

mode (Next Gen Services Service-Set System Logging) | 601

name (Next Gen Services Global System Logging) | 602

nat-options (Next Gen Services) | 603

nat-rule-sets (Service Set Next Gen Services) | 604

next-hop-service | 605

no-remote-trace (Next Gen Services Global System Logging) | 606

no-translation (Source NAT Next Gen Services) | 607

no-world-readable (Next Gen Services Global System Logging) | 608

off (Destination NAT Next Gen Services) | 609

open-timeout | 610

pcp-rules | 611

ping-death (IDS Screen Next Gen Services) | 612

policy (Services CoS Next Gen Services) | 613

policy (Stateful Firewall Rules Next Gen Services) | 615

pool (Destination NAT Next Gen Services) | 616

pool (Source NAT Next Gen Services) | 617

pool (NAT Rule Next Gen Services) | 618

pool-default-port-range (Source NAT Next Gen Services) | 619

pool-utilization-alarm (Source NAT Next Gen Services) | 620

port (Source NAT Next Gen Services) | 621

port-forwarding (Destination NAT Next Gen Services) | 622

port-forwarding-mappings (Destination NAT Rule Next Gen Services) | 623

port-round-robin (Source NAT Next Gen Services) | 624

ports-per-session | 625

preserve-parity (Source NAT Next Gen Services) | 626

preserve-range (Source NAT Next Gen Services) | 627

profile (Traffic Load Balancer) | 628

profile (Web Filter) | 632

protocol (Applications) | 635

range (Source NAT Next Gen Services) | 637

rate | 638

real-service (Traffic Load Balancer) | 639

reassembly-timeout | 640

record-route-option (IDS Screen Next Gen Services) | 641

redistribute-all-traffic (Aggregated Multiservices) | 642

redundancy-event (Services Redundancy Daemon) | 643

redundancy-options (Aggregated Multiservices) | 645

redundancy-options (Stateful Synchronization) | 646

redundancy-policy (Interchassis Services Redundancy) | 648

redundancy-set | 650

redundancy-set-id (Service Set) | 652

rejoin-timeout (Aggregated Multiservices) | 653

rpc-program-number | 654

rtlog (Next Gen Services Global System Logging) | 655

rule (Destination NAT Next Gen Services) | 656

rule (Services CoS Next Gen Services) | 657

rule (PCP) | 659

rule (Source NAT Next Gen Services) | 661

rule-set (Services CoS Next Gen Services) | 662

rule-set (Softwires Next Gen Services) | 663

secure-nat-mapping (Source NAT Next Gen Services) | 664

security-intelligence | 665

security-intelligence-policy | 667

security-option (IDS Screen Next Gen Services) | 668

server (pcp) | 669

service-domain | 671

service-interface (Services Interfaces) | 672

services-options (Next Gen Services Interfaces) | 673

service-set (Interfaces) | 677

service-set (Services) | 678

service-set-options (Next Gen Services Services) | 682

session-limit | 683

session-limit (Service Set Next Gen Services) | 684

session-timeout (Service Set Next Gen Services) | 685

severity (Next Gen Services Service-Set Remote System Logging) | 686

sip (Services CoS Next Gen Services) | 688

size (Next Gen Services Global System Logging) | 689

snmp-command | 690

snmp-trap-thresholds (Next Gen Services) | 691

softwire-name (Next Gen Services) | 692

softwires (Next Gen Services) | 694

softwire-name (Next Gen Services) | 695

softwire-options | 697

software-types (Next Gen Services) | 698

softwires-rule-set (Service Set Next Gen Services) | 702

source-address (Next Gen Services Service-Set Remote System Logging) | 703

source-address (NAT Next Gen Services) | 704

source-address-name (NAT Next Gen Services) | 705

source-port | 706

source-route-option (IDS Screen Next Gen Services) | 707

stateful-firewall-rules (Service Set Next Gen Services) | 708

stateful-firewall-rule-set (Next Gen Services) | 709

stateful-firewall-rule-sets (Service Set Next Gen Services) | 710

stream (Next Gen Services Service-Set Remote System Logging) | 711

stream-option (IDS Screen Next Gen Services) | 712

strict-source-route-option (IDS Screen Next Gen Services) | 713

syn-ack-ack-proxy (IDS Screen Next Gen Services) | 714

syn-fin (IDS Screen Next Gen Services) | 715

syn-frag (IDS Screen Next Gen Services) | 716

syslog (Services CoS) | 717

syslog (Next Gen Services Service-Set System Logging) | 718

tcp-no-flag (IDS Screen Next Gen Services) | 719

tcp-session (Service Set Next Gen Services) | 720

tcp-tickles (Service Set Next Gen Services) | 721

tear-drop (IDS Screen Next Gen Services) | 722

then (Services CoS Next Gen Services) | 723

then (Stateful Firewall Rule Next Gen Services) | 725

timestamp-option (IDS Screen Next Gen Services) | 726

traceoptions (Next Gen Services Service-Set Flow) | 727

traceoptions (Traffic Load Balancer) | 731

traceoptions (Next Gen Services Global System Logging) | 734

traceoptions (Next Gen Services Softwires) | 735

traffic-load-balance (Traffic Load Balancer) | 737

transport (Next Gen Services Syslog Message Security) | 739

ttl-threshold | 740

unknown-protocol (IDS Screen Next Gen Services) | 741

url-filter | 742

url-filter-profile | 744
url-filter-template | 745
uuid | 747
v6rd | 748
video (Application Profile) | 749
video (Application Profile) | 750
virtual-service (Traffic Load Balancer) | 751
voice | 753
voice (Application Profile) | 754
web-filter | 755
web-filter-profile | 757
winnuke (IDS Screen Next Gen Services) | 758
world-readable (Next Gen Services Global System Logging) | 759

Operational Commands

Operational Commands | 761

clear log (Next Gen Services) | 764
clear services alg statistics | 765
clear services nat source mappings | 766
clear services sessions | 769
clear services sessions analysis | 773
clear services stateful-firewall flows | 774
clear services stateful-firewall sip-call | 777
clear services stateful-firewall sip-register | 780
clear services stateful-firewall statistics | 783
clear services subscriber analysis | 784
clear services web-filter statistics profile | 785
request services web-filter update dns-filter-database | 787
request services web-filter validate dns-filter-file-name | 788
request system disable unified-services | 789
request system enable unified-services | 790
show interfaces load-balancing (Aggregated Multiservices) | 791
show log | 796
show services alg conversations | 804

show services alg statistics | 812

show services cos statistics (Next Gen Services) | 829

show services inline software statistics | 833

show services inline ip-reassembly statistics | 838

show services nat destination pool | 844

show services nat destination rule | 846

show services nat destination summary | 849

show services nat ipv6-multicast-interfaces | 851

show services nat resource-usage source-pool | 854

show services nat source deterministic | 856

show services nat source mappings address-pooling-paired | 859

show services nat source mappings endpoint-independent | 863

show services nat source mappings pcg | 866

show services nat source mappings summary | 868

show services nat source pool | 870

show services nat source port-block | 876

show services nat source rule | 879

show services nat source rule-application | 882

show services nat source summary | 884

show services pcg statistics | 886

show services policies | 890

show services policies detail | 893

show services policies hit-count | 896

show services policies interface | 897

show services policies service-set | 898

show services redundancy-group | 899

show services screen ids-option (Next Gen Services) | 909

show services screen-statistics service-set (Next Gen Services) | 911

show services security-intelligence category summary | 916

show services security-intelligence update status | 919

show services service-sets cpu-usage | 920

show services service-sets memory-usage | 922

show services service-sets plug-ins | 925

show services service-sets statistic screen-drops (Next Gen Services) | 926

show services service-sets statistic screen-session-limit-counters (Next Gen Services) | 936

show services service-sets statistics integrity-drops | 943

show services service-sets statistics packet-drops | 949

show services service-sets statistics syslog | 951

show services service-sets statistics tcp | 959

show services service-sets summary | 961

show services sessions (Next Gen Services) | 963

show services sessions (Aggregated Multiservices) | 977

show services sessions analysis | 986

show services sessions analysis (USF) | 991

show services sessions count | 996

show services sessions service-set | 997

show services sessions service-set | 998

show services sessions softwire | 1000

show services sessions utilization | 1005

show services softwire | 1006

show services softwire flows | 1008

show services softwire statistics | 1013

show services stateful-firewall conversations | 1024

show services stateful-firewall flow-analysis | 1029

show services stateful-firewall flows | 1035

show services stateful-firewall sip-call | 1042

show services stateful-firewall sip-register | 1048

show services stateful-firewall statistics | 1052

show services stateful-firewall statistics application-protocol sip | 1063

show services subscriber analysis | 1067

show services tcp-log | 1070

show services traffic-load-balance statistics | 1071

show services web-filter dns-resolution profile | 1086

show services web-filter dns-resolution-statistics profile template | 1090

show services web-filter secintel-policy status | 1096

show services web-filter statistics dns-filter-template | 1099

show services web-filter statistics profile | 1102

show system unified-services status | 1108

About the Documentation

IN THIS SECTION

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

Use this guide to understand and configure Next Gen Services on MX240, MX480, and MX960 routers.

Documentation and Release Notes

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

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

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

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

[Table 1 on page xxxi](#) 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 xxxi defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces or emphasizes important new terms. Identifies guide names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>Junos OS CLI User Guide</i> RFC 1997, <i>BGP Communities Attribute</i>

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

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

GUI Conventions


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

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

Documentation Feedback

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

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



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

Requesting Technical Support

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

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

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

Self-Help Online Tools and Resources

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

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

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

Creating a Service Request with JTAC

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

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

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

1

PART

Overview

[Next Gen Services Overview | 2](#)

[Configuration Overview | 10](#)

[Global System Logging Overview and Configuration | 71](#)

[Next Gen Services SNMP MIBS and Traps | 88](#)

Next Gen Services Overview

IN THIS CHAPTER

- [Next Gen Services Overview | 2](#)

Next Gen Services Overview

IN THIS SECTION

- [MX Series 5G Universal Router Services Overview | 2](#)
- [Adaptive Services Overview | 3](#)
- [Next Gen Services | 4](#)
- [Summary of Services Supported on MX Series 5G Universal Routers | 4](#)
- [Next Gen Services Documentation | 6](#)
- [Enabling Next Gen Services | 7](#)
- [Compatibility with Other Services Cards | 7](#)
- [Configuring the MX-SPC3 Services Card | 8](#)
- [Methods for Applying Services to Traffic | 8](#)

This topic provides an overview of Next Gen Services and includes the following topics

MX Series 5G Universal Router Services Overview

MX Series 5G Universal routers support several types of Services interfaces, which provide specific capabilities for inspecting, monitoring and manipulating traffic as it transits an MX Series router. Services can be categorized into Adaptive Services and Next Gen Services, with each category providing Inline services interfaces and Multiservices interfaces options. [Table 3 on page 3](#) lists the cards that provide these services.

NOTE: The MX-SPC3 replaces MS- type cards providing a significant overall performance improvement together with high-end scale and capacity.

Table 3: MX Series 5G Universal Router Services

MX Series 5G Universal Routing Platform					
Adaptive Services				Next Gen Services	
MPC	MS-DPC	MS-MPC	MS-MIC	MPC	MX-SPC3
si-1/0/0	ms-1/0/0	ms-1/0/0	ms-1/0/0	si-1/0/0	vms-1/0/0
Inline services				Inline services	

- Adaptive Services can run on MS-DPC, MS-MPC, and MS-MIC cards using Multiservices (MS) PICs or Adaptive Services (AS) PICs.
- Next Gen Services can run on MPC cards and the MX-SPC3 security services card.

Inline services are configured on MX Series Modular Port Concentrators (MPC)s. Inline services interfaces, are virtual physical interfaces that reside on the Packet Forwarding Engine. They provide high performance processing on traffic transiting the MPC, and allow you to maximize your chassis slot capacity and utilization.

Multiservices Security cards (MS-DPC, MS-MPC, MS-MIC or MX-SPC3), provide services that can be applied to any traffic transiting the MX chassis beyond just an individual MPC. They also provide dedicated processing to support a variety of security features at scale and high performance.

Adaptive Services Overview

Adaptive Services run inline on MPCs and on MS-DPC, MS-MPC, and MS-MIC Multiservice security cards. Adaptive Services (AS) PICs and Multiservices PICs enable you to perform multiple services on the same PIC by configuring a set of services and applications. The AS and Multiservices PICs offer a range of services that you can configure in one or more service sets.

NOTE: On Juniper Networks MX Series 5G Universal Routing Platforms, the MS-DPC provides essentially the same capabilities as the MS-MPC. The interfaces on both platforms are configured in the same way.

For more information about Adaptive Services including inline services, see [Adaptive Services Overview](#).

Inline Services

Adaptive Services also use *inline services interfaces* to provide *inline* services. Inline services interfaces are virtual interfaces that reside on the Packet Forwarding Engine.

You configure inline services only on MPCs using the naming convention **si-fpc/pic/port** rather than the **ms-fpc/pic/port** naming convention.

Next Gen Services

Next Gen Services provide the combined capabilities of MX and SRX security services enabling you to inspect, monitor and manipulate traffic as it transits the MX Series router. Next Gen Services are supported both inline on Modular Port Concentrators (MPCs) and the MX-SPC3 security services card in MX240, MX480 and MX960 routers. Please refer to [Table 4 on page 5](#), which provides a summary of Next Gen Services that are supported both inline and on the MX-SPC3 card. Both Inline and MX-SPC3 based services can be used at the same time.

You configure Next Gen Services on the MX-SPC3 security services card using the *virtual multiservices* naming convention: **vms-fpc/pic/port**.

Summary of Services Supported on MX Series 5G Universal Routers

[Table 4 on page 5](#) provides a summary of the services supported under Next Gen Services.

Table 4: Summary of Services Supported on MX Series 5G Universal Routing Platform

Next Gen Services: Inline (si-) Interface and MX-SPC3				
Service Feature	Inline Services		MX-SPC3	
	Junos OS Release	Sub-Service	Junos OS Release	Sub-Service
CGNAT	19.3R2	Basic-NAT44 and NAT66	19.3R2	Basic-NAT44
		Static Destination NAT		Basic-NAT66
		Twice-NAT44 Basic		Dynamic-NAT44
		6rd Softwires		Static Destination NAT
		NPTv6		Basic-NAT-PT
				NAPT-PT
				NAPT44
				NAPT66
				Port Block Allocation
				Deterministic-nat44 and nat64
				End Point Independent Mapping (EIM)/End Point Independent Filtering (EIF)
				Persistent NAT – Application Pool Pairing (APP)
				Twice-NAT44 – Basic, Dynamic and NAPT
				NAT64
				XLAT-464
				NPTv6
			20.1R1	Port Control Protocol (PCP) – v1 and v2
				DS-Lite
	20.2R1	MAP-E		NAT46

Table 4: Summary of Services Supported on MX Series 5G Universal Routing Platform (*continued*)

Next Gen Services: Inline (si-) Interface and MX-SPC3				
Service Feature	Inline Services		MX-SPC3	
	Junos OS Release	Sub-Service	Junos OS Release	Sub-Service
Traffic Load Balancer	19.3R2		19.3R2	
SecIntel (SkyATP IP Threat Feeds)	19.3R2		N/A	
Stateful Firewall Services	N/A		19.3R2	
Intrusion Detection Services (IDS)	N/A		19.3R2	
DNS Request Filtering	N/A		19.3R2	
Aggregated Multiservices Interfaces	N/A		19.3R2	
Inter-chassis High Availability	N/A		19.3R2	CGNAT, Stateful Firewall, IDS
URL Filtering	N/A		20.1R1	
JFlow	20.1R1		N/A	
RPM and TWAMP	20.1R1		N/A	
Video Monitoring	20.1R1		N/A	

Next Gen Services Documentation

You can run Next Gen Services on the MX240, MX480, and MX960 if you have the MX-SPC3 services card installed in the router. Refer to our [TechLibrary](#) for all MX router documentation. For Next Gen Services, refer to the following documentation:

- To learn about and configure Next Gen Services, see *Next Gen Services Interfaces User Guide for Routing Devices* (this guide).
- For details on installing or replacing the MX-SPC3 card, see [MX Series 5G Universal Routing Platform Interface Module Reference](#).
- To monitor flows and sample traffic — See the [Monitoring, Sampling, and Collection Services Interfaces Feature Guide](#), which describes how to configure traffic flow monitoring, packet flow capture, traffic sampling for accounting or discard, port mirroring to an external device, and real-time performance monitoring.
- [Broadband Subscriber Services User Guide](#)

Enabling Next Gen Services

To run Next Gen Services, you must enable it on the MX Series router. This enables the operating system to run its own operating system (OS) for Next Gen Services.

There are specific steps you'll need to take if you're migrating your services from legacy services cards to the MX-SPC3. The Next Gen Services CLI differs from these legacy services. For more information, see [“Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3”](#) on [page 10](#).

Compatibility with Other Services Cards

The MX-SPC3 services card is compatible end-to-end with the MX Series Switch Fabrics, Routing Engines and MS-MPC line cards as described in [Table 5 on page 7](#).

Table 5: MX-SPC3 Services Card Compatibility with MX Series Switch Fabrics, Routing Engines and MPC Line Cards

Switch Fabric	Route Engine	MPC Line Cards
SCBE	RE-S-1800X4-16G-UPG-BB	MPC2E-3D
	RE-S-1800X4-32G-UB	MPC2-3D-NG
		MPC3E and MPC3E-3D-NG
		MPC4E-3D
		MPC-3D-16XGE

Table 5: MX-SPC3 Services Card Compatibility with MX Series Switch Fabrics, Routing Engines and MPC Line Cards (*continued*)

Switch Fabric	Route Engine	MPC Line Cards
SCBE2	RE-S-1800X4-16G-UPG-BB	MPC2E-3D
	RE-S-1800X4-32G-UB	MPC2-3D-NG
	RE-S-X6-64G-UB	MPC3E and MPC3E-3D-NG
		MPC4E-3D
		MPC5E and MPC5EQ
		MPC7E, MPC7EQ, and MPC-3D-16XGE
		MPC-3D-16XGE

Configuring the MX-SPC3 Services Card

The interfaces on the MX-SPC3 services card are referred to as a virtual multi service (vms) PIC. When you configure an MX-SPC3 interface, you specify the interface as a **vms-** interface as follows:

```
user@host# set services service-set service-set-name interface-service service-interface
vms-slot-number/pic-number/0.logical-unit-number
```

Aside from the CLI differences, you need to be aware of the basic hardware differences between multiservices (MS) type (MS-DPC, MS-MPC, and MS-MIC) cards and the MX-SPC3 services card. MS type cards contain four CPU complexes whereas the MX-SPC3 card, while more powerful, contains two CPU complexes. Each CPU complex services a single PIC, meaning that MS type cards support four PICs whereas the MX-SPC3 supports two PICs. MS type cards use special multiservices (MS) and adaptive services (AS) PICs, whereas the PICs on the MX-SPC3 card are integrated.

Because the number of PICs directly affects the number of interfaces, you might need to add logical units to each interface on the MX-SPC3 to increase the number of interfaces to four. For example, if you currently use all four interfaces on the MS type card and you have a service set per interface, you can create two logical units per interface on the MX-SPC3 to bring the total number of interfaces to four, and then reassociate the four service sets to these four logical interfaces.

Methods for Applying Services to Traffic

When you configure Next Gen Services, you can apply those services with either of the following methods:

- Apply the configured services to traffic that flows through a particular interface on the MX router.

- Apply the configured services to traffic that is destined for a particular next hop.

RELATED DOCUMENTATION

[Enabling and Disabling Next Gen Services | 67](#)

[Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)

[Adaptive Services Overview](#)

Configuration Overview

IN THIS CHAPTER

- [Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)
- [Next Gen Services Feature Configuration Overview | 59](#)
- [How to Configure Services Interfaces for Next Gen Services | 60](#)
- [How to Configure Interface-Style Service Sets for Next Gen Services | 62](#)
- [How to Configure Next-Hop Style Service Sets for Next Gen Services | 63](#)
- [How to Configure Service Set Limits for Next Gen Services | 65](#)
- [Enabling and Disabling Next Gen Services | 67](#)

Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3

IN THIS SECTION

- [Overview | 10](#)
- [Interfaces | 12](#)
- [Service Set | 15](#)
- [Stateful Firewall | 18](#)
- [Carrier Grade Network Address Translation \(CGNAT\) | 23](#)
- [Intrusion Detection System \(IDS\) | 49](#)
- [Migrate from the MS Card to the MX-SPC3 | 57](#)

Overview

Next Gen Services on the MX-SPC3 require you to configure services differently from what you are accustomed to with Adaptive Services, which run on MS type cards (MS-MPC, MS-MIC and MS-DPC).

Configuring the MX-SPC3 services card more closely aligns with the way you configure the SRX Series services gateway. Once you are familiar with this more unified approach, you should be able to configure services on these two platforms in a more seamless fashion, ultimately resulting in less training overhead and lower risk of configuration error.

Aside from the CLI differences, you need to be aware of the basic hardware differences between multiservices (MS) type (MS-DPC, MS-MPC, and MS-MIC) cards and the MX-SPC3 services card. MS type cards contain four CPU complexes whereas the MX-SPC3 card, while more powerful, contains two CPU complexes. Each CPU complex services a single PIC, meaning that MS type cards support four PICs whereas the MX-SPC3 supports two PICs. MS type cards use special multiservices (MS) and adaptive services (AS) PICs, whereas the PICs on the MX-SPC3 card are integrated.

Because the number of PICs directly affects the number of interfaces ([Table 6 on page 11](#)), you might need to add logical units to each interface on the MX-SPC3 to increase the number of interfaces to four. For example, if you currently use all four interfaces on the MS type card and you have a service set per interface, you can create two logical units per interface on the MX-SPC3 to bring the total number of interfaces to four, and then reassociate the four service sets to these four logical interfaces.

Table 6: Hardware Differences: MS Type Cards versus MX-SPC3 Card

	MS-Cards	MX-SPC3
Number of CPU complexes	4	2
Number of PICs per CPU complex	1	1
Number of interfaces per PIC	1	1
Total number of interfaces on card	4	2

NOTE: See the [MX Series 5G Universal Routing Platform Interface Module Reference](#) for more information on the MX-SPC3 hardware.

The following sections provide an overview of the basic configuration differences between services on the MS type cards and services on the MX-SPC3 card. The intent of these sections is to help you get started by using basic examples to illustrate the major changes. These examples show a subset of the CLI configuration options and do not replace the more formal treatment of the subject matter found in the Next Gen Services Interfaces User Guide for Routing Devices and the Junos OS CLI Reference Guide.

The configuration examples in these sections are presented side-by-side so you can easily see the differences between the two. The examples are intended to show you how to configure existing MS type card features on the MX-SPC3. The examples are not intended to show you how to configure new features only found on the MX-SPC3. For legibility and ease of comparison, the order of statements presented might differ slightly from the actual order of statements displayed in the CLI.

If you have a large set of existing adaptive services, we recognize that these changes might be an inconvenience to you. To help you migrate from MS type cards to the MX-SPC3, we suggest that you proceed as follows:

- Look through the examples in this guide to get an overall view of the changes required.
- Look through the set of configuration examples in knowledge base article KB35348.
- Look through this guide and the Junos OS CLI Reference Guide to understand all the features, configuration options, and syntax.
- Contact JTAC for help with your migration.

You do not need to make these configuration changes if you continue to run adaptive services on the MS type cards. However, once you deploy the MX-SPC3 on a router, you must replace all MS type cards on that router and reconfigure your services to align with the Next Gen Services configuration paradigm.

Interfaces

MS type cards use the interface naming convention **ms-1/0/0**, whereas you specify MX-SPC3 interfaces using the virtual multiservices or **vms-1/0/0** interface naming convention. There are no changes to the names of **ams** and **mams** interfaces.

In addition, a number of parameters that are configured under **services-options** on an **ms** interface are configured under **service-set-options** in a service set.

Table 7 on page 12 shows examples of these changes.

Table 7: Interfaces and Service Options

MS Type Cards	MX-SPC3
<pre>[edit interfaces] ms-5/1/0 { <...> }</pre>	<pre>[edit interfaces] # Change interface name to vms. vms-5/1/0 { <...> }</pre>

Table 7: Interfaces and Service Options (continued)

MS Type Cards	MX-SPC3
<pre>[edit interfaces] ms-5/1/0 { services-options { open-timeout 40; close-timeout 40; inactivity-tcp-timeout 10; inactivity-asymm-tcp-timeout 10; tcp-tickles 8; ignore-errors tcp; } }</pre>	<pre>[edit services] service-set sset1 { service-set-options { # Set tcp parameters under tcp-session. tcp-session { open-timeout 40; close-timeout 40; inactivity-tcp-timeout 10; inactivity-asymm-tcp-timeout 10; tcp-tickles 8; ignore-errors tcp; } } }</pre>
<pre>[edit interfaces] ms-5/1/0 { services-options { inactivity-non-tcp-timeout 40; session-timeout 10; } }</pre>	<pre>[edit services] service-set sset1 { # Set non-tcp parameters directly under # service-set-options. service-set-options { inactivity-non-tcp-timeout 40; session-timeout 10; } }</pre>
<pre>[edit interfaces] ms-5/1/0 { services-options { fragment-limit 32; reassembly-timeout 3; } }</pre>	<p>These parameters are hardcoded as follows:</p> <ul style="list-style-type: none"> • fragment-limit 62 • reassembly-timeout 2

Table 7: Interfaces and Service Options (continued)

MS Type Cards	MX-SPC3
<pre> [edit interfaces] ms-5/1/0 { services-options { session-limit { maximum 100; cpu-load-threshold 12; rate 10; } } } </pre>	<pre> [edit services] # Maximum number of sessions can be # specified per service-set. service-set sset1 { service-set-options { session-limit { maximum 100; } } } [edit interfaces] # All session-limit parameters continue to be # configurable per interface. If the maximum # number of sessions is different from the associated # service-set, the smaller number takes effect. vms-5/1/0 { services-options { session-limit { maximum 100; cpu-load-threshold 12; rate 10; } } } </pre>
<pre> [edit interfaces] ms-5/1/0 { services-options { pba-interim-logging-interval 10; } } </pre>	<pre> [edit interfaces] # Set interim-logging-interval under the nat branch. nat { source { pool src-pool { port { block-allocation { interim-logging-interval 10; } } } } } </pre>

Table 7: Interfaces and Service Options (*continued*)

MS Type Cards	MX-SPC3
<pre>[edit interfaces] ms-5/1/0 { services-options { syslog { host { <...> } } } }</pre>	See service-set syslog stream host .
<pre>[edit interfaces] ms-5/1/0 { services-options { syslog { message-rate-limit 10; } } }</pre>	<pre>[edit services] service-set sset1 { syslog { event-rate 10; } }</pre>
<pre>[edit interfaces] ms-5/1/0 { services-options { ignore-errors alg; disable-global-timeout-override; trio-flow-offload { minimum-bytes 1000; } } }</pre>	Not supported

Service Set

Table 8 on page 16 shows minor changes in the way some **service-set** parameters are configured.

Table 8: Service Set

MS Type Cards	MX-SPC3
<pre>[edit services] service-set sset1 { tcp-mss 1460; service-set-options { tcp-non-syn drop-flow-send-rst; tcp-fast-open drop; } }</pre>	<pre>[edit services] service-set sset1 { service-set-options { # Set tcp parameters under tcp-session. tcp-session { tcp-mss 1460; tcp-non-syn drop-flow-send-rst; tcp-fast-open drop; } } }</pre>
<pre>[edit services] service-set sset1 { replicate-services { replication-threshold 180; } }</pre>	<pre>[edit interfaces] # Set replication-threshold on the interface. vms-5/1/0 { redundancy-options { replication-threshold 180; } }</pre>
<pre>[edit services] service-set sset1 { syslog { host 10.1.1.1 { port 514; } } }</pre>	<pre>[edit services] service-set sset1 { syslog # Process security logs in the dataplane. mode stream; stream s1 { # Specify host to send security logs to. host { 10.1.1.1; port 514; } } } }</pre>

Table 8: Service Set (continued)

MS Type Cards	MX-SPC3
<pre>[edit services] service-set sset1 { syslog { host local; } }</pre>	<pre>[edit services] service-set sset1 { syslog # Process security logs in the control plane, # saving logs to local file specified by rtlog. mode event; } } rtlog { traceoptions { # Specify filename for logs. file rtlog size 1g; flag all; } }</pre>
<pre>[edit services] service-set sset1 { service-order <...> }</pre>	Service order is fixed.
<pre>[edit services] service-set sset1 { sampling-service <...> }</pre>	J-Flow logging is supported inline.
<pre>[edit services] service-set sset1 { tag-rule-sets <...> tag-rules <...> hcm-profile <...> hcm-url-rule-sets <...> hcm-url-rules <...> service-set-options { bypass-traffic-on-pic-failure; } }</pre>	Currently unsupported

Stateful Firewall

Rules and Policies

Stateful firewall rules on the MX-SPC3 are structured slightly differently from stateful firewall rules for services on the MS type cards. On the MX-SPC3, you enclose the rules within a **policies** wrapper, and you define the match terms and actions for the rule in a **policy** contained within the rule.

Just like a stateful firewall service on the MS type card, you create a service set to associate an interface with a rule set. A rule set contains references to one or more rules. Rules are applied sequentially in the order that you list them until a match occurs and an action taken.

Each rule contains one or more pairs of match terms and actions. On the MX-SPC3, each pair of match terms and actions is called a policy. Policies are applied sequentially in the order that you specify them until a match occurs and an action taken.

Table 9 on page 18 shows the configuration differences between stateful firewall rules on the MS card and the MX-SPC3. In particular, note the different definitions for the **permit/deny/reject** actions.

Table 9: Stateful Firewall Rules and Policies

MS Card	MX-SPC3
[edit services]	[edit services]
service-set s1 { stateful-firewall-rule-sets rule-set-basic-sfw; interface-service { service-interface ms-1/1/0; } }	service-set s1 { stateful-firewall-rule-sets rule-set-basic-sfw; interface-service { service-interface vms-1/1/0; } }
stateful-firewall {	# Enclose stateful firewall rules within the policies wrapper. policies {

Table 9: Stateful Firewall Rules and Policies (continued)

MS Card	MX-SPC3
<pre> rule Rule1 { match-direction input; term ping-https-apps { from { source-address { any } destination-address { any } applications [junos-icmp-ping junos-https]; } then { accept/reject/discard skip-ids; syslog; } } term accept { then { accept; } } } # end Rule1 </pre>	<pre> stateful-firewall-rule Rule1 { match-direction input; # Define match terms and actions in a policy. policy ping-https-apps { # Unlike the from statement, the match statement (and # source-address, destination-address, and application) # are mandatory. match { source-address any; destination-address any; application [junos-icmp-ping junos-https]; } then { # permit = allow # deny = silently drop # reject = drop and send ICMP unreachable or TCP RST permit/deny/reject # skip-ids is not supported. One possible way of # achieving this same goal is to create two # service-sets, one with IDS and one without IDS, # and route your next-hop-service # traffic to the desired service set via the associated # inside or outside interface. log; } } } policy accept { match { source-address any; destination-address any; application any; } then { permit; } } } # end Rule1 </pre>

Table 9: Stateful Firewall Rules and Policies (*continued*)

MS Card	MX-SPC3
<pre> rule Rule2 { match-direction output; term local { from { source-address { 10.1.3.2/32; } application-sets APPL-SET1; } then { accept; } } } # end Rule2 </pre>	<pre> stateful-firewall-rule Rule2 { match-direction output; policy local { match { source-address 10.1.3.2/32; destination-address any; # application can refer to an application set. application APPL-SET1; } then { permit; } } } # end Rule2 </pre>
<pre> rule-set rule-set-basic-sfw { rule Rule1; rule Rule2; } } # end stateful-firewall </pre>	<pre> # Use the stateful-firewall-rule-set element to list the # firewall rules in the order that you want them applied. stateful-firewall-rule-set rule-set-basic-sfw { stateful-firewall-rule Rule1; stateful-firewall-rule Rule2; } } # end policies </pre>

Address Lists and Ranges

Stateful firewall rules can contain match terms that refer to address ranges and lists.

On the MS card, you use **source-address-range** and **destination-address-range** elements to specify address ranges and **prefix-list** elements under **policy-options** to specify address lists. The **prefix-list** element is not for use solely for stateful firewall rules. You also use the **prefix-list** element to specify address lists for use within routing policies.

On the MX-SPC3, the **prefix-list** element is not used for stateful firewall rules. You use an **address-book** under **services** to define address lists and ranges for use within stateful firewall rules. The **prefix-list** element still exists, but is used exclusively for routing policies. You therefore need to configure both **address-book** and **prefix-list** elements if you are specifying address lists for stateful firewall rules and address lists for routing policies.

Table 10 on page 21 shows the differences between how you specify addresses for stateful firewall rules on the MS card versus the MX-SPC3.

Table 10: Addresses

MS Card	MX-SPC3
<pre> [edit] policy-options { prefix-list p1 { 10.1.22.45/32; 192.168.0.11/32; } } [edit services] stateful-firewall { rule sfw-rule { match-direction input; term banned-addresses { from { source-prefix-list { p1; } source-address-range { low 10.1.22.100 high 10.1.22.109; } } then { reject; syslog; } } } } <...> </pre>	<pre> [edit services] # Define address lists and address ranges in an address book. address-book { global { address-set p1 { address p1-a; address p1-b; } address p1-a 10.1.22.45/32; address p1-b 192.168.0.11/32; address p2 { address-range 10.1.22.100/32 { to { 10.1.22.109/32; } } } } } # end address-book policies { stateful-firewall-rule sfw-rule { match-direction input; policy banned-addresses { match { # Refer to the addresses defined in the address book. source-address [p1 p2]; destination-address any; application any; } then { deny; log; } } } } <...> </pre>

Applications

The MX-SPC3 supports more built-in Junos applications than the MS card. You can match on these built-in applications when you create a stateful firewall rule.

To see the complete list of built-in applications, use the **show groups junos-defaults applications** configuration mode command. For example:

```
[edit]
# show groups junos-defaults applications | match junos
application junos-ftp {
  application junos-ftp-data {
    application junos-tftp {
      application junos-twamp {
        application junos-rtsp {
          application junos-netbios-session {

<...>
```

Traceoptions and Counters

Stateful firewalls for Next Gen Services on the MX-SPC3 support additional capabilities to help debug and count traffic:

- **traceoptions** - Use to trace policy-related events such as policy lookups and rules-based events. The events are captured in the specified file for viewing.
- **count** - Use to count traffic-related events such as incoming/outgoing bytes and packets. View the counters using show commands:
 - **show services policies detail** - the output includes traffic-related counters when you specify the **count** option in your policy
 - **show services policies hit-count** - the hit count is always available regardless of whether you use the **count** option in your policy or not

[Table 11 on page 23](#) shows how to use the **traceoptions** and **count** elements:

Table 11: Traceoptions and Count

MS Card	MX-SPC3
Not supported	<pre> [edit services] policies { # Enable traceoptions to trace policy-related events. traceoptions { file policylogs size 10m files 5; flag all; } stateful-firewall-rule Rule1 { match-direction input; policy my-policy { match { source-address any; destination-address any; application [junos-dns-udp junos-dns-tcp]; } then { permit # Enable counting of traffic events. count; } } } # end my-policy ... </pre>

Carrier Grade Network Address Translation (CGNAT)

Configuring NAT for Next Gen Services on the MX-SPC3 is different from configuring NAT on legacy services on the MS card in a number of ways:

- On the MX-SPC3, you configure source NAT separately from destination NAT. You configure source NAT in the source branch of the configuration tree and you configure destination NAT in the destination branch of the configuration tree. Source NAT and destination NAT each has its own sets of address pools and rules in its respective branch of the configuration tree.
- On the MX-SPC3, if you configure both source NAT and destination NAT, destination NAT applies first, and then source NAT applies to the destination NAT translated result. In other words, you write the source NAT rule not based on the original packet, but based on the destination NAT translated result.
- On the MX-SPC3, you do not explicitly configure a **translation-type**. The type of translation is determined implicitly by your configuration.
- On the MX-SPC3, port translation is the default behavior for dynamic mappings (where different pre-NAT addresses might map to the same post-NAT address over time). If you do not explicitly include the **port**

statement in a pool definition, port translation takes place with a port range [1024, 65535], and the port is selected in a round robin fashion. If you do not want port translation to take place, you must add a **port** statement with the **no-translation** option. This default does not apply to static mappings where a pre-NAT address always maps to the same post-NAT address.

Table 12 on page 24 through Table 24 on page 45 show examples of how the different translation types are configured on the MX-SPC3.

Table 12: Example: Basic-NAT44

MS Card	MX-SPC3
[edit services]	[edit services]
<pre> service-set sset1 { nat-rules rule-basic-nat44; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-basic-nat44; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { source { </pre>
<pre> pool src-pool { address 10.10.10.0/24; } </pre>	<pre> pool src-pool { address { 10.10.10.0/24; } # host-address-base indicates a type of static mapping # where the base address 10.45.1.0/32 maps to the # lowest address in the pool, namely 10.10.10.0/32, # and the other addresses map sequentially from there # e.g. 10.45.1.1 maps to 10.10.10.1, and so on. # Since this is a static mapping, there is no port translation # by default. # Note that host-address-base does not have to be the # lowest address allowed by the subsequent source rule. # Any packet with a source address allowed by the source rule # but is lower than the host-address-base is discarded. host-address-base 10.45.1.0/32; } </pre>

Table 12: Example: Basic-NAT44 (continued)

MS Card	MX-SPC3
<pre> rule rule-basic-nat44 { match-direction input; term t1 { from { source-address { 10.45.1.0/24 } } then { translated { source-pool src-pool; translation-type { basic-nat44; } } } } } </pre>	<pre> rule-set rule-basic-nat44 { match-direction input; rule r1 { match { source-address 10.45.1.0/24; } then { source-nat { pool { src-pool; } } } } } </pre>
<pre> } # end nat </pre>	<pre> } # end source } # end nat </pre>

Table 13: Example: Basic-NAT66

MS Card	MX-SPC3
<pre> [edit services] service-set sset1 { nat-rules rule-basic-nat66; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> [edit services] service-set sset1 { nat-rule-sets rule-basic-nat66; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { source { </pre>
<pre> pool src-pool { address 2001:DB8:2222::0/128; } } </pre>	<pre> pool src-pool { address { 2001:DB8:2222::0/128; } } } } </pre>

Table 13: Example: Basic-NAT66 (continued)

MS Card	MX-SPC3
<pre> rule rule-basic-nat66 { match-direction input; term t1 { from { source-address { 2001:DB8:1111::0/128; } } then { translated { source-pool src-pool; translation-type { basic-nat66; } } } } } </pre>	<pre> rule-set rule-basic-nat66 { match-direction input; rule r1 { match { source-address 2001:DB8:1111::0/128; } then { source-nat { pool { src-pool; } } } } } </pre>
<pre> } # end nat </pre>	<pre> } # end source } # end nat </pre>

Table 14: Example: Dynamic-NAT44

MS Card	MX-SPC3
<pre> [edit services] service-set sset1 { nat-rules rule-dynamic-nat44; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> [edit services] service-set sset1 { nat-rule-sets rule-dynamic-nat44; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { source { </pre>

Table 14: Example: Dynamic-NAT44 (continued)

MS Card	MX-SPC3
<pre> pool src-pool { address-range low 10.10.10.2 high 10.10.10.10; } </pre>	<pre> pool src-pool { address { 10.10.10.2/32 to 10.10.10.10/32; } # Since this is implicitly a dynamic mapping, # there is port translation by default , so we need to # explicitly specify that we don't want port translation. port { no-translation; } } </pre>
<pre> rule rule-dynamic-nat44 { match-direction input; term t0 { from { applications junos-icmp-all; } then { no-translation; } } term t1 { from { destination-address { 10.99.0.2/32; } source-address-range { low 10.45.0.2 high 10.45.0.10; } } then { translated { source-pool src-pool; translation-type { dynamic-nat44; } } } } } </pre>	<pre> rule-set rule-dynamic-nat44 { match-direction input; rule r0 { match { source-address 0.0.0.0/32; application junos-icmp-all; } then { source-nat { off; } } } rule r1 { match { source-address-name addr1; destination-address 10.99.0.2/32; } then { source-nat { pool { src-pool; } } } } } </pre>

Table 14: Example: Dynamic-NAT44 (continued)

MS Card	MX-SPC3
<pre> } # end nat </pre>	<pre> } # end source } # end nat </pre>
	<pre> address-book { global { address addr1 { address-range 10.45.0.2/32 { to { 10.45.0.10/32; } } } } } </pre>

Table 15: Example: NAPT-44

MS Card	MX-SPC3
[edit services]	[edit services]
<pre> service-set sset1 { nat-rules rule-napt44; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-napt44; interface-service { service-interface vms-2/0/0; } } </pre>
nat {	nat {
	source {
<pre> pool src-pool { address 10.10.10.0/24; port { automatic; } } </pre>	<pre> pool src-pool { address { 10.10.10.0/24; } </pre> <div style="background-color: #ff4500; color: white; padding: 5px;"> <pre> # Since this is implicitly a dynamic mapping, # and there is no explicit port statement # to indicate otherwise, the default port # mapping behavior takes effect. </pre> </div> <pre> } </pre>

Table 15: Example: NAPT-44 (continued)

MS Card	MX-SPC3
<pre> rule rule-napt44 { match-direction input; term t1 { from { source-address { 10.45.1.0/24 } application-sets accept-algs; } then { translated { source-pool src-pool; translation-type { napt44; } } } } } </pre>	<pre> rule-set rule-napt44 { match-direction input; rule r1 { match { source-address 10.45.1.0/24; application accept-algs; } then { source-nat { pool { src-pool; } } } } } </pre>
<pre> } # end nat </pre>	<pre> } # end source } # end nat </pre>

Table 16: Example: napt-66

MS Card	MX-SPC3
<pre> [edit services] </pre>	<pre> [edit services] </pre>
<pre> service-set sset1 { nat-rules rule-napt66; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-napt66; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { source { </pre>

Table 16: Example: napt-66 (continued)

MS Card	MX-SPC3
<pre> pool src-pool { address 2001:DB8:2222::0/112; port { range low 20000 high 30000; } } </pre>	<pre> pool src-pool { address { 2001:DB8:2222::0/112; } port { range { 20000; to { 30000; } } } } </pre>
<pre> rule rule-napt66 { match-direction input; term t1 { from { source-address { 2001:DB8:1111::0/96; } } then { translated { source-pool src-pool; translation-type { napt66; } } } } } </pre>	<pre> rule-set rule-napt66 { match-direction input; rule r1 { match { source-address 2001:DB8:1111::0/96; } then { source-nat { pool { src-pool; } } } } } </pre>
<pre> } # end nat </pre>	<pre> } # end source } # end nat </pre>

Table 17: Example: Deterministic NAT-44

MS Card	MX-SPC3
[edit services]	[edit services]

Table 17: Example: Deterministic NAT-44 (continued)

MS Card	MX-SPC3
<pre> service-set sset1 { nat-rules rule-dnat-44; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-dnat-44; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { destination { </pre>
<pre> pool dest-pool { address 10.10.10.2/32; } </pre>	<pre> pool dest-pool { address { 10.10.10.2/32; } } </pre>
<pre> rule rule-dnat-44 { match-direction input; term t1 { from { destination-address { 10.45.0.2/32 } } then { translated { destination-pool dest-pool; translation-type { dnat-44; } } } } } } </pre>	<pre> rule-set rule-dnat-44 { match-direction input; rule r1 { match { destination-address 10.45.0.2/32; } then { destination-nat { pool { dest-pool; } } } } } } </pre>
<pre> } # end nat </pre>	<pre> } # end destination } # end nat </pre>

Table 18: Example: Stateful-NAT464

MS Card	MX-SPC3
[edit services]	[edit services]

Table 18: Example: Stateful-NAT464 (continued)

MS Card	MX-SPC3
<pre> service-set sset1 { nat-rules rule-stateful-nat464; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-stateful-nat464-src; nat-rule-sets rule-stateful-nat464-dest; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { source { </pre>
<pre> pool src-pool { address 10.10.10.0/24; port { automatic; } } } </pre>	<pre> pool src-pool { address { 10.10.10.0/24; } port { automatic { round-robin; } } } } </pre>

Table 18: Example: Stateful-NAT464 (continued)

MS Card	MX-SPC3
<pre> rule rule-stateful-nat464 { match-direction input; term t1 { from { source-address { 2001:DB8:1111::0/96; } destination-address { 2001:DB8:2222::0/96; } applications [junos-icmp-all junos-icmp-ping junos-traceroute junos-traceroute-ttl 1]; } then { translated { source-pool src-pool; clat-prefix 2001:DB8:1111::0/96; destination-prefix 2001:DB8:2222::0/96; translation-type { stateful-nat464; } } } } } </pre>	<pre> # This source rule applies after the destination rule. rule-set rule-stateful-nat464-src { match-direction input; rule r1 { match { source-address 2001:DB8:1111::0/96; # Since destination NAT happens first, the # destination IPv6 prefix has been stripped off, # resulting in an IPv4 destination address. destination-address 0.0.0.0/32; application [junos-icmp-all junos-icmp-ping junos-traceroute junos-traceroute-ttl 1]; } then { source-nat { pool { src-pool; } clat-prefix 2001:DB8:1111::0/96; } } } } </pre>
} # end nat	} # end source
	destination {

Table 18: Example: Stateful-NAT464 (continued)

MS Card	MX-SPC3
	<pre># This destination rule applies before the source rule. rule-set rule-stateful-nat464-dest { match-direction input; rule r1 { match { destination-address 2001:DB8:2222::0/96; } then { destination-nat { destination-prefix 2001:DB8:2222::0/96; } } } }</pre>
	<pre>} # end destination } # end nat</pre>

Table 19: Example: Stateful-NAT64

MS Card	MX-SPC3
[edit services]	[edit services]
<pre>service-set sset1 { nat-rules rule-stateful-nat64; interface-service { service-interface ms-1/2/0; } }</pre>	<pre>service-set sset1 { nat-rule-sets rule-stateful-nat64-src; nat-rule-sets rule-stateful-nat64-dest; interface-service { service-interface vms-2/0/0; } }</pre>
<pre>nat {</pre>	<pre>nat { source {</pre>

Table 19: Example: Stateful-NAT64 (continued)

MS Card	MX-SPC3
<pre>pool src-pool { address 10.10.10.0/24; port { automatic; random-allocation; } } mapping-timeout 500; }</pre>	<pre>pool src-pool { address { 10.10.10.0/24; } port { automatic { random-allocation; } } } mapping-timeout 500; }</pre>

Table 19: Example: Stateful-NAT64 (continued)

MS Card	MX-SPC3
<pre> rule rule-stateful-nat64 { match-direction input; term t1 { from { destination-address { 2001:DB8:2222::0/64; } } then { translated { source-pool src-pool; destination-prefix 2001:DB8:2222::0/64; translation-type { stateful-nat64; } } } } term t2 { from { destination-address { 2001:DB8:3333::0/64; } } then { translated { source-pool src-pool; destination-prefix 2001:DB8:3333::0/64; translation-type { stateful-nat64; } } } } } </pre>	<pre> # This source rule applies after the destination rule. rule-set rule-stateful-nat64-src { match-direction input; rule r1 { match { source-address 0::0/128; # Since destination NAT applies first, the # destination address is now IPv4. destination-address 0.0.0.0/32; } then { source-nat { pool { src-pool; } } } } } </pre>
} # end nat	} # end source
	destination {

Table 19: Example: Stateful-NAT64 (continued)

MS Card	MX-SPC3
	<pre># This destination rule applies before the source rule. rule-set rule-stateful-nat64-dest { match-direction input; rule r1 { match { destination-address 2001:DB8:2222::0/64; } then { destination-nat { destination-prefix 2001:DB8:2222::0/64; } } } rule r2 { match { destination-address 2001:DB8:3333::0/64; } then { destination-nat { destination-prefix 2001:DB8:3333::0/64; } } } }</pre>
	<pre>} # end destination } # end nat</pre>

Table 20: Example: Twice-Basic-NAT-44

MS Card	MX-SPC3
[edit services]	[edit services]
<pre>service-set sset1 { nat-rules rule-twice-basic-nat-44; interface-service { service-interface ms-1/2/0; } }</pre>	<pre>service-set sset1 { nat-rule-sets rule-twice-basic-nat-44-src; nat-rule-sets rule-twice-basic-nat-44-dest; interface-service { service-interface vms-2/0/0; } }</pre>

Table 20: Example: Twice-Basic-NAT-44 (continued)

MS Card	MX-SPC3
<pre> nat { pool src-pool { address 10.98.10.0/24; } pool dest-pool { address 10.99.10.0/24; } </pre>	<pre> nat { source { pool src-pool { address { 10.98.10.0/24; } # host-address-base indicates a type of static mapping where # the base address 10.10.10.0/32 maps to the lowest # address in the pool, namely 10.98.10.0/32, # and the other addresses map sequentially from there # e.g. 10.10.10.1 maps to 10.98.10.1, and so on. # Since this is a static mapping, there is no port translation # by default. # Note that host-address-base does not have to be the # lowest address allowed by the subsequent source rule. # Any packet with a source address allowed by the source rule # but is lower than the host-address-base is discarded. host-address-base 10.10.10.0/32; } } </pre>
<pre> rule rule-twice-basic-nat-44 { match-direction input; term t1 { from { source-address { 10.10.10.0/24; } destination-address { 10.20.10.0/24; } } then { translated { source-pool src-pool; destination-pool dest-pool; translation-type { twice-basic-nat-44; } } } } } </pre>	<pre> # This source rule applies after the destination rule. rule-set rule-twice-basic-nat-44-src { match-direction input; rule r1 { match { source-address 10.10.10.0/24; # Since destination NAT happens first, the destination # address refers to the NAT'd address. destination-address 10.99.10.0/24; } then { source-nat { pool { src-pool; } } } } } </pre>

Table 20: Example: Twice-Basic-NAT-44 (continued)

MS Card	MX-SPC3
} # end nat	} # end source
	destination {
	pool dest-pool { address { 10.99.10.0/24; } }
	# This destination rule applies before the source rule. rule-set rule-twice-basic-nat-44-dest { match-direction input; rule r1 { match { destination-address 10.20.10.0/24; } then { destination-nat { pool { dest-pool; } } } } }
	} # end destination
	} # end nat

Table 21: Example: Twice-Dynamic-NAT-44

MS Card	MX-SPC3
[edit services]	[edit services]
service-set sset1 { nat-rules rule-twice-dynamic-nat-44; interface-service { service-interface ms-1/2/0; } }	service-set sset1 { nat-rule-sets rule-twice-dynamic-nat-44-src; nat-rule-sets rule-twice-dynamic-nat-44-dest; interface-service { service-interface vms-2/0/0; } }

Table 21: Example: Twice-Dynamic-NAT-44 (continued)

MS Card	MX-SPC3
<pre> nat { pool src-pool { address 10.98.10.0/24; } pool dest-pool { address 10.99.10.0/24; } </pre>	<pre> nat { source { pool src-pool { address { 10.98.10.0/24; } port { no-translation; } } } </pre>
<pre> rule rule-twice-dynamic-nat-44 { match-direction input; term t1 { from { source-address { 10.10.10.0/24; } destination-address { 10.20.10.0/24; } } then { translated { source-pool src-pool; destination-pool dest-pool; translation-type { twice-dynamic-nat-44; } } } } } </pre>	<pre> # This source rule applies after the destination rule. rule-set rule-twice-dynamic-nat-44-src { match-direction input; rule r1 { match { source-address 10.10.10.0/24; # Since destination NAT happens first, the destination # address refers to the NAT'd address. destination-address 10.99.10.0/24; } then { source-nat { pool { src-pool; } } } } } </pre>
<pre> } # end nat </pre>	<pre> } # end source </pre>
	<pre> destination { </pre>

Table 21: Example: Twice-Dynamic-NAT-44 (continued)

MS Card	MX-SPC3
	<pre> pool dest-pool { # By default, address mapping in destination pools is static. address { 10.99.10.0/24; } } </pre>
	<pre> # This destination rule applies before the source rule. rule-set rule-twice-dynamic-nat-44-dest { match-direction input; rule r1 { match { destination-address 10.20.10.0/24; } then { destination-nat { pool { dest-pool; } } } } } </pre>
	<pre> } # end destination } # end nat </pre>

Table 22: Example: Twice-NAPT-44

MS Card	MX-SPC3
[edit services]	[edit services]
<pre> service-set sset1 { nat-rules rule-twice-napt-44; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-twice-napt-44-src; nat-rule-sets rule-twice-napt-44-dest; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { source { </pre>

Table 22: Example: Twice-NAPT-44 (continued)

MS Card	MX-SPC3
<pre> pool src-pool { address 10.98.10.0/24; port { automatic; secured-port-block-allocation block-size 256 max-blocks-per-address 1 active-block-timeout 300; } } pool dest-pool { address 10.99.10.2/32; } </pre>	<pre> pool src-pool { address { 10.98.10.0/24; } port { automatic { round-robin; } } block-allocation { block-size 256; maximum-blocks-per-host 1; active-block-timeout 300; } } </pre>
<pre> rule rule-twice-napt-44 { match-direction input; term t1 { from { source-address { 10.10.10.0/24; } destination-address { 10.20.10.2/32; } } then { translated { source-pool src-pool; destination-pool dest-pool; translation-type { twice-napt-44; } } } } } </pre>	<pre> # This source rule applies after the destination rule. rule-set rule-twice-napt-44-src { match-direction input; rule r1 { match { source-address 10.10.10.0/24; # Since destination NAT happens first, the # destination address refers to the NAT'd address. destination-address 10.99.10.2/32; } then { source-nat { pool { src-pool; } } } } } </pre>
<pre> } # end nat </pre>	<pre> } # end source </pre>
	<pre> destination { </pre>

Table 22: Example: Twice-NAPT-44 (continued)

MS Card	MX-SPC3
	<pre> pool dest-pool { address { 10.99.10.2/32; } } </pre>
	<div># This destination rule applies before the source rule.</div> <pre> rule-set rule-twice-napt-44-dest { match-direction input; rule r1 { match { source-address 10.10.10.0/24; destination-address 10.20.10.2/32; } then { destination-nat { pool { dest-pool; } } } } } </pre>
	<pre> } # end destination } # end nat </pre>

Table 23: Example: Deterministic-NAPT44

MS Card	MX-SPC3
[edit services]	[edit services]
<pre> service-set sset1 { nat-rules rule-deterministic-napt44; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-deterministic-napt44; interface-service { service-interface vms-2/0/0; } } </pre>
<pre> nat { </pre>	<pre> nat { source { </pre>

Table 23: Example: Deterministic-NAPT44 (continued)

MS Card	MX-SPC3
<pre> pool src-pool { address 10.10.10.0/24; port { range low 1024 high 19999; deterministic-port-block-allocation block-size 256; } mapping-timeout 120; } </pre>	<pre> pool src-pool { address { 10.10.10.0/24; } port { range { 1024; to { 19999; } } } deterministic { block-size 256; # host address specifies the subnet that you # want to apply to this pool. host address 10.2.0.0/20; } } mapping-timeout 120; } </pre>
<pre> rule rule-deterministic-napt44 { match-direction input; term t1 { from { source-address { 10.2.0.0/18; } } then { translated { source-pool src-pool; translation-type { deterministic-napt44; } mapping-type endpoint-independent; } } } } </pre>	<pre> rule-set rule-deterministic-napt44 { match-direction input; rule r1 { match { source-address 10.2.0.0/18; } then { source-nat { pool { src-pool; } mapping-type endpoint-independent; } } } } </pre>

Table 23: Example: Deterministic-NAPT44 (continued)

MS Card	MX-SPC3
} # end nat	} # end source } # end nat

Table 24: Example: Deterministic-NAPT64

MS Card	MX-SPC3
[edit services]	[edit services]
<pre> service-set sset1 { nat-rules rule-deterministic-napt64; interface-service { service-interface ms-1/2/0; } } </pre>	<pre> service-set sset1 { nat-rule-sets rule-deterministic-napt64-src; nat-rule-sets rule-deterministic-napt64-dest; interface-service { service-interface vms-2/0/0; } } </pre>
nat {	nat {
<pre> pool src-pool { address 10.98.10.0/24; port { automatic; random-allocation; } deterministic-port-block-allocation block-size 256; } } </pre>	<pre> source { pool src-pool { address { 10.98.10.0/24; } port { automatic { random-allocation; } deterministic { block-size 256; host address 2001:DB8:1111::1/120; } } } } } </pre>

Table 24: Example: Deterministic-NAPT64 (continued)

MS Card	MX-SPC3
<pre> rule rule-deterministic-napt64 { match-direction input; term t1 { from { source-address { 2001:DB8:1111::1/120; } } then { translated { destination-prefix 2001:DB8:2222::/96; source-pool src-pool; translation-type { deterministic-napt64; } } } } } </pre>	<pre> # This source rule applies after the destination rule. rule-set rule-deterministic-napt64-src { match-direction input; rule r1 { match { source-address 2001:DB8:1111::1/120; # Since destination NAT happens first, the destination # address refers to the NAT'd address. destination-address 0.0.0.0/32; } then { source-nat { pool { src-pool; } } } } } </pre>
} # end nat	} # end source
	destination {
	<pre> pool dest-pool { address { 10.99.10.2/32; } } </pre>

Table 24: Example: Deterministic-NAPT64 (continued)

MS Card	MX-SPC3
	<pre># This destination rule applies before the source rule. rule-set rule-destination-napt64-dest { match-direction input; rule r1 { match { destination-address 2001:DB8:2222::/96; } then { destination-nat { destination-prefix 2001:DB8:2222::/96; } } } }</pre>
	<pre>} # end destination } # end nat</pre>

Table 25: Example: napt-pt

MS Card	MX-SPC3
[edit services]	[edit services]
<pre>service-set sset1 { nat-rules rule-napt-pt; interface-service { service-interface ms-1/2/0; } }</pre>	<pre>service-set sset1 { nat-rule-sets rule-napt-pt-src; nat-rule-sets rule-napt-pt-dest; interface-service { service-interface vms-2/0/0; } }</pre>
nat {	nat {
<pre> pool src-pool { address 10.10.10.2/32; } pool dest-pool { address 10.99.10.2/32; } }</pre>	<pre> source { pool src-pool { address { 10.10.10.2/32; } } } }</pre>

Table 25: Example: napt-pt (continued)

MS Card	MX-SPC3
<pre> rule rule-napt-pt { match-direction input; term t1 { from { source-address { 2001:DB8:1111::2/128; } destination-address { 2001:DB8:2222::2/128; } } then { translated { source-pool src-pool; destination-pool dest-pool; translation-type { napt-pt; } } } } } </pre>	<pre> rule-set rule-napt-pt-src { match-direction input; rule r1 { match { source-address 2001:DB8:1111::2/128; destination-address 10.99.10.0/24; } then { source-nat { pool { src-pool; } } } } } </pre>
} # end nat	} # end source
	destination {
	<pre> pool dest-pool { address { 10.99.10.2/32; } } </pre>

Table 25: Example: napt-pt (continued)

MS Card	MX-SPC3
	<pre>rule-set rule-napt-pt-dest { match-direction input; rule r1 { match { destination-address 2001:DB8:2222::2/128; } then { destination-nat { pool { dest-pool; } } } } }</pre>
	<pre>} # end destination } # end nat</pre>

Intrusion Detection System (IDS)

IDS rules for Next Gen Services on the MX-SPC3 are defined under the **screen** branch. There are minor differences in the naming of the various elements, but the main change is in the behavior for detecting packets with IPv4 options and IPv6 extensions:

- For the IDS service on the MS Card, the default behavior is to detect and drop packets with IPv4 options and IPv6 extensions. If you want to allow these packets, you have to allow them explicitly through configuration.
- For the IDS Next Gen Service on the MX-SPC3, the default behavior is to allow packets with IPv4 options and IPv6 extensions. If you want to detect and drop these packets, you have to disallow them explicitly through configuration.

Table 26 on page 50 shows examples of the configuration differences.

Table 26: IDS Rules

MS Card	MX-SPC3
<pre>[edit services] service-set sset1 { ids-rules r1; ids-rules r2; }</pre>	<pre>[edit services] service-set sset1 { # Replace ids-rules with ids-option. ids-option ids1; ids-option ids2; }</pre>
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { <...> } } }</pre>	<pre>[edit services] # Define ids rules under the screen branch. screen { # Replace rule with ids-option. ids-option ids1 { match-direction input; # Flatten hierarchy by removing term and placing # contents directly under ids-option. <...> } }</pre>
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { then { allow-ip-options [loose-source-route route-record router-alert security stream-id strict-source-route timestamp]; } } } }</pre>	<pre>[edit services] screen { ids-option ids1 { match-direction input; # By default, all ip options are allowed. } }</pre>

Table 26: IDS Rules (continued)

MS Card	MX-SPC3
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { then { <no allow-ip-options configured> } } } }</pre>	<pre>[edit services] screen { ids-option ids1 { match-direction input; # Explicitly specify the disallowed options. ip { loose-source-route-option; record-route-option; security-option; stream-option; strict-source-route-option; timestamp-option; # router-alert option for IPv4 is not supported. } } }</pre>
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { then { allow-ipv6-extension-header [ah dstopts esp fragment hop-by-hop mobility routing]; } } } }</pre>	<pre>[edit services] screen { ids-option ids1 { match-direction input; # By default, all ipv6 extensions are allowed. } }</pre>

Table 26: IDS Rules (*continued*)

MS Card	MX-SPC3
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { then { <no allow-ipv6-extension-header configured> } } } }</pre>	<pre>[edit services] screen { ids-option ids1 { match-direction input; ip { # Explicitly specify the disallowed extensions. ipv6-extension-header { AH-header; ESP-header; fragment-header; hop-by-hop-header; mobility-header; routing-header; # dstoptions is not supported. } } } }</pre>
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { then { aggregation { source-prefix 24; destination-prefix 24; source-prefix-ipv6 64; destination-prefix-ipv6 64; } } } } }</pre>	<pre>[edit services] screen { ids-option ids1 { match-direction input; aggregation { source-prefix-mask 24; destination-prefix-mask 24; source-prefix-v6-mask 64; destination-prefix-v6-mask 64; } } }</pre>

Table 26: IDS Rules (*continued*)

MS Card	MX-SPC3
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { then { icmp-fragment-check; icmp-large-packet-check; } } } }</pre>	<pre>[edit services] screen { ids-option ids1 { match-direction input; # Group icmp checks under icmp. icmp { fragment; large; } } }</pre>
<pre>[edit services] ids { rule r1 { match-direction input; term t1 { then { land-attack-check; tcp-winnuke-check; tcp-syn-fragment-check; tcp-syn-defense; } } } }</pre>	<pre>[edit services] screen { ids-option ids1 { match-direction input; # Group tcp checks under tcp. tcp { land; winnuke; syn-frag; # tcp-syn-defense is not supported. } } }</pre>

Table 26: IDS Rules (*continued*)

MS Card	MX-SPC3
<pre> [edit services] ids { rule r1 { match-direction input; term t1 { then { session-limit { by-source { maximum 100; rate 10; packets 1k; } by-destination { maximum 100; rate 10; packets 1k; } } } } } } </pre>	<pre> [edit services] screen { ids-option ids1 { match-direction input; limit-session { by-source { maximum-sessions 100; session-rate 10; packet-rate 1k; } by-destination { maximum-sessions 100; session-rate 10; packet-rate 1k; } } } } </pre>

Table 26: IDS Rules (*continued*)

MS Card	MX-SPC3
<pre> [edit services] ids { rule r1 { match-direction input; term t1 { then { session-limit { by-source { by-protocol { tcp { maximum 100; rate 10; packets 1k; } udp { maximum 100; rate 10; packets 1k; } icmp { maximum 100; rate 10; packets 1k; } } } } } } } } </pre>	<pre> [edit services] screen { ids-option ids1 { match-direction input; limit-session { by-source { by-protocol { tcp { maximum-sessions 100; session-rate 10; packet-rate 1k; } udp { maximum-sessions 100; session-rate 10; packet-rate 1k; } icmp { maximum-sessions 100; session-rate 10; packet-rate 1k; } } } } } } </pre>

Table 26: IDS Rules (*continued*)

MS Card	MX-SPC3
<pre> [edit services] ids { rule r1 { match-direction input; term t1 { then { session-limit { by-destination { by-protocol { tcp { maximum 100; rate 10; packets 1k; } udp { maximum 100; rate 10; packets 1k; } icmp { maximum 100; rate 10; packets 1k; } } } } } } } } </pre>	<pre> [edit services] screen { ids-option ids1 { match-direction input; limit-session { by-destination { by-protocol { tcp { maximum-sessions 100; session-rate 10; packet-rate 1k; } udp { maximum-sessions 100; session-rate 10; packet-rate 1k; } icmp { maximum-sessions 100; session-rate 10; packet-rate 1k; } } } } } } </pre>

Migrate from the MS Card to the MX-SPC3

Use this procedure to configure a router to support Next Gen Services.

You typically use this procedure to migrate a router supporting legacy services on the MS card to a router supporting Next Gen Services on the MX-SPC3, but this procedure applies even if the router that you are migrating from does not contain MS card cards.

Because Next Gen Services configuration is not compatible with legacy service provisioning, migrating a router to support Next Gen Services on the MX-SPC3 requires you to completely deprovision and reprovision your router . Furthermore:

- You cannot install an MX-SPC3 card in a router that has MS cards.
- You cannot configure Next Gen Services on a router equipped with MS cards.
- You cannot configure legacy services on a router equipped with MX-SPC3 cards.

In other words, a router can run with either MS cards or MX-SPC3 cards but not both at the same time.

NOTE: This procedure is service affecting. You are setting the router to factory default configuration.

1. Upgrade the router to release 19.3R2.
2. Back up the current router configuration to a remote host.
3. Set the router to factory default configuration.
 - a. Load the router with the factory default configuration:

```
root# load factory-default
```

- b. Configure the management interface with the same IP address as you had before you loaded the factory default configuration:

```
root# set interfaces fxp0 unit 0 family inet address <mgt-ip-address>
```

- c. Configure SSH so that you can continue to access the router. For example:

```
root# set system services ssh root-login allow
root# set system services ssh max-sessions-per-connection 32
```

```
root# set system root-authentication plain-text-password
New password:
Retype new password:
```

d. Commit the changes.

4. Enable Next Gen Services on the router.

Junos OS provides a system-wide operational parameter that you enable if you want to configure Next Gen Services on a router. By default, this parameter is not enabled.

From operational mode:

```
root> request system enable unified-services
Before enabling unified services, please move to baseline configuration.
Are above conditions satisfied ? [yes,no]
```

NOTE: This setting is persistent and survives a reboot.

5. Reboot the router.

```
root> request system reboot
```

6. Replace the MS card cards with MX-SPC3 cards.

7. Reprovision your router.

As a starting point, you can restore the backup from step 2 but you might need to change this configuration to be compatible with Next Gen Services before you can commit.

SEE ALSO

[Next Gen Services Overview | 2](#)

[Enabling and Disabling Next Gen Services | 67](#)

Next Gen Services Feature Configuration Overview

IN THIS SECTION

- [Service Rules and Rule Sets | 59](#)
- [Service Sets | 59](#)
- [Services Interfaces | 60](#)

To configure services with Next Gen Services, you need to configure the following objects:

- Service rules
- Service sets
- Services interfaces

Service Rules and Rule Sets

Service rules specify a set of matching conditions and a set of actions to apply to traffic when it matches the conditions. For example, a stateful firewall rule can specify a destination address that must be matched, and take the action of dropping packets that have that destination address.

Service rule sets consist of a group of services rules that belong to the same category. For example, a stateful firewall rule set consists of stateful firewall rules.

Service Sets

A service set specifies one or more service rules or rule sets to apply to traffic. The service set also specifies a services interface, which indicates where the services processing is performed.

A service set is either an interface-style service set or a next-hop-style service set.

Interface-Style Service Set

The service set applies the service rules to all traffic that flows through a particular interface.

Next-Hop-Style Service Set

The service set applies the service rules to traffic that is destined for a particular next hop. You must redirect the next-hop traffic to the services interface that the service set uses.

Services Interfaces

A services interface indicates where a service is applied to traffic. Services interfaces are not physical links to external devices.

If a service is performed on an MX-SPC3 services card, the service interface has the format:

```
vms-slot-number/pic-number/port-number
```

If a service is performed on a line card's PFE (inline services), the service interface has the format *si-slot-number/pic-number/0*.

RELATED DOCUMENTATION

[Next Gen Services Overview | 2](#)

[How to Configure Services Interfaces for Next Gen Services | 60](#)

[Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)

How to Configure Services Interfaces for Next Gen Services

To configure services interfaces:

1. Configure the services interface name.

```
[edit]
user@host# set interfaces interface-name
```

Where the *interface-name* one of the following:

- *vms-slot-number/pic-number/port-number* for an MX-SPC3 services card
- *si-slot-number/pic-number/0* for a line card PFE (inline services interface)

2. Configure the unit and family for the interface.

- If you are using the services interface in an interface service set:

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

- If you are using the services interface in a next-hop service set, configure inside and outside interface units:

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family (inet | inet6)
user@host# set interfaces interface-name unit logical-unit-number service-domain inside
user@host# set interfaces interface-name unit logical-unit-number family (inet | inet6)
user@host# set interfaces interface-name unit logical-unit-number service-domain outside
```

For example:

```
[edit]
user@host# set interfaces vms-1/0/0 unit 100 family inet
user@host# set interfaces vms-1/0/0 unit 100 service-domain inside
user@host# set interfaces vms-1/0/0 unit 1000 family inet
user@host# set interfaces vms-1/0/0 unit 1000 service-domain outside
```

3. When neither NAT nor the **max-sessions-per-subscriber** statement at the **[edit service-set service-set-name service-set-options]** hierarchy level are configured, enable the creation of subscribers if you want to track subscribers.

```
[edit interfaces interface-name services-options]
user@host# set enable-subscriber-analysis
```

4. Configure CPU resource restrictions for the services interface.

```
[edit interfaces interface-name services-options session-limit]
user@host# set cpu-load-threshold percentage
```

When the CPU usage exceeds the value (percentage of the total available CPU resources), the system reduces the rate of new sessions so that the existing sessions are not affected by low CPU availability. The CPU utilization is constantly monitored, and if the CPU usage remains above the configured **cpu-load-threshold** value for a continuous period of 5 seconds, Junos OS reduces the session rate value configured at **edit interfaces *interface-name* services-options session-limit rate** by 10%. This is repeated until the CPU utilization comes down to the configured limit.

5. Configure the maximum number of sessions allowed simultaneously on a services card.

```
[edit interfaces interface-name services-options session-limit]
user@host# set maximum number
```

If you specify the maximum number of sessions to be zero, it indicates that the configuration is not effective. You must specify a value higher than zero for the maximum number of sessions.

6. Configure the maximum number of new sessions allowed per second on a services card.

```
[edit interfaces interface-name services-options session-limit]
user@host# set rate rate
```

RELATED DOCUMENTATION

[Next Gen Services Overview | 2](#)

[How to Configure Next-Hop Style Service Sets for Next Gen Services | 63](#)

[How to Configure Service Set Limits for Next Gen Services | 65](#)

[Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)

How to Configure Interface-Style Service Sets for Next Gen Services

To configure an interface service set:

1. Configure the service set name.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Specify the service interface that the service set uses to apply services.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

3. Specify the service rules that the service set applies to traffic.

For example:

```
[edit services service-set ss1]
user@host# set nat-rule-sets internal-nat
```


4. (Optional) Enable the service set to process unidirectional traffic.

```
[edit services service-set service-set-name service-set-options]
user@host# set enable-asymmetric-traffic-processing
```

5. Enable service-processing at routing engine (RE).

```
[edit services service-set service-set-name service-set-options]
user@host# set routing-engine-services
```

6. Apply the service set to an interface that is passing traffic. You can apply a service filter to apply the service set to only certain traffic on the interface.

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

For details about configuring the service-filter, see *Guidelines for Configuring Service Filters*.

The **input** option applies the service set to the input side of the interface, and the **output** option applies the service set to the output side of the interface. If you are using a bidirectional service rule in the service set, then the same service set must be used for input and output.

RELATED DOCUMENTATION

[Next Gen Services Overview | 2](#)

[How to Configure Interface-Style Service Sets for Next Gen Services | 62](#)

[How to Configure Service Set Limits for Next Gen Services | 65](#)

[Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)

How to Configure Next-Hop Style Service Sets for Next Gen Services

To configure a next-hop service set:

1. Configure the service set name.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Specify the services interface inside unit and outside unit for the service set.

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name.unit-number outside-service-interface
interface-name.unit-number
```

The inside-service-interface must be a service interface logical unit that is configured with **service-domain inside**. The outside-service-interface must be a service interface logical unit that is configured with **service-domain outside**.

3. Specify the service rules that the service set applies to traffic.

For example:

```
[edit services service-set SS1]
user@host# set nat-rule-sets internal-nat
```

4. (Optional) Enable the service set to process unidirectional traffic.

```
[edit services service-set service-set-name service-set-options]
user@host# set enable-asymmetric-traffic-processing
```

5. Configure a static route to force traffic to the inside or outside interface of the next-hop service set.

For example, if you want traffic with the destination 198.51.100.33 to be processed by the service set with the inside interface vms-1/0/0.100:

```
[edit routing-options]
user@host# set static route 198.51.100.33 next-hop vms-1/0/0.100
```

RELATED DOCUMENTATION

[Next Gen Services Overview | 2](#)

[How to Configure Interface-Style Service Sets for Next Gen Services | 62](#)

[How to Configure Service Set Limits for Next Gen Services | 65](#)

[Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)

How to Configure Service Set Limits for Next Gen Services

To configure service set limits:

1. Set the maximum number of session setups allowed per second for the service set. After this setup rate is reached, any additional session setup attempts are dropped. If you do not include the **max-session-creation-rate** statement, the session setup rate is not limited.

```
[edit services service-set service-set-name ]
user@host# set max-session-setup-rate (number | numberk)
```

If you use the *numberk* format, 1k=1000.

2. Enable packets to bypass without creating a new session when the flow in the service set exceeds the limit that is set by the **max-flows** statement at the **[edit services service-set service-set-name]** hierarchy level.

```
[edit services service-set service-set-name service-set-options]
user@host# bypass-traffic-on-exceeding-flow-limits
```

3. To limit the session open information in you system logs, you can disable it from being collected.

```
[edit services service-set service-set-name service-set-options]
user@host# set disable-session-open-syslog
```

4. Configure the maximum number of sessions allowed from a single subscriber.

```
[edit services service-set service-set-name service-set-options]
user@host# set max-sessions-per-subscriber session-number
```

5. Specify the maximum number of sessions allowed simultaneously on the service set. If you specify the maximum number of sessions to be zero, it indicates that the configuration is not effective. You must specify a value higher than zero for the maximum number of sessions.

```
[edit services service-set service-set-name service-set-options]
user@host# set session-limit maximum number
```

6. Configure the session lifetime for the service set in seconds. The session is closed after this amount of time, even if traffic is running on the session.

```
[edit services service-set service-set-name service-set-options]
user@host# set session-timeout seconds
```

7. Specify the inactivity timeout period for non-TCP established sessions.

```
user@host# set inactivity-non-tcp-timeout seconds
```

8. Configure the TCP session parameters for the service-set.

1. Set the timeout period for the Transmission Control Protocol (TCP) session tear-down.

```
[edit services service-set service-set-name services-options]
user@host# set close-timeout seconds
```

The default value is 1 second. The range is 2 through 300 seconds.

2. Configure the inactivity timeout period for asymmetric TCP established sessions

```
[edit services service-set service-set-name service-set-options tcp-session]
user@host# set inactivity-asymm-tcp-timeout seconds
```

3. Configure the number of seconds that a unidirectional TCP session can be inactive before it is closed.

```
[edit services service-set service-set-name service-set-options tcp-session]
user@host# set inactivity-tcp-timeout seconds
```

The default value is 30 seconds. The range is 4 through 86,400 seconds. Any value you configure in the application protocol definition overrides the value specified here; for more information, see [“Configuring Application Properties for Next Gen Services” on page 421](#).

4. Set the timeout period for Transmission Control Protocol (TCP) session establishment, for use with SYN-cookie defenses against network intrusion.

```
[edit services service-set service-set-name service-set-options ]
user@host# set open-timeout seconds
```

The default value is 5 seconds. The range of possible values is from 4 through 224 seconds. Any value you configure in the intrusion detection service (IDS) definition overrides the value specified here; for more information, see [“Configuring Network Attack Protection With IDS Screens for Next Gen Services” on page 276](#).

RELATED DOCUMENTATION

[Next Gen Services Overview | 2](#)

[How to Configure Interface-Style Service Sets for Next Gen Services | 62](#)

[Next Gen Services Feature Configuration Overview | 59](#)

[Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)

Enabling and Disabling Next Gen Services

IN THIS SECTION

- [Loading the Software Images on RE-S-X6-64G-UB | 67](#)
- [Enabling Next Gen Services on an MX Series Router | 68](#)
- [Disabling Next Gen Services on an MX Series Router | 69](#)
- [Determining Whether Next Gen Services is Enabled on an MX Series Router | 70](#)

To use Next Gen Services, you must first enable it on the MX Series router. This topic describes how to enable Next Gen Services, how to disable Next Gen Services, and how to determine whether Next Gen Services is enabled or disabled on your system.

Loading the Software Images on RE-S-X6-64G-UB

The Next-Gen Services MX-SPC3 services card can exhibit inconsistent behavior when the vmhost image is installed on the Next-Generation Routing Engine: RE-S-X6-64G-UB (NG-RE).

This behavior can result in you encountering one of the following:

- The MX-SPC3 card remains in **Present** state and does not come online
- The MX-SPC3 comes online successfully with different a software image (either a previously installed image or the pre-loaded image from manufacturing)

To work around this problem, you must install the **jpfe-spc3*** package manually on the NG-RE. To install this package manually, follow one of these procedures, depending on whether or not you have enabled Next Gen Services (**unified-services**) mode:

If **unified-services** are enabled:

1. Download the **jpfe-spc3*** package from: [Downloads](#)
2. Load vmhost* image on the RE
3. After the RE comes up, copy package jpfe-spc3-mx-x86-32-19.4R1.9.tgz to the /var/tmp directory
4. Load **jpfe-spc3*** package:

```
user@host> request system software add /var/tmp/jpfe-spc3-mx-x86-32-19.4R1.9.tgz reboot
```

If **unified-services** are disabled:

1. Download the **jpfe-spc3*** package from: [Downloads](#)
2. Load vmhost* image on the RE
3. After it comes up, enable **unified-servicesmode**.

```
user@host> request system enable unified-services and restart
```

4. After the RE comes up, copy package jpfe-spc3-mx-x86-32-19.4R1.9.tgz to the /var/tmp directory.
5. Load jpfe-spc3*.

```
user@host> request system software add /var/tmp/jpfe-spc3-mx-x86-32-19.4R1.9.tgz reboot
```

Enabling Next Gen Services on an MX Series Router

There are specific steps you'll need to take if you're migrating your services from MS-MPC cards to the MX-SPC3 services cards. The Next Gen Services CLI differs from these legacy services.

The following procedure is a general procedure for enabling and disabling Next Gen Services.

Before you do anything, you'll need to back up your configuration.

For more details on the differences between the configuration of the MX-SPC3 services card and legacy services cards, see [“Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3” on page 10](#) and plan your migration appropriately.

You can run Next Gen Services on the MX240, MX480 and MX960 using the MX-SPC3 services card. To use Next Gen Services on the MX Series, you must first enable Next Gen Services:

1. Delete any router configuration that is for services. This includes configuration under the **[edit services]** hierarchy, configuration for services interfaces, and any configuration that refers to services interfaces.
2. Enable Next Gen Services.

```
user@host> request system enable unified-services
```

3. When the following message appears, enter **yes**.

```
Before enabling unified services, please move to baseline configuration.
Are above conditions satisfied ? [yes,no]
```

4. Reboot the MX Series chassis.

```
user@host> request system reboot
```

You can also enable the Next Gen Services on a Guest network function (GNF), by using the CLI **request system enable unified-services** at the GNF level. For more information, see *Next Gen Services on Junos node slicing*.

Disabling Next Gen Services on an MX Series Router

To disable Next Gen Services on the MX Series:

1. Delete any router configuration that is for services. This includes configuration under the **[edit services]** hierarchy, configuration for services interfaces, and any configuration that refers to services interfaces.
2. Disable Next Gen Services.

```
user@host> request system disable unified-services
```

3. When the following message appears, enter **yes**.

```
Before disabling unified services, please move to baseline configuration.
Are above conditions satisfied ? [yes,no]
```

```
Unified-Services downgrade staged. Please reboot with 'request system reboot'
command to complete the downgrade
```

```
WARNING: cli has been replaced by an updated version:
CLI release 20190829.221548_builder.r1052644 built by builder on 2019-08-29
22:27:13 UTC
Restart cli using the new version ? [yes,no] (yes)
```

4. Reboot the MX Series chassis.

```
user@host> request system reboot
```

Determining Whether Next Gen Services is Enabled on an MX Series Router

To determine whether Next Gen Services is enabled:

- Enter the following command:

```
user@host> show system unified-services status
```

One of the following messages appears:

- **Enabled**—Next Gen Services is enabled and ready to use.
- **Unified Services : Upgrade staged , please reboot with 'request system reboot' to enable unified services.**—You must perform a system reboot before Next Gen Services is enabled.
- **Disabled**—Next Gen Services is disabled.
- **Unified Services : Upgrade staged , please reboot with 'request system reboot' to disable unified services.**—You must perform a system reboot before Next Gen Services is disabled.

RELATED DOCUMENTATION

[Next Gen Services Overview | 2](#)

[Next Gen Services Feature Configuration Overview | 59](#)

[Configuration Differences Between Adaptive Services and Next Gen Services on the MX-SPC3 | 10](#)

Global System Logging Overview and Configuration

IN THIS CHAPTER

- [Understanding Next Gen Services CGNAT Global System Logging | 71](#)
- [Enabling Global System Logging for Next Gen Services | 73](#)
- [Configuring Local System Logging for Next Gen Services | 74](#)
- [Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)
- [System Log Error Messages for Next Gen Services | 79](#)
- [Configuring Syslog Events for NAT Rule Conditions with Next Gen Services | 86](#)

Understanding Next Gen Services CGNAT Global System Logging

IN THIS SECTION

- [Next Gen Services CGNAT Global System Logging | 71](#)
- [Modes of Operation for Next Gen Services System Logging | 72](#)
- [Understanding Stream Mode | 72](#)
- [System Logging Configuration Overview | 72](#)
- [Disabling Session Open Information in Syslogs | 73](#)

All CGNAT services supported under Next Gen Services use global system logging. This topic describes global system logging for Next Gen Services CGNAT services and how to configure it.

Next Gen Services CGNAT Global System Logging

The CGNAT services supported under Next Gen Services support global system logging for syslog messages. You configure syslog messaging for these services under the **service-set** hierarchy. You can send logs to either the local routing engine (RE) or one or more remote servers (each of these is identified as a stream). You can configure files to log system messages and also assign attributes, such as severity levels, to

messages. Reboot requests are recorded to the system log files, which you can view with the **show log** command.

In the case of an AMS bundle, each PIC establishes a TCP connection with the log server and the external collector receives log messages from all the AMS members.

Modes of Operation for Next Gen Services System Logging

You can save logs for Next Gen Services locally, which is called: event mode, or send the log messages to one or more external servers, called: stream mode.

In event mode, after the log message is recorded, the log is stored within a log file which is then stored in the database table of the local routing engine (RE) for further analysis.

When configured in stream mode, log messages are streamed to one or more remote log servers. Each remote log server is assigned a stream from which it receives logs.

Understanding Stream Mode

When configured in stream mode, Next Gen Services log messages are streamed to a remote device.

For stream mode log forwarding, you can configure which transport protocol is used between MX-SPC3 services card and the log server. You can use either UDP, TCP, or TLS as the transport protocol.

When the device is configured in stream mode, you can configure a maximum of eight system log hosts to stream to.

System Logging Configuration Overview

Configuring system logging for Next Gen Services involves several main steps and considerations:

- Global system logging — Next Gen Services system logging uses a global logging option that you need to enable in order to collect system log messages.

To enable global system logging for Next Gen Services, set the **traceoptions** option under the **edit services rtlog** hierarchy.

- For Next Gen Services, syslogs are always set at the **service-set** level regardless of whether you are running event mode or stream mode.

You must configure system logging for each service-set for which you want to collect logs. Each **service-set** uses a separate TCP connection in stream mode.

As a log client, Next Gen Services initiates TCP/TLS connections to the remote log server. By default, we connect to port 514 for TCP logging [RFC 6587], and port 6514 for TLS logging [RFC 5425]. You can also specify port numbers for TCP and TLS logging using CLI.

- If you are using AMS bundles, syslogs are generated from each member interface of AMS group

Disabling Session Open Information in Syslogs

You can stop open session information from cluttering up your syslogs by disabling session open information from being collected:

```
user@host# set services service-set ss1 service-set-options disable-session-open-syslog
```

RELATED DOCUMENTATION

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

Enabling Global System Logging for Next Gen Services

To configure either event mode or stream mode system logging for Next Gen Services, you must first globally enable logging:

1. Enable system logging for Next Gen Services.

```
[edit]
user@host# edit services rtlogtraceoptions
```

2. Specify the groups from which to inherit configuration data.

```
[edit services rtlog traceoptions]
user@host# set apply-groups group-names
```

3. Specify which groups not to inherit configuration data from.

```
[edit services rtlog traceoptions]
user@host# set apply-groups-except group-names
```

4. Configure information about the files that contain trace logging information.

```
[edit services rtlog traceoptions]
```

```
user@host# set file filename
```

5. Define tracing operations for individual service-sets. To specify more than one tracing operation, include multiple flag statements.

```
[edit services rtlog traceoptions]
user@host# set flag flag, flag...
```

6. (Optional) If you prefer not to perform any system logging, you can disable it.

```
[edit services rtlog traceoptions]
user@host# set no-remote-trace
```

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

Configuring Local System Logging for Next Gen Services

To send Next Gen Services log messages to a file on the local router, you'll need to configure system logging for **event** mode. This procedure describes this configuration process.

You must enable global system logging for Next Gen Services in order to perform event mode system logging. See, "[Enabling Global System Logging for Next Gen Services](#)" on page 73.

NOTE: For Next Gen Services, syslogs are always set at the **service-set** level. You must perform this procedure for each service-set for which you want to collect logs.

To configure event mode logging for Next Gen Services:

1. Specify the filename to send log messages to.

```
user@host# set system syslog file filename
```

2. Specify the name of the service-set for which you want to log messages.

```
user@host# edit services service-set service-set-name syslog
```

For example specify the service-set name to ss1.

```
user@host# edit services service-set ss1 syslog
```

3. Specify the security transport protocol for syslog messages.

```
[edit services service-set ss1 syslog]
user@host# set transport protocol tls | tcp | udp
```

4. Enable event mode system logging for the service-set.

```
[edit services service-set ss1 syslog]
user@host# set mode event
```

5. Specify the rate at which log messages are sent per second.

```
[edit services service-set ss1 syslog]
user@host# set event-rate 100
```

6. Specify a local tag name for the log messages.

```
[edit services service-set ss1 syslog]
user@host# set local-log-tag SYSLOG
```

7. Specify the categories for which you want to collect events.

```
[edit services service-set ss1 syslog]
user@host# set local-category category, category
```

For example, to collect logs for stateful firewall, sessions and NAT:

```
[edit services service-set ss1 syslog]
user@host# set local-category sfw, session, nat
```

RELATED DOCUMENTATION

[Enabling Global System Logging for Next Gen Services | 73](#)

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

Configuring System Logging to One or More Remote Servers for Next Gen Services

To send system log messages about Next Gen Services to one or more remote servers, you can configure system logging for **stream** mode. This procedure describes the configuration process.

NOTE: Next Gen Services system log messages are configured and collected at the **service-set** level.

In this procedure, you'll configure a stream for the log messages between each service set and each remote server that you want to send log messages.

Complete this procedure for each **service-set** and each remote server for which you want to collect logs and send logs.

You must enable global system logging for Next Gen Services in order to perform stream logging. See, [“Enabling Global System Logging for Next Gen Services” on page 73](#).

To configure stream mode system logging for Next Gen Services:

1. Specify the names of the service-set for which you want to collect log messages.

```
user@host# edit services service-set service-set-name syslog
```

For example specify the service-set name to ss1.

```
user@host# edit services service-set ss1 syslog
```

2. Specify the security transport protocol for syslog messages.

```
[edit services service-set ss1 syslog]
user@host# set transport protocol tls |tcp | udp
```

3. (Optional) Specify the syslog source address.

```
[edit services service-set ss1 syslog]
user@host# set source-address 50.0.0.10
```

BEST PRACTICE: The syslog source address can be any arbitrary IP address. It does not have to be an IP address that is assigned to the device. Rather, this IP address is used on the syslog collector to identify the syslog source. The best practice is to configure the source address as the IP address of the interface that the traffic is sent out on.

4. Specify a local tag name for the log messages.

```
[edit services service-set ss1 syslog]
user@host# set local-log-tag SYSLOG
```

5. Enable stream mode system logging for the service-set.

```
[edit services service-set ss1 syslog]
user@host# set modestream
```

6. Specify a name for the stream.

```
[edit services service-set ss1 syslog]
user@host# set stream stream-name
```

For example, let's call the stream: stream-aa

```
[edit services service-set ss1 syslog]
user@host# edit stream stream-aa
```

7. Specify the categories for which you want to collect events.

```
[edit services service-set ss1 syslog stream stream-aa]
user@host# set category
```

For example, to collect logs for stateful firewall, sessions and NAT:

```
[edit services service-set ss1 syslog stream stream-aa]
user@host# set category sfw, session, nat
```

8. Specify the file format for the log.

```
[edit services service-set ss1 syslog stream stream-aa]
user@host# set format sd-syslog
```

9. Specify the IP address of syslog server to receive log messages,

```
[edit services service-set ss1 syslog stream stream-aa]
user@host# set host address
```

10. Specify the level of severity for the stream.

```
[edit services service-set ss1 syslog stream stream-aa]
user@host# set severity severity-level
```

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging](#) | 71

[Enabling Global System Logging for Next Gen Services](#) | 73

[Configuring Local System Logging for Next Gen Services](#) | 74

System Log Error Messages for Next Gen Services

IN THIS SECTION

- [Session Open Logs | 79](#)
- [Session Close Logs | 81](#)
- [NAT Out of Address Logs | 82](#)
- [NAT Out of Ports Logs | 82](#)
- [NAT Rule Match Logs | 82](#)
- [NAT Pool Release Logs | 82](#)
- [NAT Port Block Allocation Logs | 83](#)
- [NAT Port Block Allocation Interim Logs | 83](#)
- [NAT Port Block Release Logs | 83](#)
- [Deterministic NAT Logs | 84](#)
- [Stateful Firewall Rule Accept Logs | 84](#)
- [Stateful Firewall Rule Reject Logs | 85](#)
- [Stateful Firewall Rule Discard Logs | 85](#)
- [Stateful Firewall Rule No Rule Drop Logs | 85](#)
- [Stateful Firewall No Policy Drop Logs | 86](#)

This topic describes Next Gen Services MX-SPC3 services card system log error messages and provides a comparison of these messages with the MS-MPC services card.

Session Open Logs

Following are example session open logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

JSERVICES_SESSION_OPEN application source-interface-name source-address source-port
source-nat-information destination-address destination-port destination-nat-information protocol-name
software-information;

MX-SPC3 Services Card

RT_FLOW_SESSION_CREATE_USF Prefix service-set-name source-interface-name source-address
source-port destination-address destination-port service-name nat-source-address nat-source-port

nat-destination-address nat-destination-port src-nat-rule-type src-nat-rule-name dst-nat-rule-type
dst-nat-rule-name protocol-name policy-name application software-information;

Sample MX-SPC3 Output

A sample output is as follows:

```
<14>1 2018-06-26T17:23:06.269-07:00 booklet RT_FLOW - RT_FLOW_SESSION_CREATE_USF
[junos@2636.1.1.1.2.25 prefix="SYSLOG-PREFIX" service-set-name="JNPR-NH-SSET3"
source-address="50.0.0.10" source-port="1" destination-address="60.0.0.10" destination-port="21219"
connection-tag="0" service-name="icmp" nat-source-address="100.0.0.1" nat-source-port="1024"
nat-destination-address="60.0.0.10" nat-destination-port="21219" nat-connection-tag="0"
src-nat-rule-type="source rule" src-nat-rule-name="SRC-NAT-RULE1" dst-nat-rule-type="N/A"
dst-nat-rule-name="N/A" protocol-id="1" policy-name="p1" source-zone-name="JNPR-NH-SSET3-ZoneIn"
destination-zone-name="JNPR-NH-SSET3-ZoneOut" session-id-32="160000001" username="N/A"
roles="N/A" packet-incoming-interface="vms-2/0/0.100" application="UNKNOWN"
nestedapplication="UNKNOWN" encrypted="UNKNOWN" application-category="N/A"
application-sub-category="N/A" application-risk="-1"] Prefix PADDY3 svc-set-name JNPR-NH-SSET3:
session created 50.0.0.10/1->60.0.0.10/21219 0x0 icmp 100.0.0.1/1024->60.0.0.10/21219 0x0 source
rule SRC-NAT-RULE1 N/A N/A 1 p1 JNPR-NH-SSET3-ZoneIn JNPR-NH-SSET3-ZoneOut 160000001
N/A(N/A) vms-2/0/0.100 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
```

Session Open Logs With NAT

MS-MPC Services Card

```
SYSLOG_MSMPC{SS_TEST}JSERVICES_SESSION_OPEN: application:ike-esp-nat, xe-2/2/1.0 24.0.0.2:1234
[85.0.0.1:1024] -> 25.0.0.2:1234 (UDP)
```

MX-SPC3 Services Card

```
Aug 3 02:04:28 mobst480i RT_FLOW: RT_FLOW_SESSION_CREATE_USF: Tag svc-set-name sset1:
session created 90.0.0.2/1->30.0.0.2/4323 0x0 icmp 50.0.0.3/1024->30.0.0.2/4323 0x0 source rule
rule1 N/A N/A 1 p1 sset1-ZoneIn sset1-ZoneOut 160000015 N/A(N/A) vms-2/0/0.1 UNKNOWN
UNKNOWN UNKNOWN N/A N/A -1 N/A
```

Session Open Logs Without NAT

MS-MPC Services Card

```
SYSLOG_MSMPC{SS_TEST}JSERVICES_SESSION_OPEN: application:ike-esp-nat, xe-2/2/1.0 24.0.0.2:1234
-> 25.0.0.2:1234 (UDP)
```

MX-SPC3 Services Card

```
RT_FLOW - RT_FLOW_SESSION_CREATE_USF [junos@2636.1.1.1.2.25 tag="SYSLOG_SFW"
service-set-name="ss1" source-address="20.1.1.2" source-port="12000" destination-address="30.1.1.2"
destination-port="22000" connection-tag="0" service-name="None" nat-source-address="20.1.1.2"
nat-source-port="12000" nat-destination-address="30.1.1.2" nat-destination-port="22000"
nat-connection-tag="0" src-nat-rule-type="N/A" src-nat-rule-name="N/A" dst-nat-rule-type="N/A"
dst-nat-rule-name="N/A" protocol-id="6" policy-name="policy1" source-zone-name="ss1-ZoneIn"
```

```
destination-zone-name="ss1-ZoneOut" session-id-32="190000004" username="N/A" roles="N/A"
packet-incoming-interface="xe-5/3/2.0" application="UNKNOWN" nested-application="UNKNOWN"
encrypted="UNKNOWN" application-category="N/A" application-sub-category="N/A"
application-risk="-1" application-characteristics="N/A"] Tag SYSLOG_SFW svc-set-name ss1: session
created 20.1.1.2/12000->30.1.1.2/22000 0x0 None 20.1.1.2/12000->30.1.1.2/22000 0x0 N/A N/A
N/A N/A 6 policy1 ss1-ZoneIn ss1-ZoneOut 190000004 N/A(N/A) xe-5/3/2.0 UNKNOWN UNKNOWN
UNKNOWN N/A N/A -1 N/A
```

Session Close Logs

Following are example session close logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

```
JSERVICES_SESSION_CLOSE application source-interface-name source-address source-port
source-nat-information destination-address destination-port destination-nat-information protocol-name
software-information;
```

MX-SPC3 Services Card

```
RT_FLOW_SESSION_CLOSE_USF Prefix service-set-name source-interface-name source-address
source-port destination-address destination-port service-name nat-source-address nat-source-port
nat-destination-address nat-destination-port src-nat-rule-type src-nat-rule-name dst-nat-rule-type
dst-nat-rule-name protocol-name policy-name; software-information;
```

Sample MX-SPC3 Output

A sample output follows:

```
<14>1 2018-06-27T09:24:00.058-07:00 booklet RT_FLOW - RT_FLOW_SESSION_CLOSE_USF
[junos@2636.1.1.1.2.25 prefix="SYSLOG-PREFIX" service-set-name="JNPR-NH-SSET3" reason="idle
Timeout" source-address="50.0.0.10" source-port="1" destination-address="60.0.0.10"
destination-port="30170" connection-tag="0" service-name="icmp" nat-source-address="100.0.0.1"
nat-source-port="1024" nat-destination-address="60.0.0.10" nat-destination-port="30170"
nat-connection-tag="0" src-nat-rule-type="source rule" src-nat-rule-name="SRC-NAT-RULE1"
dst-nat-rule-type="N/A" dst-nat-rule-name="N/A" protocol-id="1" policy-name="p1"
source-zone-name="JNPR-NH-SSET3-ZoneIn" destination-zone-name="JNPR-NH-SSET3-ZoneOut"
session-id-32="160000001" packets-from-client="1" bytes-from-client="84" packets-from-server="0"
bytes-from-server="0" elapsed-time="4" application="UNKNOWN" nested-application="UNKNOWN"
username="N/A" roles="N/A" packet-incoming-interface="vms-2/0/0.100" encrypted="UNKNOWN"
application-category="N/A" application-sub-category="N/A" application-risk="-1"] Prefix PADDY-DEF
svc-set-name JNPR-NH-SSET3: session closed idle Timeout: 50.0.0.10/1->60.0.0.10/30170 0x0 icmp
100.0.0.1/1024->60.0.0.10/30170 0x0 source rule SRC-NAT-RULE1 N/A N/A 1 p1
JNPR-NH-SSET3-ZoneIn JNPR-NH-SSET3-ZoneOut 160000001 1(84) 0(0) 4 UNKNOWN UNKNOWN
N/A(N/A) vms-2/0/0.100 UNKNOWN N/A N/A -1
```

NAT Out of Address Logs

Following are example NAT Out of Address logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

JSERVICES_NAT_OUTOF_ADDRESSES: nat-pool-name

MX-SPC3 Services Card:

Aug 10 10:06:13 champ RT_NAT: RT_SRC_NAT_OUTOF_ADDRESSES: nat-pool-name src_pool1 is out of addresses

NAT Out of Ports Logs

Following are example NAT Out of Ports logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

{NPU-1-PFX1}[j]services-nat]: JSERVICES_NAT_OUTOF_PORTS: natpool NAT-POOL-NPU1-PFX3 is out of ports

MX-SPC3 Services Card

jul 31 03:08:30 esst480h RT_NAT: RT_SRC_NAT_OUTOF_PORTS: nat-pool-name nat_pool1 is out of ports

NAT Rule Match Logs

Following are example NAT rule match logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

SYSLOG_MSMPC{SS_TEST}[j]services-nat]: JSERVICES_NAT_RULE_MATCH: proto 17 (UDP) application: any, xe-2/2/1.0:24.0.0.2:1234 -> 25.0.0.2:1234, Match NAT rule-set: (null), rule: NAT_RULE_TEST, term: t

MX-SPC3 Services Card

RT_NAT: RT_NAT_RULE_MATCH: protocol-id 17 protocol-name udp application Unknown interface-name ge-2/0/9.0 source-address 11.1.1.2 source-port 2000 destination-address 12.1.1.2 destination-port 5000 rule-set-name rule-set rule-name nat-rule

NAT Pool Release Logs

Following are example NAT Rule Match logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

SYSLOG_MSMPC{SS_TEST}[jservices-nat]: JSERVICES_NAT_POOL_RELEASE: natpool release 85.0.0.1:1024[1]

MX-SPC3 Services Card

RT_NAT: RT_SRC_NAT_POOL_RELEASE: nat-pool-name nat-pool address 112.1.1.4 port 1024 count 1

NAT Port Block Allocation Logs

Following are example NAT port block allocation logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card-Example 1

SYSLOG_MSMPC{ss1}[jservices-nat]: JSERVICES_NAT_PORT_BLOCK_ALLOC: 11.1.1.2 -> 112.1.1.4:42494-42503 0x59412760

MX-SPC3 Services Card-Example 1

Aug 9 23:01:59 esst480r RT_NAT: RT_SRC_NAT_PBA_ALLOC: Subscriber 20.1.1.5 used/maximum [1/1] blocks, allocates port block [49774-49923] from 100.0.0.1 in source pool p1 lsys_id: 0

MS-MPC Services Card-Example 2

SYSLOG_MSMPC{ss1}[jservices-nat]: JSERVICES_NAT_PORT_BLOCK_RELEASE: 2001:2010:0:0:0:0:2 -> 161.161.16.1:56804-56813 0x597ef2c3

MX-SPC3 Services Card-Example 2

RT_NAT: RT_SRC_NAT_PBA_ALLOC: Subscriber 11.1.1.2 used/maximum [1/2] blocks, allocates port block [13934-13943] from 112.1.1.1 in source pool nat-pool lsys_id: 0

NAT Port Block Allocation Interim Logs

Following are example interim logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

SYSLOG_MSMPC{ss1}[jservices-nat]: JSERVICES_NAT_PORT_BLOCK_ACTIVE: 11.1.1.2 -> 112.1.1.4:42494-42503 0x59412760

MX-SPC3 Services Card

RT_NAT: RT_SRC_NAT_PBA_INTERIM: Subscriber 50.0.0.3 used/maximum [1/1] blocks, allocates port block [5888-6015] from 202.0.0.1 in source pool JNPR-CGNAT-PUB-POOL lsys_id: 0

NAT Port Block Release Logs

Following are example NAT port block release logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

JSERVICES_NAT_PORT_BLOCK_RELEASE source-address nat-source-address nat-source-port-range-start
nat-source-port-range-end object-create-time;

MX-SPC3 Services Card

RT_NAT: RT_SRC_NAT_PBA_RELEASE: Subscriber 11.1.1.2 used/maximum [2/3] blocks, releases port
block [3839-3843] from 112.1.2.1 in source pool nat-pool lsys_id: 0

Deterministic NAT Logs*MS-MPC Services Card*

{ss1}[jsservices-nat]: JSERVICES_DET_NAT_CONFIG: Deterministic NAT Config
[2001:2010::-2001:2010::ff]:[161.161.16.1-161.161.16.254]:0:200:0:1024-65535

Stateful Firewall Rule Accept Logs

Following are example stateful firewall rule accept logs for MS-MPC services cards versus MX-SPC3
services processing card:

MS-MPC Services Card

Sep 20 01:36:51 mobst480b (FPC Slot 5, PIC Slot 0) 2017-09-20 08:36:19:
SYSLOG_MSMPC{SS_TEST}[jsservices-sfw]: JSERVICES_SFW_RULE_ACCEPT: proto 17 (UDP) application:
any, interface: xe-2/2/1.0, 24.0.0.2:1234 -> 25.0.0.2:1234, Match SFW allow rule-set: (null), rule:
SFW_RULE_TEST, term: t

MX-SPC3 Services Card

expo RT_FLOW: RT_FLOW_SESSION_POLICY_ACCEPT_USF: Tag SYSLOGMSG svc-set-name ss1:session
created with policy accept 20.1.1.2/5->30.1.1.2/15100 0x0 icmp R11 1 sfw_policy1 ss1-ZoneIn
ss1-ZoneOut 160000010 N/A(N/A) xe-5/3/2.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A

Sample MX-SPC3 Output

Here's a sample output for MX-SPC3 card:

```
<14>1 2018-06-27T09:23:56.808-07:00 booklet RT_FLOW - RT_FLOW_SESSION_POLICY_ACCEPT_USF
[junos@2636.1.1.1.2.25 prefix="PADDY-DEF" service-set-name="JNPR-NH-SSET3"
source-address="50.0.0.10" source-port="1" destination-address="60.0.0.10" destination-port="30170"
connection-tag="0" service-name="icmp" rule-name="To be implemented" rule-set-name="To be
implemented" protocol-id="1" policy-name="p1" source-zone-name="JNPR-NH-SSET3-ZoneIn"
destination-zone-name="JNPR-NH-SSET3-ZoneOut" session-id-32="160000001"
username="N/A"roles="N/A" packet-incoming-interface="vms-2/0/0.100" application="UNKNOWN"
nested-application="UNKNOWN"encrypted="UNKNOWN" application-category="N/A"
application-sub-category="N/A" application-risk="-1"] Prefix PADDY-DEF svc-set-name JNPR-NH-SSET3:
session created 50.0.0.10/1->60.0.0.10/30170 0x0 icmp To be implemented To be implemented 1 p1
```

JNPR-NH-SSET3-ZoneIn JNPR-NH-SSET3-ZoneOut 160000001 N/A(N/A) vms-2/0/0.100 UNKNOWN
UNKNOWN UNKNOWN N/A N/A -1

Stateful Firewall Rule Reject Logs

Following are example stateful firewall rule reject logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

Sep 20 01:42:02 mobst480b (FPC Slot 5, PIC Slot 0) 2017-09-20 08:41:31:

SYSLOG_MSMPC{SS_TEST}[jsservices-sfw]: JSERVICES_SFW_RULE_REJECT: proto 17 (UDP) application: any, 24.0.0.2:1234 -> 25.0.0.2:1234, Match SFW reject rule-set: (null), rule: SFW_RULE_TEST, term: t

MX-SPC3 Services Card

expo RT_FLOW: RT_FLOW_SESSION_RULE_REJECT_USF: Tag SYSLOGMSG svc-set-name ss1: session denied 20.1.1.2/5->30.1.1.2/15183 0x0 icmp R11 1(8) sfw_policy1 ss1-ZoneIn ss1-ZoneOut UNKNOWN UNKNOWN N/A(N/A) xe-5/3/2.0 No Rejected by policy 160000030 N/A N/A -1 N/A

Stateful Firewall Rule Discard Logs

Following are example stateful firewall rule discard logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

Sep 20 01:43:57 mobst480b (FPC Slot 5, PIC Slot 0) 2017-09-20 08:43:26:

SYSLOG_MSMPC{SS_TEST}[jsservices-sfw]: JSERVICES_SFW_RULE_DISCARD: proto 17 (UDP) application: any, 24.0.0.2:1234 -> 25.0.0.2:1234, Match SFW drop rule-set: (null), rule: SFW_RULE_TEST, term: t

MX-SPC3 Services Card

RT_FLOW - RT_FLOW_SESSION_RULE_DISCARD_USF [junos@2636.1.1.1.2.25 tag="SYSLOG_SFW" service-set-name="ss1" source-address="20.1.1.2" source-port="10000" destination-address="30.1.1.2" destination-port="20000" connection-tag="0" service-name="None" rule-name="R1" rule-set-name="" protocol-id="17" icmp-type="0" policy-name="policy1" source-zone-name="ss1-ZoneIn" destination-zone-name="ss1-ZoneOut" application="UNKNOWN" nested-application="UNKNOWN" username="N/A" roles="N/A" packet-incoming-interface="xe-5/3/2.0" encrypted="No" reason="Denied by policy" session-id-32="190000014" application-category="N/A" application-sub-category="N/A" application-risk="-1" application-characteristics="N/A"] Tag SYSLOG_SFW svc-set-name ss1: session denied 20.1.1.2/10000->30.1.1.2/20000 0x0 None R1 17(0) policy1 ss1-ZoneIn ss1-ZoneOut UNKNOWN UNKNOWN N/A(N/A) xe-5/3/2.0 No Denied by policy 190000014 N/A N/A -1 N/A

Stateful Firewall Rule No Rule Drop Logs

Following are example stateful firewall rule no rule drop logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

Sep 20 01:43:57 mobst480b (FPC Slot 5, PIC Slot 0) 2017-09-20 08:43:26:

SYSLOG_MSMPC{SS_TEST}[jservices-sfw]: JSERVICES_SFW_NO_RULE_DROP: proto 17 (UDP) application: any, 24.0.0.2:1234 -> 25.0.0.2:1234

MX-SPC3 Services Card

RT_FLOW_SESSION_NO_RULE_DROP_USF Prefix service-set-name protocol-id protocol-name
source-interface-name separator source-address source-port destination-address destination-port
event-type;

Stateful Firewall No Policy Drop Logs

Following are example stateful firewall logs for MS-MPC services cards versus MX-SPC3 services processing card:

MS-MPC Services Card

JSERVICES_SFW_NO_POLICY source-address destination-address;

MX-SPC3 Services Card

RT_FLOW_SESSION_NO_POLICY_USF Prefix service-set-name source-address destination-address;

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

Configuring Syslog Events for NAT Rule Conditions with Next Gen Services

To configure syslog events to be generated when traffic matches NAT rule conditions for Next Gen Services NAT:

1. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

The following are logs collected:

Out of addresses logs — If the allocation request fails to be handled as the public IP addresses in the No-PAT pool are used up, the out of addresses syslog is generated.

Out of ports logs — If the allocation request fails to be handled as the public IPs and ports in the NAPT pool are used up, the out of ports syslog is generated.

NAT Rule Match Logs — If the packet matches the NAT rule, the NAT rule match syslog is generated.

Pool resource release logs — If the public IP and port succeeds to be released to the NAPT pool, the pool release syslog is generated.

RELATED DOCUMENTATION

[Network Address Port Translation \(NAPT\) Overview | 125](#)

[Configuring Network Address Port Translation for Next Gen Services | 126](#)

Next Gen Services SNMP MIBS and Traps

IN THIS CHAPTER

- [Next Gen Services SNMP MIBs and Traps | 88](#)

Next Gen Services SNMP MIBs and Traps

IN THIS SECTION

- [Service-Set Related SNMP MIBs | 88](#)
- [Summary Mapping of MX-SPC3 CLI Services Operational Commands to SNMP MIBs | 94](#)
- [NAT SNMP MIBs | 97](#)
- [SNMP Traps | 100](#)

This topic describes the SNMP MIBS and traps for Next Gen Services with the MX-SPC3 services. As a reference, it also compares MX-SPC3 services card MIBS and traps with the MPC services card.

Service-Set Related SNMP MIBs

[Table 27 on page 89](#), [Table 28 on page 90](#), and [Table 29 on page 91](#) describe the MIB objects in the service-set related SNMP MIB tables supported in **jnxSPMIB**. This MIB is supported for both MS-MPC services cards and MX-SPC3 services cards with the exception of the following:

- The MX-SPC3 services card supports counters, such as memory usage and cpu usage, at the per service-set and per pic level, whereas MS-MPC services cards support these counters at the service level, for example, stateful firewall (SFW) and NAT).

The MX-SPC3 card uses the **jnxSpSvcSetTable** MIB for these counters and the MS-MPC services card uses the **jnxSpSvcSetSvcTypeTable** MIB for these counters.

- In [Table 27 on page 89](#) the **jnxSpSvcSetTable**, the object **jnxSpSvcSetSvcType** field will show a value of “ALL” since no per service-type specific counters are supported.

Table 27: Service-Set SNMP MIB Table (jnxSpSvcSetTable)

MIB Object	jnxSpSvcSet Entry Number	Description
jnxSpSvcSetIfName	jnxSpSvcSetEntry 4	The name of the interface identifying the AS PIC. If more than one interface is associated with the AS PIC, the name associated with the lower layer interface is used.
jnxSpSvcSetIfIndex	jnxSpSvcSetEntry 5	An index number associated with the interface name.
jnxSpSvcSetMemoryUsage	jnxSpSvcSetEntry 6	Amount of memory used by the service set, in bytes.
jnxSpSvcSetCpuUtil	jnxSpSvcSetEntry 7	Amount of CPU processing used by the service set, expressed as a percentage of total CPU usage. J Series Services Routers do not have a dedicated CPU for services. CPU usage on these routers appears as 0.
jnxSpSvcSetSvcStyle	jnxSpSvcSetEntry 8	Type of service for the service set. Service types include: <ul style="list-style-type: none"> • Unknown—The service type is not known. • Interface-service—The service is interface based. • Next-hop-service—The service is next-hop based.
jnxSpSvcSetMemLimitPktDrops	jnxSpSvcSetEntry 9	Number of packets dropped because the service set exceeded its memory limits (operating in the Red zone).
jnxSpSvcSetCpuLimitPktDrops	jnxSpSvcSetEntry 10	Number of packets dropped because the service set exceeded the average CPU limits (when total CPU usage exceeds 85 percent).
jnxSpSvcSetFlowLimitPktDrops	jnxSpSvcSetEntry 11	Number of packets dropped because the service set exceeded the flow limit.

Table 27: Service-Set SNMP MIB Table (jnxSpSvcSetTable) (continued)

MIB Object	jnxSpSvcSet Entry Number	Description
jnxSpSvcSetMemoryUsage64		Amount of memory used by the service set, in bytes.
jnxSpSvcSetMemLimitPktDrops64		Number of packets dropped because the service set exceeded its memory limits (operating in the Red zone).
jnxSpSvcSetCpuLimitPktDrops64		Number of packets dropped because the service set exceeded the average CPU limits (when total CPU usage exceeds 85 percent).
jnxSpSvcSetFlowLimitPktDrops64		Number of packets dropped because the service set exceeded the flow limit.
jnxSpSvcSetSessCount		Number of valid sessions in the service-set.

Table 28: Service-Set Service Type SNMP MIB Table (jnxSpSvcSetSvcTypeTable)

MIB Object	(jnxSpSvcSetSvcType Entry Number	Description
jnxSpSvcSetSvcTypeIndex	jnxSpSvcSetSvcTypeEntry 1	An integer used to identify the service type.
jnxSpSvcSetSvcTypeIfName	jnxSpSvcSetSvcTypeEntry 2	The name of the interface identifying the AS PIC. If more than one interface is associated with the AS PIC, the name associated with the lower layer interface is used.
jnxSpSvcSetSvcTypeName	jnxSpSvcSetSvcTypeEntry 3	The name of the service type.
jnxSpSvcSetSvcTypeSvcSets	jnxSpSvcSetSvcTypeEntry 4	Number of service sets configured on the AS PIC that use this service type.
jnxSpSvcSetSvcTypeMemoryUsage	jnxSpSvcSetSvcTypeEntry 5	Amount of memory used by this service type, expressed in bytes.

Table 28: Service-Set Service Type SNMP MIB Table (jnxSpSvcSetSvcTypeTable) (continued)

MIB Object	(jnxSpSvcSetSvcType Entry Number	Description
jnxSpSvcSetSvcTypePctMemoryUsage	jnxSpSvcSetSvcTypeEntry 6	Amount of memory used by this service type, expressed as a percentage of total memory.
jnxSpSvcSetSvcTypeCpuUtil	jnxSpSvcSetSvcTypeEntry 7	Amount of CPU processing used by the service set, expressed as a percentage of total CPU usage. J Series Services Routers do not have a dedicated CPU for services. CPU usage on these routers appears as 0.

Table 29: Service-Set Interface SNMP MIB Table (jnxSpSvcSetIfTable)

MIB Object	jnxSpSvcSetIf Entry Number	Description
jnxSpSvcSetIfTableName	jnxSpSvcSetIfEntry 1	The name of the interface used to identify the AS PIC. If more than one interface is associated with the AS PIC, the name associated with the lower layer interface is used.
jnxSpSvcSetIfSvcSets	jnxSpSvcSetIfEntry 2	The number of service sets configured on the AS PIC.
jnxSpSvcSetIfMemoryUsage	jnxSpSvcSetIfEntry 3	Amount of memory used by the AS PIC, expressed in bytes.
jnxSpSvcSetIfPctMemoryUsage	jnxSpSvcSetIfEntry 4	Amount of memory used by the AS PIC, expressed as a percentage of total memory.
jnxSpSvcSetIfPolMemoryUsage	jnxSpSvcSetIfEntry 5	Amount of policy memory used by the AS PIC, expressed in bytes.

Table 29: Service-Set Interface SNMP MIB Table (jnxSpSvcSetIfTable) (continued)

MIB Object	jnxSpSvcSetIf Entry Number	Description
jnxSpSvcSetIfPctPolMemoryUsage	jnxSpSvcSetIfEntry 6	Amount of policy memory used by the AS PIC, expressed as a percentage of the total.
jnxSpSvcSetIfMemoryZone	jnxSpSvcSetIfEntry 7	<p>The memory usage zone currently occupied by the AS PIC. The definitions of each zone are:</p> <ul style="list-style-type: none"> • Green—All new flows are allowed. • Yellow—Unused memory is reclaimed. All new flows are allowed. • Orange—New flows are allowed only for service sets that use less than their equal share of memory. • Red—No new flows are allowed.
jnxSpSvcSetIfCpuUtil	jnxSpSvcSetIfEntry 8	<p>Amount of CPU processing used by the AS PIC, expressed as a percentage of total CPU usage.</p> <p>J Series Services Routers do not have a dedicated CPU for services. CPU usage on these routers appears as 0.</p>
jnxSpSvcSetIfMemoryUsage64		Amount of policy memory used by the AS PIC, expressed in bytes.
jnxSpSvcSetIfPolMemoryUsage64		Amount of policy memory used by the AS PIC, expressed as a percentage of the total.
jnxSpSvcSetIfNumTotalSessActive		Total number of active sessions in the PIC.
jnxSpSvcSetIfPeakTotalSessActive		Number of active sessions in the PIC at any time.
jnxSpSvcSetIfNumCreatedSessPerSec		Number of created sessions per second in the PIC

Table 29: Service-Set Interface SNMP MIB Table (jnxSpSvcSetIfTable) (continued)

MIB Object	jnxSpSvcSetIf Entry Number	Description
jnxSpSvcSetIfNumDeletedSessPerSec		Number of deleted sessions per second in the PIC
jnxSpSvcSetIfNumTotalTcpSessActive		Number of active sessions (TCP, UDP and other)in the PIC
jnxSpSvcSetIfNumTotalUdpSessActive		
jnxSpSvcSetIfNumTotalOtherSessActive		
jnxSpSvcSetIfPeakTotalTcpSessActive		Number of active sessions (TCP, UDP, and others) in the PIC
jnxSpSvcSetIfPeakTotalUdpSessActive		
jnxSpSvcSetIfPeakTotalOtherSessActive		
jnxSpSvcSetIfPeakCreatedSessPerSec		Number of created sessions per sec in the PIC
jnxSpSvcSetIfPeakDeletedSessPerSec		Number of deleted sessions per sec in the PIC
jnxSpSvcSetIfNumTotalTcpIpv4SessActive		Total number of active sessions (TCP, UDP and other) for IPv4 and IPv6 in the PIC
jnxSpSvcSetIfNumTotalTcpIpv6SessActive		
jnxSpSvcSetIfNumTotalUdpIpv4SessActive		
jnxSpSvcSetIfNumTotalUdpIpv6SessActive		
jnxSpSvcSetIfNumTotalOtherIpv4SessActive		
jnxSpSvcSetIfNumTotalOtherIpv6SessActive		
jnxSpSvcSetIfNumTotalTcpGatedSessActive		Number of TCP and UDP gated sessions in the PIC
jnxSpSvcSetIfNumTotalUdpGatedSessActive		
jnxSpSvcSetIfNumTotalTcpRegSessActive		Number of TCP and UDP regular sessions in the PIC
jnxSpSvcSetIfNumTotalUdpRegSessActive		
jnxSpSvcSetIfNumTotalTcpTunSessActive		Number of TCP and UDP tunneled sessions in the PIC
jnxSpSvcSetIfNumTotalUdpTunSessActive		

Table 29: Service-Set Interface SNMP MIB Table (jnxSpSvcSetIfTable) (continued)

MIB Object	jnxSpSvcSetIf Entry Number	Description
jnxSpSvcSetIfSessPktRecv		Number of packets received in session handling
jnxSpSvcSetIfSessPktXmit		Number of packets transmitted as a part of session handling
jnxSpSvcSetIfSessSlowPathDiscard		Number of packets discarded in slow path
jnxSpSvcSetIfSessSlowPathForward		Number of packets forwarded in slow path
jnxSpSvcSetIfMspNumCreatedSubsPerSec		Number of subscribers created per sec
jnxSpSvcSetIfMspNumDeletedSubsPerSec		Number of Subscribers deleted per sec
jnxSpSvcSetIfMspNumTotalSubsActive		Number of active subscribers
jnxSpSvcSetIfMspPeakCreatedSubsPerSec		Peak number of created subscribers per sec in the PIC
jnxSpSvcSetIfMspPeakDeletedSubsPerSec		Peak number of deleted subscribers per sec in the PIC
jnxSpSvcSetIfMspPeakTotalSubsActive		Peak number of total active subscribers in the PIC

Summary Mapping of MX-SPC3 CLI Services Operational Commands to SNMP MIBs

Table 30 on page 95 summarizes the mapping of the MX-SPC3 services card operations commands to the respective SNMP MIB.

Table 30: Summary Mapping of MX-SPC3 CLI Services Set Command to SNMP MIBs

CLI Command	Variable Name	MIB Tables	MIB Object
show services service-sets cpu-usage	cpu-utilization-percent	jnxSpSvcSetTable	jnxSpSvcSetCpuUtil
show services service-sets memory-usage	bytes-used		jnxSpSvcSetMemoryUsage64
show services service-sets memory-usage zone	mem-zone		jnxSpSvcSetIfMemoryZone
show services service-sets statistics packet-drops	cpulimit-drops		jnxSpSvcSetCpuLimitPktDrops64
	flowlimit-drops	jnxSpSvcSetFlowLimitPktDrops64	jnxSpSvcSetFlowLimitPktDrops64
	memlimit-drops		jnxSpSvcSetMemLimitPktDrops64
show services service-sets summary	service-set-bytes-used	jnxSpSvcSetIfTable	jnxSpSvcSetIfMemoryUsage64
	service-set-cpu-utilization		jnxSpSvcSetIfCpuUtil
	service-set-percent-bytes-used		jnxSpSvcSetIfPctMemoryUsage
	service-set-percent-poly-bytes-used		jnxSpSvcSetIfPctPolMemoryUsage
	service-set-policy-bytes-used		jnxSpSvcSetIfPolMemoryUsage64
	service-sets-configured		jnxSpSvcSetIfSvcSets
show services sessions count	sess-count	jnxSpSvcSetTable	jnxSpSvcSetSessCount

Table 30: Summary Mapping of MX-SPC3 CLI Services Set Command to SNMP MIBs (continued)

CLI Command	Variable Name	MIB Tables	MIB Object
show services sessions analysis	num-total-session-active	jnxSpSvcSetIfTable	jnxSpSvcSetIfNumTotalSessActive
	peak-total-session-active		jnxSpSvcSetIfPeakTotalSessActive
	num-created-session-per-sec		jnxSpSvcSetIfNumCreatedSessPerSec
	num-deleted-session-per-sec		jnxSpSvcSetIfNumDeletedSessPerSec
	num-total-tcp-session-active		jnxSpSvcSetIfNumTotalTcpSessActive
	num-total-udp-session-active		jnxSpSvcSetIfNumTotalUdpSessActive
	peak-total-tcp-session-active		jnxSpSvcSetIfPeakTotalTcpSessActive
	peak-total-udp-session-active		jnxSpSvcSetIfPeakTotalUdpSessActive
	num-total-other-session-active		jnxSpSvcSetIfNumTotalOtherSessActive
	peak-created-session-per-second		jnxSpSvcSetIfPeakCreatedSessPerSec
	peak-deleted-session-per-second		jnxSpSvcSetIfPeakDeletedSessPerSec
	peak-total-other-session-active		jnxSpSvcSetIfPeakTotalOtherSessActive
	num-total-tcp-ipv4-session-active		jnxSpSvcSetIfNumTotalTcpIpv4SessActive
	num-total-tcp-ipv6-session-active		jnxSpSvcSetIfNumTotalTcpIpv6SessActive
	num-total-udp-ipv4-session-active		jnxSpSvcSetIfNumTotalUdpIpv4SessActive
	num-total-udp-ipv6-session-active		jnxSpSvcSetIfNumTotalUdpIpv6SessActive
	num-total-tcp-gated-session-active		jnxSpSvcSetIfNumTotalTcpGatedSessActive
	num-total-udp-gated-session-active		jnxSpSvcSetIfNumTotalUdpGatedSessActive
	num-total-other-ipv4-session-active		jnxSpSvcSetIfNumTotalOtherIpv4SessActive
	num-total-other-ipv6-session-active		jnxSpSvcSetIfNumTotalOtherIpv6SessActive
	num-total-tcp-reg-ipv4-session-active		jnxSpSvcSetIfNumTotalTcpRegIpv4SessActive

Table 30: Summary Mapping of MX-SPC3 CLI Services Set Command to SNMP MIBs (continued)

CLI Command	Variable Name	MIB Tables	MIB Object
	num-total-up-reg-session-active	jnxSpSvcSetIfTable	jnxSpSvcSetIfNumTotalUpRegSessActive
	num-total-tun-up-reg-session-active		jnxSpSvcSetIfNumTotalTupTunSessActive
	num-total-up-tun-reg-session-active		jnxSpSvcSetIfNumTotalUdpTunSessActive
	session-pkts-received		jnxSpSvcSetIfSessPktRecv
	session-pkts-transmitted		jnxSpSvcSetIfSessPktXmit
	session-slow-path-discard		jnxSpSvcSetIfSessSlowPathDiscard
	session-slow-path-forward		jnxSpSvcSetIfSessSlowPathForward
show services subscriber analysis	misp-num-created-sub-per-sec		jnxSpSvcSetIfMspNumCreatedSubsPerSec
	misp-num-deleted-sub-per-sec		jnxSpSvcSetIfMspNumDeletedSubsPerSec
	misp-num-total-sub-active		jnxSpSvcSetIfMspNumTotalSubsActive
	misp-peak-created-sub-per-second		jnxSpSvcSetIfMspPeakCreatedSubsPerSec
	misp-peak-deleted-sub-per-second		jnxSpSvcSetIfMspPeakDeletedSubsPerSec
	misp-peak-total-sub-active		jnxSpSvcSetIfMspPeakTotalSubsActive

NAT SNMP MIBs

This section describes the **jnxSrcNatStatsTable** MIB objects.

[Table 31 on page 98](#) describes the source NAT SNMP MIB objects for the MS-MPC services card. This table exposes the source NAT translation attributes of the translated addresses.

[Table 32 on page 99](#) describes the source NAT SNMP MIB objects for the MX-SPC3 services card. This table contains information on source IP address translation only.

Table 31: MS-MPC Services Card Source NAT SNMP MIB Table (jnxSrcNatStatsTable)

jnxSrcNatStatsTable	MIB Object	Description
	jnxNatSrcPoolName	The name of dynamic source IP address pool
	jnxNatSrcXlatedAddrType	V4 or V6. The type of dynamic source IP address allocated from the address pool used in the NAT translation
	jnxNatSrcPoolType	The source port pool type indicates whether the address translation is done with port or without the port, or if it is a static translation. Ex napt-44, nat64 etc
	jnxNatSrcNumPortAvail	The number of ports available with this pool
	jnxNatSrcNumPortInuse	The number of ports in use for this NAT address entry
	jnxNatSrcNumAddressAvail	The total number of addresses available in this pool
	jnxNatSrcNumAddressInUse	The number of addresses in use from this pool
	jnxNatSrcNumSessions	The number of sessions are in use based on this NAT address entry
jnxNatRuleTable		This table monitors NAT rule hits
	jnxNatRuleName	NAT rule name
	jnxNatRuleType	NAT types: Static Source, Static Destination, Dynamic Source and NAPT. Ex: napt44 etc
	jnxNatRuleTransHits	The number of hits on this NAT rule
jnxNatPoolTable		This table monitors NAT pool hits
	jnxNatPoolName	NAT Pool name

Table 31: MS-MPC Services Card Source NAT SNMP MIB Table (jnxSrcNatStatsTable) (continued)

jnxSrcNatStatsTable	MIB Object	Description
	jnxNatPoolType	NAT types: Static Source, Static Destination, Dynamic Source and NAPT. Ex: napt44 etc
	jnxNatPoolTransHits	The number of hits on this NAT Pool

Table 32: MX-SPC3 Source NAT SNMP MIB Table (jnxNatObjects)

jnxJsSrcNatStatsTable	MIB Object	Description
	jnxJsNatSrcPoolName	The name of dynamic source IP address pool
	jnxJsNatSrcXlatedAddrType	New MIB. The type of dynamic source IP address allocated from the address pool used in the NAT translation. Value is v4 or v6
	jnxJsNatSrcPoolType	withPAT or withoutPAT or static
	jnxJsNatSrcNumPortAvail	New MIB. The number of ports available with this pool
	jnxJsNatSrcNumPortInuse	The number of ports in use for this NAT address entry
	jnxJsNatSrcNumSessions	The number of sessions are in use based on this NAT address entry
	jnxJsNatSrcNumAddressAvail	New MIB. The total number of addresses available in this pool
	jnxJsNatSrcNumAddressInuse	New MIB. The number of addresses in use from this pool
jnxJsNatRuleTable		This table monitors NAT rule hits
	jnxJsNatRuleName	NAT rule name
	jnxJsNatRuleType	NAT types: Source, Destination and Static
	jnxJsNatRuleTransHits	The number of hits on this NAT rule. Status is deprecated. New - jnxJsNatRuleHits

Table 32: MX-SPC3 Source NAT SNMP MIB Table (jnxNatObjects) (continued)

jnxJsSrcNatStatsTable	MIB Object	Description
	jnxJsNatRuleHits	The number of hits on this NAT rule,
	jnxJsNatRuleNumOfSessions	The number of sessions on this NAT rule
	jnxJsNatTransType	New MIB. Details below
jnxJsNatPoolTable		This table monitors NAT pool hits
	jnxJsNatPoolName	NAT Pool name
	jnxJsNatPoolType	NAT types: Source, Destination and Static
	jnxJsNatPoolTransHits	The number of hits on this NAT pool. Status is deprecated. New - jnxJsNatPoolHits
	jnxJsNatPoolHits	The number of hits on this NAT pool to deprecate jnxJsNatRuleTransHits.

SNMP Traps

Table 33 on page 100 describes the SNMP traps supported by both the MS-MPC services card and the MX-SPC3 services card.

Table 33: SNMP Traps

Trap	Description
SPD_TRAP_OIDS(jnxSpSvcSetZoneEntered)	jnxSpSvcSetZoneEntered — Indicates that an AS PIC has entered a more severe memory usage zone from a less severe memory usage zone. The zone entered is identified by JnxSpSvcSetIfMemoryZone
SPD_TRAP_OIDS(jnxSpSvcSetZoneExited)	jnxSpSvcSetZoneExited — Indicates that an AS PIC has exited a more severe memory usage zone to a less severe memory usage zone. The zone entered is identified by JnxSpSvcSetIfMemoryZone.
SPD_TRAP_OIDS(jnxSpSvcSetCpuExceeded)	jnxSpSvcSetCpuExceeded — Indicates that an AS PIC has over 85% CPU usage.

Table 33: SNMP Traps (*continued*)

Trap	Description
SPD_TRAP_OIDS(jnxSpSvcSetCpuOk)	jnxSpSvcSetCpuOk — Indicates that an AS PIC has returned to less than 85%CPU usage.
SPD_TRAP_OIDS(jnxSpSvcSetFlowLimitUtilized)	jnxSpSvcSetFlowLimitUtilized — Indicates a service-set has reached its upper limit of flows threshold of a maximum flows allowed for a service set.

Configuring SNMP Trap Generation

This sections describes how to configure the MS-MPC service card versus the MX-SPC3 services card to generate SNMP traps.

Configuring SNMP Trap for NAT Ports in a Source NAT Pool

If the current usage is above the raise threshold or below the clear threshold, we will generate a SNMP trap.

Configuring SNMP Traps for NAT Ports in a Source NAT Pool on an MS-MPC

```
user@host# set services nat pool NAT_POOL_TEST snmp-trap-thresholds address-port low 50
user@host# set services nat pool NAT_POOL_TEST snmp-trap-thresholds address-port high 75
```

Configuring SNMP Traps for NAT Ports in a Source NAT Pool on an MX-SPC3

```
user@host# set services nat source pool NAT_POOL_TEST pool-utilization-alarm raise-threshold 50
user@host# set services nat source pool NAT_POOL_TEST pool-utilization-alarm clear-threshold 40
```

Configuring SNMP Trap for Sessions

This is infra trap which configures SNMP flow thresholds for all flows for a service set or flows for all NAT pools configured for a service set.

Configuring a Sessions SNMP Trap on an MS-MPC

```
user@host# set services service-set SS_TEST max-flows 2m
user@host# set services service-set SS_TEST snmp-trap-thresholds flow low 50
user@host# set services service-set SS_TEST snmp-trap-thresholds flow high 75
```

Configuring a Sessions SNMP Trap on an MX-SPC3

```
user@host# set services service-set ss1 service-set-options session-limit maximum 2000
user@host# set services service-set ss1 snmp-trap-thresholds session low 50
```

```
user@host# set services service-set ss1 snmp-trap-thresholds session high 60
```

Example-Configuration for MX-SPC3 NAT for Three SNMP MIB Tables

Example Configuration

```
user@host> show services | display set
```

```
Configuration
=====
show services | display set

set services service-set ss1_nh_style1 nat-rule-sets rset1
set services service-set ss1_nh_style1 nat-rule-sets rset2
set services service-set ss1_nh_style1 nat-rule-sets rset5
set services service-set ss1_nh_style1 next-hop-service inside-service-interface
vms-0/0/0.1
set services service-set ss1_nh_style1 next-hop-service outside-service-interface
vms-0/0/0.2
set services nat source pool src_pool2_v6 address 300::0/128
set services nat source pool src_pool1 address 50.0.0.0/29
set services nat source rule-set rset1 rule nr1 match source-address 10.0.0.0/32
set services nat source rule-set rset1 rule nr1 match destination-address
20.0.0.0/32
set services nat source rule-set rset1 rule nr1 match application any
set services nat source rule-set rset1 rule nr1 then source-nat pool src_pool1
set services nat source rule-set rset1 match-direction input
set services nat source rule-set rset2 rule nr2_v6 match source-address 200::0/34
set services nat source rule-set rset2 rule nr2_v6 match destination-address
400::0/34
set services nat source rule-set rset2 rule nr2_v6 match application any
set services nat source rule-set rset2 rule nr2_v6 then source-nat pool src_pool2_v6
set services nat source rule-set rset2 match-direction input
set services nat destination pool src_pool5_dnat address 20.0.0.0/30
set services nat destination rule-set rset5 rule nr5_dnat match destination-address
21.0.0.0/30
set services nat destination rule-set rset5 rule nr5_dnat match application any
set services nat destination rule-set rset5 rule nr5_dnat then destination-nat
pool src_pool5_dnat
set services nat destination rule-set rset5 match-direction input
set services nat traceoptions file nat-trace.txt
set services nat traceoptions flag all
```


show snmp mib walk jnxJsSrcNatStatsTable

```
user@host>show snmp mib walk jnxJsSrcNatStatsTable
```

```
jnxJsNatSrcPoolName.2.112.49.0.0.0.0.0 = p1
jnxJsNatSrcXlatedAddrType.2.112.49.0.0.0.0.0 = 1
jnxJsNatSrcPoolType.2.112.49.0.0.0.0.0 = 1
jnxJsNatSrcNumPortInuse.2.112.49.0.0.0.0.0 = 0
jnxJsNatSrcNumSessions.2.112.49.0.0.0.0.0 = 0
jnxJsNatSrcNumPortAvail.2.112.49.0.0.0.0.0 = 10
jnxJsNatSrcNumAddressAvail.2.112.49.0.0.0.0.0 = 1
jnxJsNatSrcNumAddressInuse.2.112.49.0.0.0.0.0 = 0
```

show snmp mib walk jnxJsNatPoolTable

```
user@host>show snmp mib walk jnxJsNatPoolTable
```

```
jnxJsNatPoolName.9.115.114.99.95.112.111.111.108.49.1 = src_pool1
jnxJsNatPoolName.14.115.114.99.95.112.111.111.108.53.95.100.110.97.116.2 =
src_pool5_dnat
jnxJsNatPoolType.9.115.114.99.95.112.111.111.108.49.1 = 1
jnxJsNatPoolType.14.115.114.99.95.112.111.111.108.53.95.100.110.97.116.2 = 2
jnxJsNatPoolTransHits.9.115.114.99.95.112.111.111.108.49.1 = 0
jnxJsNatPoolTransHits.14.115.114.99.95.112.111.111.108.53.95.100.110.97.116.2 = 0
jnxJsNatPoolHits.9.115.114.99.95.112.111.111.108.49.1 = 0
jnxJsNatPoolHits.14.115.114.99.95.112.111.111.108.53.95.100.110.97.116.2 = 0
jnxJsNatPoolUtil.9.115.114.99.95.112.111.111.108.49.1 = 0
jnxJsNatPoolUtil.14.115.114.99.95.112.111.111.108.53.95.100.110.97.116.2 = 0
```

show snmp mib walk jnxJsNatRuleTable

```
user@host>show snmp mib walk jnxJsNatRuleTable
```

```
jnxJsNatRuleName.3.110.114.49.1 = nr1
jnxJsNatRuleName.6.110.114.50.95.118.54.1 = nr2_v6
jnxJsNatRuleName.8.110.114.53.95.100.110.97.116.2 = nr5_dnat
jnxJsNatRuleType.3.110.114.49.1 = 1
jnxJsNatRuleType.6.110.114.50.95.118.54.1 = 1
jnxJsNatRuleType.8.110.114.53.95.100.110.97.116.2 = 2
jnxJsNatRuleTransHits.3.110.114.49.1 = 0
jnxJsNatRuleTransHits.6.110.114.50.95.118.54.1 = 0
jnxJsNatRuleTransHits.8.110.114.53.95.100.110.97.116.2 = 0
jnxJsNatRuleHits.3.110.114.49.1 = 0
jnxJsNatRuleHits.6.110.114.50.95.118.54.1 = 0
```

This section provides some example trace logs for these SNMP traps.

[illegible]

[SNMP MIB Explorer](#)

[Explore System Log Messages](#)

2

PART

Carrier Grade NAT (CGNAT)

Deterministic NAT Overview and Configuration | **109**

Dynamic Address-Only Source NAT Overview and Configuration | **120**

Network Address Port Translation Overview and Configuration | **125**

NAT46 | **134**

Stateful NAT64 Overview and Configuration | **138**

IPv4 Connectivity Across IPv6-Only Network Using 464XLAT Overview and Configuration | **148**

IPv6 NAT Protocol Translation (NAT PT) | **158**

Stateless Source Network Prefix Translation for IPv6 Overview and Configuration | **161**

Transitioning to IPv6 Using Softwires | **165**

Transitioning to IPv6 Using DS-Lite Softwires | **171**

Reducing Traffic and Bandwidth Requirements Using Port Control Protocol | **185**

Transitioning to IPv6 Using Mapping of Address and Port with Encapsulation (MAP-E) | **195**

Monitoring and Troubleshooting Softwires | **207**

Port Forwarding Overview and Configuration | **211**

Port Translation Features Overview and Configuration | **219**

Static Source NAT Overview and Configuration | **223**

Static Destination NAT Overview and Configuration | **228**

Twice NAT Overview and Configuration | **233**

Twice NAT Overview and Configuration | **243**

Class of Service Overview and Configuration | **255**

Deterministic NAT Overview and Configuration

IN THIS CHAPTER

- Deterministic NAPT Overview for Next Gen Services | 109
- Configuring Deterministic NAPT for Next Gen Services | 114

Deterministic NAPT Overview for Next Gen Services

Under Next Gen Services with the MX-SPC3, you can configure both Deterministic NAPT44 and NAPT64 services. Next Gen Services deterministic NAPT services use an algorithm to allocate blocks of destination ports.

Next Gen Services deterministic NAPT44 service ensures that the original source IPv4 address and port always map to the same post-NAT IPv4 address and port range, and that the reverse mapping of a given translated external IPv4 address and port are always mapped to the same internal IPv4 address.

Next Gen Services deterministic NAPT64 service ensures that the original source IPv6 address and port always map to the same post-NAT IPv4 address and port range, and that the reverse mapping of a given translated external IPv4 address and port are always mapped to the same internal IPv6 address.

For detailed information on how to configure deterministic NAPT, see [“Configuring Deterministic NAPT for Next Gen Services” on page 114](#).

Benefits of Deterministic NAPT

- Eliminates the need for address translation logging because an IP address is always mapped to the same external IP address and port range, and the reverse mapping of a given translated external IP address and port are always mapped to the same internal IP address.

Understanding Deterministic NAPT Algorithms

The effectiveness of your implementation of deterministic NAPT depends on your analysis of your subscriber requirements. The block size you provide indicates how many ports will be made available for each incoming subscriber address from the range in the **from** clause specified in the applicable NAT rule. The allocation

algorithm computes an offset value to determine the outgoing IP address and port. A reverse algorithm is used to derive the originating subscriber address.

NOTE: In order to track subscribers without using logs, an ISP must use a reverse algorithm to derive a subscriber (source) addresses from a translated address.

The following variables are used in forward calculation (private subscriber IP address to public IP address) and reverse calculation (public IP address to private subscriber IP address):

- **Pr_Prefix**—Any pre-NAT IPv4 subscriber address.
- **Pr_Port**—Any pre-NAT protocol port.
- **Block_Size**—Number of ports configured to be available for each **Pr_Prefix**.
If **block-size** is configured as zero, the method for computing the block size is computed as follows:

$$\text{block-size} = \text{int}(64512 / \text{ceil}[(\text{Nr_Addr_PR_Prefix} / \text{Nr_Addr_PU_Prefix})])$$
 where 64512 is the maximum available port range per public IP address.
- **Base_PR_Prefix**—First usable pre-NAT IPv4 subscriber address in a **from** clause of the NAT rule.
- **Base_PU_Prefix**—First usable post-NAT IPv4 subscriber address configured in the NAT pool.
- **Pu_Port_Range_Start**—First usable post-NAT port. This is 1024.
- **Pr_Offset**—The offset of the pre-NAT IP address that is being translated from the first usable pre-NAT IPv4 subscriber address in a **from** clause of the NAT rule. $\text{Pr_Offset} = \text{Pr_Prefix} - \text{Base_Pr_Prefix}$.
- **PR_Port_Offset**—Offset of the pre-NAT IP address multiplied by the block size. $\text{PR_Port_Offset} = \text{Pr_Offset} * \text{Block_Size}$.
- **Pu_Prefix**—Post-NAT address for a given **Pr_Prefix**.
- **Pu_Start_Port**—Post-NAT start port for a flow from a given **Pr_Prefix**
- **Pu_Actual_Port**—Post-NAT port seen on a reverse flow.
- **Nr_Addr_PR_Prefix** — Number of usable pre-NAT IPv4 subscriber addresses in a **from** clause of the NAT rule.
- **Nr_Addr_PU_Prefix** — Number of usable post-NAT IPv4 addresses configured in the NAT pool.
- **Rounded_Port_Range_Per_IP** — Number of ports available for each post-NAT IP address.

$$\text{Rounded_Port_Range_Per_IP} = \text{ceil}[(\text{Nr_Addr_PR_Prefix} / \text{Nr_Addr_PU_Prefix})] * \text{Block_Size}$$

- **Pu_Offset**—Offset of the post-NAT IP address from the first usable post-NAT address. $\text{Pu_Offset} = \text{Pu_Prefix} - \text{Base_Pu_Prefix}$.
- **Pu_Port_Offset**— Offset of the post-NAT port from 1024 added to the product of the offset of the post-NAT IP address and the number of ports available for each post-NAT IP address. $\text{Pu_Port_Offset} = (\text{Pu_Offset} * \text{Rounded_Port_Range_Per_IP}) + (\text{Pu_Actual_Port} - \text{Pu_Port_Range_Start})$.

Algorithm Usage—Assume the following configurations:

```

services {
  nat {
    source {
      pool src-pool {
        address 203.0.113.0/16;
        port {
          automatic {
            random-allocation;
          }
          deterministic {
            block-size 249;
            host address 10.1.0.1/16;
          }
        }
      }
    }
    rule-set set1 {
      rule det-nat {
        match-direction-input;
        match {
          source-address 10.1.0.0/16;
        }
        then {
          source-nat {
            pool src-pool;
          }
        }
      }
    }
  }
}

```

Forward Translation

1. $\text{Pr_Offset} = \text{Pr_Prefix} - \text{Base_Pr_Prefix}$
2. $\text{Pr_Port_Offset} = \text{Pr_Offset} * \text{Block_Size}$

3. $\text{Rounded_Port_Range_Per_IP} = \text{ceil}[(\text{Nr_Addr_PR_Prefix} / \text{Nr_Addr_PU_Prefix})] * \text{Block_Size}$
4. $\text{Pu_Prefix} = \text{Base_Public_Prefix} + \text{floor}(\text{Pr_Port_Offset} / \text{Rounded_Port_Range_Per_IP})$
5. $\text{Pu_Start_Port} = \text{Pu_Port_Range_Start} + (\text{Pr_Port_Offset} \% \text{Rounded_Port_Range_Per_IP})$

Using the sample configuration and assuming a subscriber flow sourced from 10.1.1.250:5000:

1. $\text{Pr_Offset} = 10.1.1.250 - 10.1.0.1 = 505$
2. $\text{Pr_Port_Offset} = 505 * 249 = 125,745$
3. $\text{Rounded_Port_Range_Per_IP} = \text{ceil}[(65, 533 / 254)] * 249 = 259 * 249 = 64,491$
4. $\text{Pu_Prefix} = 203.0.113.1 + \text{floor}(125,745 / 64,491) = 203.0.113.1 + 1 = 203.0.113.2$
5. $\text{Pu_Start_Port} = 1,024 + (125,745 \% 64,491) = 62278$
 - 10.1.1.250 is translated to 203.0.113.2.
 - The starting port is 62278. There are 249 ports available to the subscriber based on the configured block size. The available port range spans ports 62278 through 62526 (inclusive).
 - The specific flow 10.1.1.250:5000 randomly assigns any of the ports in its range because random allocation was specified.

Reverse Translation

1. $\text{Pu_Offset} = \text{Pu_Prefix} - \text{Base_Pu_Prefix}$
2. $\text{Pu_Port_Offset} = (\text{Pu_Offset} * \text{Rounded_Port_Range_Per_IP}) + (\text{Pu_Actual_Port} - \text{Pu_Port_Range_Start})$
3. $\text{Subscriber_IP} = \text{Base_Pr_Prefix} + \text{floor}(\text{Pu_Port_Offset} / \text{Block_Size})$

The reverse translation is determined as follows. Assume a flow returning to 203.0.113.2:62278.

1. $\text{Pu_Offset} = 203.0.113.2 - 203.0.113.1 = 1$
2. $\text{Pu_Port_Offset} = (1 * 64,491) + (62,280 - 1024) = 125,747$
3. $\text{Subscriber_IP} = 10.1.0.1 + \text{floor}(125,747 / 249) = 10.1.0.1 + 505 = 10.1.1.250$

NOTE: In reverse translation, only the original private IP address can be derived, and not the original port in use. This is sufficiently granular for law enforcement requirements.

When you have configured deterministic NAT, you can use the **show services nat deterministic-nat internal-host** and **show services nat deterministic-nat nat-port-block** commands to show forward and reverse mapping. However, mappings will change if you reconfigure your deterministic port block allocation block size or the **from** clause for your NAT rule. In order to provide historical information on mappings, we recommend that you write scripts that can show specific mappings for prior configurations.

Deterministic NAT Restrictions

When you configure deterministic NAT, be aware of the following:

- For IPv6 deterministic NAT64 host address configuration, we support the last 32-bit (4 byte) change of the IPv6 host prefix. This means we only can configure /96 prefix masks for IPv6 address, which supports a maximum address number of 2^{32} for one IPv6 prefix. The host address is specified at the **[services nat source pool p1 port deterministic host]** configuration hierarchy.

- Usually, the number of address in host-range should be more than the number of address in pool.

-

BEST PRACTICE: We don't recommend the host address number be configured to exceed the total port block resource number because some hosts may not receive a port block resource successfully.

- The minimum block size for deterministic NAT is 1. If you configure a smaller block size, the commit fails. If the block size is configured to 0, the block size will be automatically calculated based on host number and translated address number. If the calculated block size is less than 1, the commit fails.
- For Next Gen Services deterministic NAT, you can configure a mix of IPv4 and IPv6 host addresses together in a NAT pool in either a host address or an address name list, However. the total host prefix number cannot exceed 1000.
- You cannot configure an address range or DNS name in a host address book name.
- The configured host address prefix and host address book name are merged together if its prefixes are overlapped. You can use the **show services nat source deterministic** operational command to show the merged prefixes.

- BEST PRACTICE:** We recommend, you keep subscriber host addresses consistent with multiple rule's matching the source address prefix, if the same deterministic NAT pool is used across multiple rules; otherwise, traffic from hosts which are not configured in the NAT pool, even it matches the NAT rule, may not allocate the port successfully.
- For Next Gen Services NAPT services, the total number of host addresses configured must be greater than or equal to the deterministic NAT port blocks available.

RELATED DOCUMENTATION

[Configuring Deterministic NAPT for Next Gen Services | 114](#)

Configuring Deterministic NAPT for Next Gen Services

IN THIS SECTION

- [Configuring the NAT Pool for Deterministic NAPT for Next Gen Services | 114](#)
- [Configuring the NAT Rule for Deterministic NAPT44 for Next Gen Services | 116](#)
- [Configuring the NAT Rule for Deterministic NAPT64 for Next Gen Services | 117](#)
- [Configuring the Service Set for Deterministic NAT for Next Gen Services | 118](#)
- [Clearing the Don't Fragment Bit | 119](#)

To configure deterministic NAPT on Next Gen Services, perform the following:

Configuring the NAT Pool for Deterministic NAPT for Next Gen Services

To configure the NAT pool for deterministic NAPT:

1. Create a pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. Configure deterministic port block allocation for the pool.

```
[edit services nat source pool nat-pool-name port]  
user@host# set deterministic
```

4. If you want the lowest and highest IPv4 addresses (the network and broadcast addresses) in the source address range of a NAT rule to be translated when the NAT pool is used, configure **include-boundary-address**.

```
[edit services nat source pool nat-pool-name port deterministic]  
user@host# set include-boundary-addresses
```

5. Configure the port block size. The range is 1 to 64,512. The default block size is 256.

```
[edit services nat source pool nat-pool-name port deterministic]  
user@host# set block-size block-size
```

6. Configure the first usable pre-NAT subscriber address, which is used in calculating the offset value for a pre-NAT address that is being translated. This offset is used to perform the deterministic NAT mapping.

```
[edit services nat source pool nat-pool-name port deterministic]  
user@host# set host address host-addr
```

7. Configure the interval at which the syslog is generated for the deterministic NAT configuration.

```
[edit services nat source pool nat-pool-name port deterministic]  
user@host# set deterministic-nat-configuration-log-interval seconds
```

8. To configure automatic port assignment for the pool, specify either random allocation or round-robin allocation.

```
[edit services nat source pool nat-pool-name port]
user@host# set automatic (random-allocation | round-robin)
```

Random allocation randomly assigns a port from the range 1024 through 65535 for each port translation. Round robin allocation first assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

9. To disable round-robin port allocation for all NAT pools that do not specify an **automatic (random-allocation | round-robin)** setting, configure the global setting.

```
[edit services nat source]
user@host# set port-round-robin disable
```

SEE ALSO

[Network Address Translation Configuration Overview](#)

Configuring the NAT Rule for Deterministic NAPT44 for Next Gen Services

To configure the NAT rule for deterministic NAPT44:

1. Configure the NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address any-unicast
```

- Specify one or more application protocols to which the NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

- Specify the NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

Configuring the NAT Rule for Deterministic NAPT64 for Next Gen Services

To configure the NAT rule for deterministic NAPT64:

- Configure the source NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

- Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
```

```
user@host# set match-direction (in | out | in-out)
```

3. Specify the IPv6 prefix for the source addresses that are translated by the NAT rule.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

4. Specify one or more application protocols to which the NAT rule applies. The number of application terms must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

5. Specify the NAT source pool that contains the addresses for translated source addresses.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

Configuring the Service Set for Deterministic NAT for Next Gen Services

To configure the service set for deterministic NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```


3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]  
user@host# set nat-rule-sets rule-set-name
```

Clearing the Don't Fragment Bit

If you configured deterministic NAPT64, specify that the don't fragment (DF) bit for IPv4 packet headers is cleared when the packet length is less than 1280 bytes.

```
[edit services nat natv6v4]  
user@host# set clear-dont-fragment-bit
```

This prevents unnecessary creation of an IPv6 fragmentation header when translating IPv4 packets that are less than 1280 bytes.

RELATED DOCUMENTATION

| [Deterministic NAPT Overview for Next Gen Services](#) | 109

Dynamic Address-Only Source NAT Overview and Configuration

IN THIS CHAPTER

- [Dynamic Address-Only Source Translation Overview | 120](#)
- [Configuring Dynamic Address-Only Source NAT for Next Gen Services | 121](#)

Dynamic Address-Only Source Translation Overview

With dynamic address-only translation, you can map a private IP source address to a public IP address. A public address is picked up dynamically from a source NAT pool, and the mapping from the original source address to the translated source address is maintained as long as there is at least one active flow that uses this mapping. The port is not mapped.

Benefits of Dynamic Address-Only Source Translation

- Allows hosts in the private network to connect with the external domain, while hiding the private network.
- Allows a few public IP addresses to be used by several private hosts

RELATED DOCUMENTATION

[Configuring Dynamic Address-Only Source NAT for Next Gen Services | 121](#)

Configuring Dynamic Address-Only Source NAT for Next Gen Services

IN THIS SECTION

- [Configuring the Source Pool for Dynamic Address-Only Source NAT | 121](#)
- [Configuring the NAT Source Rule for Dynamic Address-Only Source NAT | 122](#)
- [Configuring the Service Set for Dynamic Address-Only Source NAT | 124](#)

Configuring the Source Pool for Dynamic Address-Only Source NAT

To configure the source pool for dynamic address-only source NAT:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. Disable port translation.

```
[edit services nat source pool nat-pool-name]  
user@host# set port no-translation
```

4. Define the NAT pool utilization levels that trigger SNMP traps. The **raise-threshold** is the pool utilization percentage that triggers the trap, and the range is 50 through 100. The **clear-threshold** is the pool utilization percentage that clears the trap, and the range is 40 through 100. The utilization is based on the number of addresses that are used.

```
[edit services nat source pool nat-pool-name]
user@host# set pool-utilization-alarm raise-threshold value
user@host# set pool-utilization-alarm clear-threshold value
```

If you do not configure **pool-utilization-alarm**, traps are not created.

5. To allow the IP addresses of a NAT source pool or destination pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Source Rule for Dynamic Address-Only Source NAT

To configure the NAT source rule for dynamic address-only source NAT:

1. Configure the NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address any-unicast
```

- Specify one or more application protocols to which the NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

- Specify the NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

- Configure the address-pooling paired feature if you want to ensure assignment of the same external IP address for all sessions originating from the same internal host.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat mapping-type]
user@host# set address-pooling-paired
```

- Specify the timeout period for **address-pooling-paired** mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Dynamic Address-Only Source NAT

To configure the service set for dynamic address-only source NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

RELATED DOCUMENTATION

[Dynamic Address-Only Source Translation Overview](#) | 120

Network Address Port Translation Overview and Configuration

IN THIS CHAPTER

- [Network Address Port Translation \(NAPT\) Overview | 125](#)
- [Configuring Network Address Port Translation for Next Gen Services | 126](#)
- [Configuring Syslog Events for NAT Rule Conditions with Next Gen Services | 133](#)

Network Address Port Translation (NAPT) Overview

IN THIS SECTION

- [Benefits of NAPT | 126](#)

NAPT translates a private source IP address to an external source address and port. Multiple private IP addresses can be mapped to the same external address because each private address is mapped to a different port of the external address.

With NAPT, you can configure up to 32 external address ranges, and map up to 65,536 private addresses to each external address.

NAPT supports the following:

- Round-robin port and address allocation (see [“Round-Robin Port Allocation” on page 221](#)).
- Address pooling and endpoint independent mapping (see [“Address Pooling and Endpoint Independent Mapping for Port Translation” on page 219](#)).
- Secured port block allocation (see [“Secured Port Block Allocation for Port Translation” on page 222](#)).

Benefits of NAPT

- Allows hosts in the private network to connect with the external domain, while hiding the private network.
- Minimizes the number of public IP addresses that are allocated for NAT.

Configuring Network Address Port Translation for Next Gen Services

1. [Configuring the Source Pool for NAPT | 126](#)
2. [Configuring the NAT Source Rule for NAPT | 130](#)
3. [Configuring the Service Set for NAPT | 132](#)

Configuring the Source Pool for NAPT

To configure the source pool for NAPT:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]
user@host# set address address-prefix to address address-prefix
```

3. To configure automatic port assignment for the pool, specify either random allocation or round-robin allocation. Round-robin allocation is the default.

```
[edit services nat source pool nat-pool-name port]
user@host# set automatic (random-allocation | round-robin)
```

Random allocation randomly assigns a port from the range 1024 through 65535 for each port translation. Round-robin allocation first assigns port 1024, and uses the next higher port for each successive port assignment.

4. To disable round-robin port allocation for all NAT pools that do not specify an **automatic** (**random-allocation** | **round-robin**) setting, configure the global setting.

```
[edit services nat source]
user@host# set port-round-robin disable
```

5. To configure a range of ports to assign to a pool, perform the following:

NOTE: If you specify a range of ports to assign, the **automatic** statement is ignored.

- a. Specify the low and high values for the port. If you do not configure automatic port assignment, you must configure a range of ports.

```
[edit services nat source pool nat-pool-name port]
user@host# set range port-low to port-high
```

- b. Specify either random allocation or round-robin allocation. Round-robin allocation is the default.

```
[edit services nat source pool nat-pool-name port range]
user@host# set (random-allocation | round-robin)
```

6. Assign a port within the same range as the incoming port—either 0 through 1023 or 1024 through 65,535. This feature is not available if you configure port-block allocation.

```
[edit services nat source pool nat-pool-name port]
user@host# set preserve-range
```

7. Assign a port with the same parity (even or odd) as the incoming source port. This feature is not available if you configure port-block allocation.

```
[edit services nat source pool nat-pool-name port]
user@host# set preserve-parity
```

8. Configure a global default port range for NAT pools that use port translation. This port range is used when a NAT pool does not specify a port range and does not specify automatic port assignment. The global port range can be from 1024 through 65,535.

```
[edit services nat source]
user@host# set pool-default-port-range port-low to port-high
```

9. If you want to allocate a block of ports for each subscriber to use for NAT, configure port-block allocation:

- a. Configure the number of ports in a block. The range is 1 through 64,512 and the default is 128.

```
[edit services nat source pool nat-pool-name port]
user@host# set block-allocation block-size block-size
```

- b. Configure the interval, in seconds, for which the block is active. After the timeout, a new block is allocated, even if ports are available in the active block. If you set the timeout to 0, port blocks are filled completely before a new port block is allocated, and the last port block remains active indefinitely. The range is 0 through 86,400, and the default is 0.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set active-block-timeout timeout-interval
```

- c. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- d. Configure the maximum number of blocks that can be allocated to a user address. The range is 1 through 512, and the default is 8.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set maximum-blocks-per-host maximum-block-number
```

- e. Specify how often to send interim system logs for active port blocks and for inactive port blocks with live sessions. This increases the reliability of system logs, which are UDP-based and can get lost in the network. The range is 1800 through 86,400 seconds, and the default is 0 (interim logs are disabled).

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set interim-logging-interval timeout-interval
```

10. Specify the timeout period for endpoint independent translations that use the specified NAT pool. Mappings that are inactive for this amount of time are dropped. The range is 120 through 86,400 seconds. If you do not configure **ei-mapping-timeout**, then the **mapping-timeout** value is used for endpoint independent translations.

```
[edit services nat source pool nat-pool-name]
user@host# set ei-mapping-timeout ei-mapping-timeout
```

11. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

12. Define the NAT pool utilization levels that trigger SNMP traps. The **raise-threshold** is the pool utilization percentage that triggers the trap, and the range is 50 through 100. The **clear-threshold** is the pool utilization percentage that clears the trap, and the range is 40 through 100. For pools that use port-block allocation, the utilization is based on the number of ports that are used; for pools that do not use port-block allocation, the utilization is based on the number of addresses that are used.

```
[edit services nat source pool nat-pool-name]
user@host# set pool-utilization-alarm raise-threshold value
user@host# set pool-utilization-alarm clear-threshold value
```

If you do not configure **pool-utilization-alarm**, traps are not created.

13. To allow the IP addresses of a NAT pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**. However, pools that configure port-block allocation must not overlap with other pools.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Source Rule for NAPT

To configure the NAT source rule for NAPT:

1. Configure the NAT rule name.

```
[edit services nat source]
user@host# edit rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the source addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:

```
[edit services nat source rule-set rule-set-name rule rule-name rule rule-name]
user@host# set match source-address any-unicast
```

4. Specify one or more application protocols to which the NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

5. Specify the NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

6. Configure the address-pooling paired feature if you want to ensure assignment of the same external IP address for all sessions originating from the same internal host.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat mapping-type]
user@host# set address-pooling
```

7. If you want to ensure that the same external address and port are assigned to all connections from a given host, configure endpoint-independent mapping:
 - a. Configure the mapping type as endpoint independent.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set mapping-type endpoint-independent
```

- b. Specify prefix lists that contain the hosts that are allowed to establish inbound connections using the endpoint-independent mapping. (Prefix lists are configured at the **[edit policy-options]** hierarchy level.)

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set filtering-type endpoint-independent prefix-list [allowed-host] except [denied-host]
```

- c. Specify the maximum number of inbound flows allowed simultaneously on an endpoint-independent mapping.

```
[edit services nat source rule-set rule-set-name rule rule-name filtering-type then source-nat]
user@host# set secure-nat-mapping eif-flow-limit number-of-flows
```

- d. Specify the direction in which active endpoint-independent mapping is refreshed. By default, mapping is refreshed for both inbound and outbound active flows.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set secure-nat-mapping mapping-refresh (inbound | inbound-outbound | outbound)
```

8. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for NAT

To configure the service set for NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

RELATED DOCUMENTATION

| [Network Address Port Translation \(NAPT\) Overview](#) | 125

Configuring Syslog Events for NAT Rule Conditions with Next Gen Services

To configure syslog events to be generated when traffic matches NAT rule conditions for Next Gen Services NAT:

1. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

The following are logs collected:

Out of addresses logs — If the allocation request fails to be handled as the public IP addresses in the No-PAT pool are used up, the out of addresses syslog is generated.

Out of ports logs — If the allocation request fails to be handled as the public IPs and ports in the NAPT pool are used up, the out of ports syslog is generated.

NAT Rule Match Logs — If the packet matches the NAT rule, the NAT rule match syslog is generated.

Pool resource release logs — If the public IP and port succeeds to be released to the NAPT pool, the pool release syslog is generated.

RELATED DOCUMENTATION

[Network Address Port Translation \(NAPT\) Overview | 125](#)

[Configuring Network Address Port Translation for Next Gen Services | 126](#)

NAT46

IN THIS CHAPTER

- [NAT46 Next Gen Services Configuration Examples | 134](#)

NAT46 Next Gen Services Configuration Examples

IN THIS SECTION

- [NAT46 Support Summary | 135](#)
- [NAT46 Sample Configuration | 136](#)

Starting in Junos OS Release 20.2R1 you can run NAT46 Next Gen Services.

Starting in Junos OS Release 20.2R1, Network Address Translation and Protocol Translation (NAT-PT) [RFC2766] are supported for CGNAT Next Gen Services. NAT46 is a IPv4-to-IPv6 transition mechanism that provides a way for end-nodes in IPv6 realm to communicate with end-nodes in IPv4 realm and vice versa. This is achieved using a combination of Network Address Translation and Protocol Translation.

NAT46 is supported on both the SRX and on MX240, MX480, and MX960 for CGNAT Next Gen Services. This topic provides example configurations to help you understand how to configure NAT46 CGNAT Next Gen Services on these MX Series routers.

NOTE: These examples are for SRX devices, however, you can use these same examples to configure NAT46 Next Gen Services on MX Series devices by substituting **[edit security.....]** configuration statements with **[edit services.....]** configuration statements on the MX Series devices.

You can find these examples here: [IPv6 NAT](#)

There are four examples available:

- **Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping** — This example shows how to configure an IPv4-initiated connection to an IPv6 node using default destination address prefix static mapping.
- **Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping** — This example shows how to configure an IPv4-initiated connection to an IPv6 node using static destination address one-to-one mapping.
- **Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping** — This example shows how to configure an IPv6-initiated connection to an IPv4 node using default destination address prefix static mapping. This example does not show how to configure the NAT translation for the reverse direction.
- **Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping** — This example shows how to configure an IPv6-initiated connection to an IPv4 node using static destination address one-to-one mapping.

NAT46 Support Summary

NAT46 for Next Gen Services supports the following:

- ICMP, TCP, and UDP protocol packets.
- Static mapping is used to communicate between the IPv4 to IPv6 side of the subscriber connection.
- Bi-directional traffic flow is supported if you have other ways to convey the mapping between the IPv6 address and the dynamically allocated IPv4 address.
- NAT46 supports DNS, ICMP, and FTP ALGs.

Keep these things in mind when configuring NAT46 for Next Gen Services:

- No support of NAT64 feature described in NAT-PT (RFC 2765).
- Static NAT is not used for the source translation in any NAT scenario.
- Except DNS, FTP and ICMP, other ALGs are not supported for NAT46.
- AMS functionality is not supported for NAT46.
- Port translation is not tested with Source Address NAT (when source pool is a IPv6 prefix) for the NAT46 feature.

NAT46 Sample Configuration

```

services {
  nat {
    source {
      pool ipv6_prefix {
        address 27a6::/96;
      }
    }
    rule-set myipv6_rs {
      rule ipv6_rule {
        match {
          source-address 10.1.1.1/30 ;
          destination-address 27a6::a0a:a2d/126;
        }
        then {
          source-nat {
            pool {
              ipv6_prefix;
            }
          }
        }
      }
    }
    match-direction input;
  }
}

static {
  rule-set test_rs {
    rule test_rule {
      match {
        destination-address ip-address;
      }
      then {
        static-nat {
          prefix ip-address;
        }
      }
    }
  }
}

.....match-direction input;
}

}

service-set sset1 {
  ...
  nat-rule-sets test_rs;
}

```

```
        nat-rule-sets myipv6_rs;
        ...
    }
}
```

Release History Table

Release	Description
20.2R1	Starting in Junos OS Release 20.2R1 you can run NAT46 Next Gen Services.
20.2R1	Starting in Junos OS Release 20.2R1, Network Address Translation and Protocol Translation (NAT-PT) [RFC2766] are supported for CGNAT Next Gen Services.

RELATED DOCUMENTATION

service-set 678
<i>Configuring Service Sets for Network Address Translation</i>

Stateful NAT64 Overview and Configuration

IN THIS CHAPTER

- [Stateful NAT64 Overview | 138](#)
- [IPv4 Addresses Embedded in IPv6 Addresses | 139](#)
- [Configuring Next Gen Services Stateful NAT64 | 140](#)

Stateful NAT64 Overview

Stateful NAT64 translates IPv6 addresses to public IPv4 addresses, allowing IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP. Stateful NAT64 translates the destination IPv6 address to the embedded IPv4 address, and translates the source IPv6 address to a public IPv4 address and port from a block of IPv4 addresses that you set aside.

Stateful NAT64 supports the following:

- Round-robin port and address allocation (see [“Round-Robin Port Allocation” on page 221](#)).
- Address pooling and endpoint independent mapping (see [“Address Pooling and Endpoint Independent Mapping for Port Translation” on page 219](#)).
- Secured port block allocation (see [“Secured Port Block Allocation for Port Translation” on page 222](#)).

Benefits of Stateful NAT64

Stateful NAT64 provides a way to:

- Let IPv6-only clients contact IPv4 servers using unicast UDP, TCP, or ICMP
- Move to an IPv6 network
- Deal with IPv4 address depletion

RELATED DOCUMENTATION

IPv4 Addresses Embedded in IPv6 Addresses

Stateful NAT64 and XLAT464 embed IPv4 addresses in IPv6 addresses by using an IPv6 prefix that you specify. The prefix length you use determines how the IPv4 address is embedded.

IPv6 addresses with embedded IPv4 addresses are composed of a variable-length prefix, the embedded IPv4 address, and a variable-length suffix. Bits 64 to 71 are reserved and must be set to 0. The suffix follows the last bit of the embedded IPv4 address, and the suffix bits are ignored and should be set to 0.

The format for the IPv4-embedded IPv6 address depends on the prefix length, as shown in [Table 34 on page 139](#).

Table 34: IPv6 Address With Embedded IPv4 Address

Prefix length	Prefix bits	IPv4 address bits	Reserved bits (must be set to 0)	Suffix bits
32	0-31	32 to 63	64 to 71	72 to 127
40	0 to 39	40 to 63 and 72 to 79	64 to 71	80 to 127
48	0 to 47	48 to 63 and 72 to 87	64 to 71	88 to 127
56	0 to 55	56 to 63 and 72 to 95	64 to 71	96 to 127
64	0 to 63	72 to 103	64 to 71	104 to 127
96	0 to 95	96 to 127	64 to 71	No suffix bits

The following table shows an example of an IPv4 address embedded in an IPv6 address for various prefix lengths.

IPv6 Prefix	IPv4 Address	IPv4 Address Embedded in IPv6 Address
2001:db8::/32	192.0.2.33	2001:db8:c000:221::
2001:db8:100::/40	192.0.2.33	2001:db8:1c0:2:21::
2001:db8:122::/48	192.0.2.33	2001:db8:122:c000:2:2100::
2001:db8:122:300::/56	192.0.2.33	2001:db8:122:3c0:0:221::

IPv6 Prefix	IPv4 Address	IPv4 Address Embedded in IPv6 Address
2001:db8:122:344::/64	192.0.2.33	2001:db8:122:344:c0:2:2100::
2001:db8:122:344::/96	192.0.2.33	2001:db8:122:344::192.0.2.33

Configuring Next Gen Services Stateful NAT64

IN THIS SECTION

- [Configuring the Source Pool for Stateful NAT64 | 140](#)
- [Configuring the NAT Rules for Stateful NAT64 | 144](#)
- [Configuring the Service Set for Stateful NAT64 | 146](#)
- [Clearing the Don't Fragment Bit | 147](#)

Perform the following steps to configure Next Gen Services Stateful NAT64

Configuring the Source Pool for Stateful NAT64

To configure the source pool for Stateful NAT64:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]
user@host# set address address-prefix to address address-prefix
```

To disable round-robin port allocation for all NAT pools that do not specify an **automatic** (**random-allocation** | **round-robin**) setting, configure the global setting.

```
[edit services nat source]
user@host# set port-round-robin disable
```

3. To configure a range of ports to assign to a pool, perform the following:

NOTE: If you specify a range of ports to assign, the **automatic** statement is ignored.

- a. Specify the low and high values for the port. If you do not configure automatic port assignment, you must configure a range of ports.

```
[edit services nat source pool nat-pool-name port]
user@host# set range port-low to port-high
```

- b. Specify either random allocation or round-robin allocation. Round-robin allocation is the default.

```
[edit services nat source pool nat-pool-name port range]
user@host# set (random-allocation | round-robin)
```

4. Assign a port within the same range as the incoming port—either 0 through 1023 or 1024 through 65,535. This feature is not available if you configure port-block allocation.

```
[edit services nat source pool nat-pool-name port]
user@host# set preserve-range
```

5. Assign a port with the same parity (even or odd) as the incoming port. This feature is not available if you configure port-block allocation.

```
[edit services nat source pool nat-pool-name port]
user@host# set preserve-parity
```

6. Configure a global default port range for NAT pools that use port translation. This port range is used when a NAT pool does not specify a port range and does not specify automatic port assignment. The global port range can be from 1024 through 65,535.

```
[edit services nat source]
user@host# set pool-default-port-range port-low to port-high
```

7. Configure the source pool without port translation.

```
[edit services nat source pool nat-pool-name]
user@host# set address-pooling no-paired
```

8. Configure the maximum number of ports that can be allocated for each host. The range is 2 through 65,535.

```
[edit services nat source pool nat-pool-name]
user@host# set limit-ports-per-host number
```

9. If you want to allocate a block of ports for each subscriber to use, configure port-block allocation:

- a. Configure the number of ports in a block. The range is 1 through 64,512 and the default is 128.

```
[edit services nat source pool nat-pool-name port]
user@host# set block-allocation block-size block-size
```

- b. Configure the interval, in seconds, for which the block is active. After the timeout, a new block is allocated, even if ports are available in the active block. If you set the timeout to 0, port blocks are filled completely before a new port block is allocated, and the last port block remains active indefinitely. The range is 0 through 86,400, and the default is 0.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set active-block-timeout timeout-interval
```

- c. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- d. Configure the maximum number of blocks that can be allocated to a user address. The range is 1 through 512, and the default is 8.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set maximum-blocks-per-host maximum-block-number
```

- e. Specify how often to send interim system logs for active port blocks and for inactive port blocks with live sessions. This increases the reliability of system logs, which are UDP-based and can get lost in the network. The range is 1800 through 86,400 seconds, and the default is 0 (interim logs are disabled).

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set interim-logging-interval timeout-interval
```

10. Specify the timeout period for endpoint independent translations that use the specified NAT pool. Mappings that are inactive for this amount of time are dropped. The range is 120 through 86,400 seconds. If you do not configure **ei-mapping-timeout**, then the **mapping-timeout** value is used for endpoint independent translations.

```
[edit services nat source pool nat-pool-name]
user@host# set ei-mapping-timeout ei-mapping-timeout
```

11. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

12. To allow the IP addresses of a NAT source pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Rules for Stateful NAT64

For Stateful NAT64, you must configure a source rule and a destination rule. To configure the NAT rules for Stateful NAT64:

1. Configure the source NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the IPv6 source addresses that are translated by the NAT rule.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

4. Configure the matching destination address as 0.0.0.0/0.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match destination-address 0.0.0.0/0
```

5. Specify one or more application protocols to which the NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

6. Specify the NAT source pool that contains the addresses for translated source addresses.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

7. Configure endpoint-independent mapping, which ensures that the same external address and port are assigned to all connections from a given host.

- a. Configure the mapping type as endpoint independent.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set mapping-type endpoint-independent
```

- b. Specify prefix lists that contain the hosts that are allowed to establish inbound connections using the endpoint-independent mapping. (Prefix lists are configured at the **[edit policy-options]** hierarchy level.)

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set filtering-type endpoint-independent prefix-list [allowed-host] except [denied-host]
```

- c. Specify the maximum number of inbound flows allowed simultaneously on an endpoint-independent mapping.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set secure-nat-mapping eif-flow-limit number-of-flows
```

- d. Specify the direction in which active endpoint-independent mapping is refreshed. By default, mapping is refreshed for both inbound and outbound active flows.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set secure-nat-mapping mapping-refresh (inbound | inbound-outbound | outbound)
```

8. Configure the destination NAT rule name.

```
[edit services nat destination]
user@host# set rule-set rule-set-name rule rule-name
```

9. Specify the traffic direction to which the destination NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

10. Specify the IPv6 prefix source addresses that are translated by the destination NAT rule. Use the same value that you used for the NAT source rule.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

11. Specify the prefix that is used to embed the IPv4 destination address in the IPv6 destination address.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat destination-prefix destination-prefix
```

12. Configure the IPv6 destination address to match. This is the IPv4 destination address embedded in IPv6 by using the **destination-prefix**.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

13. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat (source | destination) rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Stateful NAT64

To configure the service set for stateful NAT64:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
```

```
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface  
interface-name
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]  
user@host# set nat-rule-sets rule-set-name
```

Clearing the Don't Fragment Bit

To prevent unnecessary creation of IPv6 fragmentation headers when translating IPv4 packets that are less than 1280 bytes, you can specify that the don't fragment (DF) bit for IPv4 packet headers is cleared when the packet length is less than 1280 bytes.

```
[edit services nat natv6v4]  
user@host# set clear-dont-fragment-bit
```

RELATED DOCUMENTATION

| [Stateful NAT64 Overview](#) | 138

IPv4 Connectivity Across IPv6-Only Network Using 464XLAT Overview and Configuration

IN THIS CHAPTER

- 464XLAT Overview | 148
- IPv4 Addresses Embedded in IPv6 Addresses | 150
- Configuring 464XLAT Provider-Side Translator for IPv4 Connectivity Across IPv6-Only Network for Next Gen Services | 151

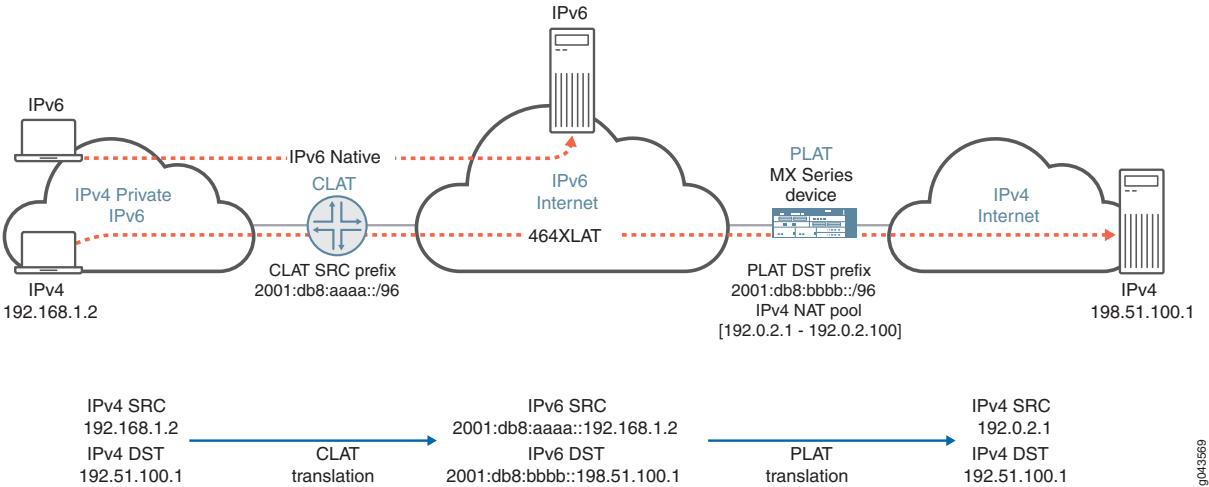
464XLAT Overview

You can configure the MX Series router as an 464XLAT Provider-Side Translator (PLAT). 464XLAT provides a simple and scalable technique for an IPv4 client with a private address to connect to an IPv4 host over an IPv6 network. 464XLAT only supports IPv4 in the client-server model, so it does not support IPv4 peer-to-peer communication or inbound IPv4 connections.

XLAT464 provides the advantages of not having to maintain an IPv4 network for this IPv4 traffic and not having to assign additional public IPv4 addresses.

A customer-side translator (CLAT), which is not a Juniper Networks product, translates the IPv4 packet to IPv6 by embedding the IPv4 source and destination addresses in IPv6 prefixes, and sends the packet over an IPv6 network to the PLAT. The PLAT translates the packet to IPv4, and sends the packet to the IPv4 host over an IPv4 network (see [Figure 1 on page 149](#)).

Figure 1: 464XLAT Wireline Flow

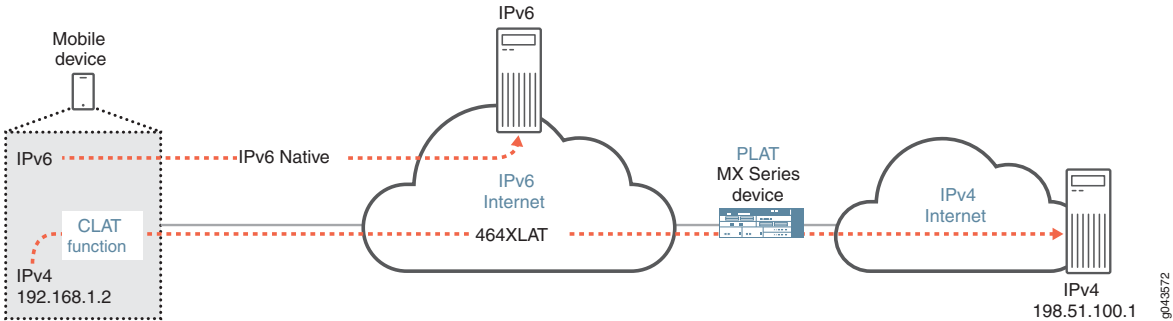


The CLAT uses a unique source IPv6 prefix for each end user, and translates the IPv4 source address to an IPv6 address by embedding it in the IPv6 /96prefix. In [Figure 1 on page 149](#), the CLAT source IPv6 prefix is 2001:db8:aaaa::/96, and the IPv4 source address 192.168.1.2 is translated to 2001:db8:aaaa::192.168.1.2. The CLAT translates the IPv4 destination address to IPv6 by embedding it in the IPv6 prefix of the PLAT (MX Series router). In [Figure 1 on page 149](#), the PLAT destination IPv6 prefix is 2001:db8:bbbb::/96, so the CLAT translates the IPv4 destination address 198.51.100.1 to 2001:db8:bbbb::198.51.100.

The PLAT translates the IPv6 source address to a public IPv4 address, and translates the IPv6 destination address to a public IPv4 address by removing the PLAT prefix.

The CLAT can reside on the end user mobile device in an IPv6-only mobile network, allowing mobile network providers to roll out IPv6 for their users *and* support IPv4-only applications on mobile devices (see [Figure 2 on page 149](#)).

Figure 2: 464XLAT Wireless Flow



464XLAT supports the following:

- Address pooling and endpoint independent mapping (see [“Address Pooling and Endpoint Independent Mapping for Port Translation” on page 219](#)).
- Secured port block allocation (see [“Secured Port Block Allocation for Port Translation” on page 222](#)

Benefits of 464XLAT

- No need to maintain an IPv4 transit network
- No need to assign additional public IPv4 addresses

RELATED DOCUMENTATION

IPv4 Addresses Embedded in IPv6 Addresses

Stateful NAT64 and XLAT464 embed IPv4 addresses in IPv6 addresses by using an IPv6 prefix that you specify. The prefix length you use determines how the IPv4 address is embedded.

IPv6 addresses with embedded IPv4 addresses are composed of a variable-length prefix, the embedded IPv4 address, and a variable-length suffix. Bits 64 to 71 are reserved and must be set to 0. The suffix follows the last bit of the embedded IPv4 address, and the suffix bits are ignored and should be set to 0.

The format for the IPv4-embedded IPv6 address depends on the prefix length, as shown in [Table 34 on page 139](#).

Table 35: IPv6 Address With Embedded IPv4 Address

Prefix length	Prefix bits	IPv4 address bits	Reserved bits (must be set to 0)	Suffix bits
32	0-31	32 to 63	64 to 71	72 to 127
40	0 to 39	40 to 63 and 72 to 79	64 to 71	80 to 127
48	0 to 47	48 to 63 and 72 to 87	64 to 71	88 to 127
56	0 to 55	56 to 63 and 72 to 95	64 to 71	96 to 127
64	0 to 63	72 to 103	64 to 71	104 to 127

Table 35: IPv6 Address With Embedded IPv4 Address (*continued*)

Prefix length	Prefix bits	IPv4 address bits	Reserved bits (must be set to 0)	Suffix bits
96	0 to 95	96 to 127	64 to 71	No suffix bits

The following table shows an example of an IPv4 address embedded in an IPv6 address for various prefix lengths.

IPv6 Prefix	IPv4 Address	IPv4 Address Embedded in IPv6 Address
2001:db8::/32	192.0.2.33	2001:db8:c000:221::
2001:db8:100::/40	192.0.2.33	2001:db8:1c0:2:21::
2001:db8:122::/48	192.0.2.33	2001:db8:122:c000:2:2100::
2001:db8:122:300::/56	192.0.2.33	2001:db8:122:3c0:0:221::
2001:db8:122:344::/64	192.0.2.33	2001:db8:122:344:c0:2:2100::
2001:db8:122:344::/96	192.0.2.33	2001:db8:122:344::192.0.2.33

Configuring 464XLAT Provider-Side Translator for IPv4 Connectivity Across IPv6-Only Network for Next Gen Services

1. [Configuring the Source Pool for 464XLAT | 151](#)
2. [Configuring the NAT Rules for 464XLAT | 153](#)
3. [Configuring the Service Set for 464XLAT | 156](#)
4. [Clearing the Don't Fragment Bit | 157](#)

Configuring the Source Pool for 464XLAT

To configure the source pool for 464XLAT:

1. Create a source NAT pool that is used to translate source IPv6 addresses to source public IPv4 addresses on PLAT.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]
user@host# set address address-prefix
```

3. If you want to allocate a block of ports for each subscriber to use, configure port-block allocation:

- a. Configure the number of ports in a block. The range is 1 through 64,512 and the default is 128.

```
[edit services nat source pool nat-pool-name port]
user@host# set block-allocation block-size block-size
```

- b. Configure the interval, in seconds, for which the block is active. After the timeout, a new block is allocated, even if ports are available in the active block. If you set the timeout to 0, port blocks are filled completely before a new port block is allocated, and the last port block remains active indefinitely. The range is 0 through 86,400, and the default is 0.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set active-block-timeout timeout-interval
```

- c. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- d. Configure the maximum number of blocks that can be allocated to a user address. The range is 1 through 512, and the default is 8.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set maximum-blocks-per-host maximum-block-number
```

- e. Specify how often to send interim system logs for active port blocks and for inactive port blocks with live sessions. This increases the reliability of system logs, which are UDP-based and can get lost in the network. The range is 1800 through 86,400 seconds, and the default is 0 (interim logs are disabled).

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set interim-logging-interval timeout-interval
```

4. Specify the timeout period for endpoint independent translations that use the specified NAT pool. Mappings that are inactive for this amount of time are dropped. The range is 120 through 86,400 seconds. If you do not configure **ei-mapping-timeout**, then the **mapping-timeout** value is used for endpoint independent translations.

```
[edit services nat source pool nat-pool-name]
user@host# set ei-mapping-timeout ei-mapping-timeout
```

5. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

Configuring the NAT Rules for 464XLAT

For 464XLAT, you must configure a source rule and a destination rule. To configure the NAT rules for 464XLAT:

1. Configure the source NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the CLAT IPv6 source prefix.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat clat-prefix clat-prefix
```

4. Configure the IPv6 source address prefix to match. This is the IPv4 source address embedded in IPv6 by using the CLAT prefix.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

5. Specify the NAT source pool that the PLAT uses for converting the IPv6 source address to a public IPv4 address.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

6. If you want to ensure that the same external address and port are assigned to all connections from a given host, configure endpoint-independent mapping:

- a. Configure the mapping type as endpoint independent.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set mapping-type endpoint-independent
```

- b. Specify prefix lists that contain the hosts that are allowed to establish inbound connections using the endpoint-independent mapping. (Prefix lists are configured at the **[edit policy-options]** hierarchy level.)

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set filtering-type endpoint-independent prefix-list [allowed-host] except [denied-host]
```

- c. Specify the maximum number of inbound flows allowed simultaneously on an endpoint-independent mapping.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set secure-nat-mapping eif-flow-limit number-of-flows
```

- d. Specify the direction in which active endpoint-independent mapping is refreshed. By default, mapping is refreshed for both inbound and outbound active flows.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set secure-nat-mapping mapping-refresh (inbound | inbound-outbound | outbound)
```

- e. Configure the address-pooling paired feature if you want to ensure assignment of the same external IP address for all sessions originating from the same internal host.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat mapping-type]
user@host# set address-pooling-paired
```

- f. Specify the timeout period for **address-pooling-paired** mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- g. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

7. Configure the destination NAT rule name.

```
[edit services nat destination]
user@host# set rule-set rule-set-name rule rule-name
```

8. Specify the traffic direction to which the destination NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

9. Configure the IPv6 source address prefix to match. Use the same value that you used for the NAT source rule.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

10. Configure the PLAT destination IPv6 prefix.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat destination-prefix address
```

11. Configure the IPv6 destination address to match. This is the IPv4 destination address embedded in IPv6 by using the PLAT destination prefix.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

Configuring the Service Set for 464XLAT

To configure the service set for 464XLAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface interface-name
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
```

```
user@host# set nat-rule-sets rule-set-name
```

Clearing the Don't Fragment Bit

Specify that the don't fragment (DF) bit for IPv4 packet headers is cleared when the packet length is less than 1280 bytes.

```
[edit services nat natv6v4]  
user@host# set clear-dont-fragment-bit
```

This prevents unnecessary creation of an IPv6 fragmentation header when translating IPv4 packets that are less than 1280 bytes.

IPv6 NAT Protocol Translation (NAT PT)

IN THIS CHAPTER

- [IPv6 NAT PT Overview | 158](#)
- [IPv6 NAT-PT Communication Overview | 159](#)

IPv6 NAT PT Overview

Starting in Junos OS Release 20.2R1 you can run IPv6 NAT-PT Next Gen Services on MX240, MX480, and MX960 routers.

IPv6 Network Address Translation-Protocol Translation (NAT-PT) provides address allocation and protocol translation between IPv4 and IPv6 addressed network devices. The translation process is based on the Stateless IP/ICMP Translation (SIIT) method; however, the state and the context of each communication are retained during the session lifetime. IPv6 NAT-PT supports Internet Control Message Protocol (ICMP), TCP, and UDP packets.

IPv6 NAT-PT supports the following types of NAT-PT:

- **Traditional NAT-PT**—In traditional NAT-PT, the sessions are unidirectional and outbound from the IPv6 network. Traditional NAT-PT allows hosts within an IPv6 network to access hosts in an IPv4 network. There are two variations to traditional NAT-PT: basic NAT-PT and NAPT-PT.

In basic NAT-PT, a block of IPv4 addresses at an IPv4 interface is set aside for translating addresses as IPv6 hosts as they initiate sessions to the IPv4 hosts. The basic NAT-PT translates the source IP address and related fields such as IP, TCP, UDP, and ICMP header checksums for packets outbound from the IPv6 domain. For inbound packets, it translates the destination IP address and the checksums.

Network Address Port Translation-Protocol Translation (NAPT-PT) can be combined with basic NAT-PT so that a pool of external addresses is used in conjunction with port translation. NAPT-PT allows a set of IPv6 hosts to share a single IPv4 address. NAPT-PT translates the source IP address, source transport identifier, and related fields such as IP, TCP, UDP, and ICMP header checksums, for packets outbound from the IPv6 network. The transport identifier can be a TCP/UDP port or an ICMP query ID. For inbound packets, it translates the destination IP address, destination transport identifier, and the IP and the transport header checksums.

- **Bidirectional NAT-PT**—In bidirectional NAT-PT, sessions can be initiated from hosts in the IPv4 network as well as the IPv6 network. IPv6 network addresses are bound to IPv4 addresses, either statically or dynamically as connections are established in either direction. The static configuration is similar to static NAT translation. Hosts in IPv4 realm access hosts in the IPv6 realm using DNS for address resolution. A DNS ALG must be employed in conjunction with bidirectional NAT-PT to facilitate name-to-address mapping. Specifically, the DNS ALG must be capable of translating IPv6 addresses in DNS queries and responses into their IPv4 address bindings, and vice versa, as DNS packets traverse between IPv6 and IPv4 realms.

NOTE: The devices partially support the bidirectional NAT-PT specification. It supports flow of bidirectional traffic assuming that there are other ways to convey the mapping between the IPv6 address and the dynamically allocated IPv4 address. For example, a local DNS can be configured with the mapped entries for IPv4 nodes to identify the addresses.

NAT-PT Operation—The devices support the traditional NAT-PT and allow static mapping for the user to communicate from IPv4 to IPv6. The user needs to statically configure the DNS server with an IPv4 address for the hostname and then create a static NAT on the device for the IPv6-only node to communicate from an IPv4-only node to an IPv6-only node based on the DNS.

Release History Table

Release	Description
20.2R1	Starting in Junos OS Release 20.2R1 you can run IPv6 NAT-PT Next Gen Services on MX240, MX480, and MX960 routers.

RELATED DOCUMENTATION

[NAT46 Next Gen Services Configuration Examples](#) | 134

IPv6 NAT-PT Communication Overview

NAT-PT communication with static mapping— Network Address Translation-Protocol Translation (NAT-PT) can be done in two directions, from IPv6 to IPv4 and vice versa. For each direction, static NAT is used to map the destination host to a local address and a source address NAT is used to translate the source address. There are two types of static NAT and source NAT mapping: one-to-one mapping and prefix-based mapping.

NAT- PT communication with DNS ALG—A DNS-based mechanism dynamically maps IPv6 addresses to IPv4-only servers. NAT-PT uses the DNS ALG to transparently do the translations. For example, a company using an internal IPv6 network needs to be able to communicate with external IPv4 servers that do not yet have IPv6 addresses.

To support the dynamic address binding, a DNS should be used for name resolution. The IPv4 host looks up the name of the IPv6 node in its local configured IPv4 DNS server, which then passes the query to the IPv6 DNS server through a device using NAT-PT.

The DNS ALG in NAT device :

- Translates the IPv6 address resolution back to IPv4 address resolution.
- Allocates an IPv6 address for the mapping.
- Stores a mapping of the allocated IPv4 address to the IPv6 address returned in the IPv6 address resolution so that the session can be established from any-IPv4 hosts to the IPv6 host.

RELATED DOCUMENTATION

| [IPv6 NAT PT Overview](#) | 158

Stateless Source Network Prefix Translation for IPv6

Overview and Configuration

IN THIS CHAPTER

- [Stateless Source Network Prefix Translation for IPv6 | 161](#)

Stateless Source Network Prefix Translation for IPv6

IN THIS SECTION

- [Stateless Source Network Prefix Translation for IPv6 for IPv6 | 161](#)
- [Configuring NPTv6 for Next Gen Services | 162](#)

Stateless Source Network Prefix Translation for IPv6 for IPv6

IN THIS SECTION

- [Benefits of Stateless Source Network Prefix Translation | 162](#)

When an IPv6 packet is going from an internal network to the external network, Stateless Source Network Prefix Translation for IPv6 (NPTv6) maps the IPv6 prefix of the source address to an IPv6 prefix of an external network. When an IPv6 packet is coming from the external network to the internal network, NPTv6 maps the IPv6 prefix of the destination address to the IPv6 prefix of the internal network.

NPTv6 uses an algorithm to translate the addresses, and does not need to maintain the state for each node or each flow in the translator. NPTv6 also removes the need to recompute the transport layer checksum.

Benefits of Stateless Source Network Prefix Translation

- For edge networks, you do not need to renumber the IPv6 addresses used inside the local network for interfaces, access lists, and system logging messages if:
 - The global prefixes used by the edge network are changed.
 - The IPv6 addresses are used inside the edge network or within other upstream networks (such as multihomed devices) when a site adds, drops, or changes upstream networks.
- IPv6 addresses used by the edge network do not need ingress filtering in upstream networks and do not need their customer-specific prefixes advertised to upstream networks.
- Connections that traverse the translation function are not disrupted by a reset or brief outage of an NPTv6 translator.

Configuring NPTv6 for Next Gen Services

1. [Configuring the Source Pool | 162](#)
2. [Configuring the NAT Rule | 162](#)
3. [Configuring the Service Set | 164](#)

Configuring the Source Pool

To configure the source pool for NPTv6:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the IPv6 prefix to which the IPv6 source address prefix is translated.

```
[edit services nat source pool nat-pool-name]
user@host# set address address-prefix
```

Configuring the NAT Rule

To configure the NAT source rule for NPTv6:

1. Configure the NAT rule name.

```
[edit]
user@host# edit services nat source rule-set rule-set-name rule rule-name
```

- Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

- Specify the IPv6 prefix of source addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

- Configure the address-pooling paired feature if you want to ensure assignment of the same external IP address for all sessions originating from the same internal host.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat mapping-type]
user@host# set address-pooling-paired
```

- Specify the timeout period for **address-pooling-paired** mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- Specify the NAT pool that contains the IPv6 prefix for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

- Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set

To configure the service set for NPTv6:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service set, which requires a single service interface, or a next-hop service set, which requires an inside and outside service interface.

- To configure an interface service set:

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface vms-slot-number/pic-number/0.logical-unit-number
```

- To configure a next-hop service set:

```
[edit services service-set service-set-name]
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface vms-slot-number/pic-number/0.logical-unit-number
outside-service-interface vms-slot-number/pic-number/0.logical-unit-number
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

4. Specify that ICMP error messages are sent if NPTv6 address translation fails.

```
[edit services service-set service-set-name nat-options nptv6]
user@host# set icmpv6-error-messages
```

Transitioning to IPv6 Using Softwires

IN THIS CHAPTER

- 6rd Softwires in Next Gen Services | 165

6rd Softwires in Next Gen Services

IN THIS SECTION


- 6rd Softwires in Next Gen Services Overview | 165
- Configuring Inline 6rd for Next Gen Services | 166

6rd Softwires in Next Gen Services Overview

Next Gen Services supports a 6rd softwire concentrator on the MX-SPC3 services card. 6rd softwires allow IPv6 end users to send traffic over an IPv4 network to reach an IPv6 network. IPv6 packets are encapsulated in IPv4 packets by a softwire initiator at the customer edge WAN, and tunneled to a 6rd softwire concentrator. A softwire is created when IPv4 packets containing IPv6 destination information are received at the softwire concentrator, which decapsulates IPv6 packets and forwards them for IPv6 routing.

6rd softwire flow is shown in [Figure 3 on page 165](#).

Figure 3: 6rd Software Flow



In the reverse path, IPv6 packets are sent to the 6rd software concentrator, which encapsulates them in IPv4 packets corresponding to the proper software and sends them to the customer edge WAN.

IPv6 flows are also created for the encapsulated IPv6 payload, and are associated with the specific software that carried them in the first place. When the last IPv6 flow associated with a software ends, the software is deleted. This simplifies configuration and there is no need to create or manage tunnel interfaces.

For more information on 6rd softwares, see RFC 5969, *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification*.

Benefits

- Rapid deployment of IPv6 service to subscribers on native IPv4 customer edge WANs.
- No need to create or manage tunnel interfaces.

Configuring Inline 6rd for Next Gen Services

IN THIS SECTION

- [Configuring a 6rd Software Concentrator | 166](#)
- [Configuring a 6rd Software Rule | 167](#)
- [Configuring Inline Services and an Inline Services Interface | 168](#)
- [Configuring the IPv4-Facing and IPv6-Facing Interfaces for 6rd | 169](#)
- [Configuring the Service Set | 169](#)

Configuring a 6rd Software Concentrator

To configure a 6rd software concentrator:

1. Configure a 6rd software concentrator name and IP address.

```
user@host# edit services softwares software-name software-name
```

For example:

```
user@host# edit services softwares software-name sw1
```

2. Configure the software type as v6rd and specify a name for it.


```
[edit services softwires software-name sw1]
user@host# set software-type v6rd name
```

For example:

```
[edit services softwires software-name sw1]
user@host# edit software-type v6rd 6rd-sw1
```

3. Configure the 6rd domain's IPv6 prefix.

```
[edit services softwires software-name sw1 software-type v6rd 6rd-sw1]
user@host# set v6rd-prefix v6rd-prefix
```

SEE ALSO

Configuring a 6rd Software Rule

To configure a 6rd software rule:

1. Specify the name of the rule set that the rule belongs to.

```
[edit services softwires]
user@host# set rule-set rule-set-name
```

2. Specify the direction of traffic to be tunneled.

```
[edit services softwires rule-set rule-set-name]
user@host# set match-direction (input | output)
```

3. Specify the name of the rule.

```
[edit services softwires rule-set rule-set-name]
user@host# set rule rule-name
```

4. Specify the software to apply if the condition is met.

```
[edit services softwires rule-set rule-set-name rule rule-name]
user@host# set then v6rd 6rd-software-name
```

Configuring Inline Services and an Inline Services Interface

Inline services run on MX line cards that can operate under Next Gen Services, for example MPC3 and MPC4 cards. This topic describes how to enable an inline service.

To enable inline services and an inline services interface:

1. Enable inline services for the FPC and PIC slot, and define the amount of bandwidth to dedicate to inline services.

```
[edit chassis fpc slot-number pic number]
user@host# set inline-services bandwidth (1g | 10g | 20g | 30g | 40g | 100g)
```

2. Configure the inline services logical interfaces. Inline interfaces use the following interface naming convention:

```
si-slot/pic/port
```

- If you are using an interface service set, configure one logical unit, and include units for IPv4 and IPv6:

```
user@host# set interfaces si-slot-number/pic-number/0 unit unit-number family inet
user@host# set interfaces si-slot-number/pic-number/0 unit unit-number family inet6
```

For example:

```
user@host# set interfaces si-0/0/0 unit 0 family inet
user@host# set interfaces si-0/0/0 unit 0 family inet6
```

- If you are using a next-hop service set, configure two logical units and define the inside and outside interfaces for IPv4 and IPv6:

```
[edit interfaces si-slot-number/pic-number/0]
user@host# set unit inside-unit-number family inet
user@host# set unit inside-unit-number family inet6
user@host# set unit inside-unit-number service-domain inside
user@host# set unit outside-unit-number family inet
user@host# set unit outside-unit-number family inet6
user@host# set unit outside-unit-number service-domain outside
```

For example:

```
user@host# set interfaces si-0/0/0 unit 1 family inet
```

```

user@host# set interfaces si-0/0/0 unit 1 family inet6
user@host# set interfaces si-0/0/0 unit 1 service-domain inside
user@host# set interfaces si-0/0/0 unit 2 family inet
user@host# set interfaces si-0/0/0 unit 2 family inet family inet6
user@host# set interfaces si-0/0/0 unit 2 service-domain outside

```

Configuring the IPv4-Facing and IPv6-Facing Interfaces for 6rd

To configure the IPv4-facing and IPv6-facing interfaces:

1. Configure the IPv4-facing interface:

- To configure an interface to use with an interface-style service set, configure input and output service and specify the service set.

```

user@host# set interfaces interface-name unit unit-number family inet service input service-set
service-set-name
user@host# set interfaces interface-name unit unit-number family inet service output service-set
service-set-name
user@host# set interfaces interface-name unit unit-number family inet address ip-address

```

- To configure an interface to use with a next-hop style service set, omit the **service input** and **service output** references.

```

user@host# set interfaces interface-name unit unit-number family inet
user@host# set interfaces interface-name unit unit-number family inet address ip-address

```

2. Configure the IPv6-facing interface.

```

user@host# set interface-name unit unit-number family inet6 address ipv6-address

```

Configuring the Service Set

To configure the service set for 6rd processing:

1. Specify a name for the service set.

```

[edit services]
user@host# edit service-set service-set-name

```

2. Configure either an interface service set, which requires a single service interface, or a next-hop service set, which requires an inside and outside service interface.

- To configure an interface service set:

```
[edit services service-set service-set-name]  
user@host# set interface-service service-interface vms-slot-number/pic-number/0.unit-number
```

- To configure a next-hop service set:

```
[edit services service-set service-set-name]  
user@host# set next-hop-service inside-service-interface vms-slot-number/pic-number/0.inside-unit-number  
          outside-service-interface vms-slot-number/pic-number/0.outside-unit-number
```

3. Specify the 6rd rule-set that contains the 6rd rule to be used with the service set.

```
[edit services service-set service-set-name]  
user@host# set softwires-rule-set software-rule-set-name
```

Transitioning to IPv6 Using DS-Lite Softwires

IN THIS CHAPTER

- [DS-Lite Softwires—IPv4 over IPv6 for Next Gen Services | 171](#)
- [Configuring Next Gen Services DS-Lite Softwires | 174](#)
- [DS-Lite Subnet Limitation | 180](#)
- [Protecting CGN Devices Against Denial of Service \(DOS\) Attacks | 184](#)

DS-Lite Softwires—IPv4 over IPv6 for Next Gen Services

IN THIS SECTION

- [DS-Lite Softwires—IPv4 over IPv6 | 172](#)

Junos OS enables service providers to transition to IPv6 by using software encapsulation and decapsulation techniques. A software is a tunnel that is created between software customer premises equipment (CPE). A software CPE can share a unique common internal state for multiple softwires, making it a very light and scalable solution. When you use softwires, you need not maintain an interface infrastructure for each software, unlike a typical mesh of generic routing encapsulation (GRE) tunnels that requires you to do so. A software initiator at the customer end encapsulates native packets and tunnels them to a software concentrator at the service provider. The software concentrator decapsulates the packets and sends them to their destination. A software is created when a software concentrator receives the first tunneled packet of a flow and prepares the packet for flow processing. The software exists as long as the software concentrator is providing flows for routing. A flow counter is maintained; when the number of active flows is 0, the software is deleted. Statistics are kept for both flows and softwires.

This topic contains the following sections:

DS-Lite Softwires—IPv4 over IPv6

When an ISP begins to allocate new subscriber home IPv6 addresses and IPv6-capable equipment, dual-stack lite (DS-Lite) provides a method for the private IPv4 addresses behind the IPv6 customer edge WAN equipment to reach the IPv4 network. DS-Lite enables IPv4 customers to continue to access the Internet using their current hardware by using a softwire initiator, referred to as a Basic Bridging Broadband (B4), at the customer edge to encapsulate IPv4 packets into IPv6 packets and tunnel them over an IPv6 network to a softwire concentrator, referred to as an Address Family Transition Router (AFTR), for decapsulation. DS-Lite creates the IPv6 softwires that terminate on the services PIC. Packets coming out of the softwire can then have other services such as NAT applied on them.

Starting in Junos OS release 20.2R1, DS-Lite is supported Next Gen Services on MX240, MX480 and MX960 routers with the MX-SPC3.

For more information on DS-Lite softwires, see the IETF draft *Dual Stack Lite Broadband Deployments Following IPv4 Exhaustion*.

NOTE: The most recent IETF draft documentation for DS-Lite uses new terminology:

- The term *softwire initiator* has been replaced by *B4*.
- The term *softwire concentrator* has been replaced by *AFTR*.

The Junos OS documentation generally uses the original terms when discussing configuration in order to be consistent with the command-line interface (CLI) statements used to configure DS-Lite.

DS-Lite and NAT in Next Gen Services

In Next Gen Services, DS-Lite changes the way NAT works with respect to the **address-pooling-paired** statement for the endpoint independent mapping (EIM), endpoint independent filtering (EIF), and port block allocation (PBA) features. In the earlier Adaptive Services implementation, all of these NAT features are subscriber-based and the subscriber is either a B4 IP address or an IPv6 prefix. In addition, for Adaptive Services, the **address-pooling-paired** association is between internal IPv4 address and NAT pool address. However in Next Gen Services DS-Lite, the **address-pooling-paired** pairing is between either the subscriber (B4 IPv6 address or IPv6 prefix) and a NAT pool address. Otherwise, the **address-pooling-paired** functionality remains the same for Next Gen Services.

NOTE: For CGNAT Next Gen Services on the MX-SPC3 security services card, when you configure DS-Lite use the following rules:

- For non-prefix based DS-Lite subscriber softwires, specify the B4 IPv6 address as the software concentrator.
- For prefix-based DS-Lite subscriber softwires, specify the IPv6 prefix address as the software concentrator. In addition for prefix-based subscriber DS-Lite softwires, you must specify the subscriber prefix length per service-set under the **[edit software-options dslite-ipv6-prefix-length dslite-ipv6-prefix-length]** configuration hierarchy.

You create EIM mappings on a per-software basis and they are bound to B4 address; which means the rule matching criteria includes B4 address. For Next Gen Services DS-Lite softwires, there is no special mapping timeout for software sessions, instead, they take the value of **inactivity-non-tcp-timeout** as their timeout value.

When a subscriber requires a port to be assigned for the first time, Port Block Allocation (PBA) ensures a block of ports is allocated to that particular subscriber. All subsequent requests from this subscriber use ports from the assigned block. A new port block is allocated when the current active block is exhausted, or after the active port block timeout interval has expired.

DS-Lite and AMS

AMS groups several PICs together and load balances traffic across all PICs that are part of the same group. In a standalone PIC configuration, all software sessions originated from any B4, which are destined to a software concentrator, are serviced on the same PIC where the software concentrator is configured. In the case of a DS-Lite in an AMS configuration, the software concentrator is hosted on all PICs in AMS group, however, software sessions from various B4 devices are distributed across member PICs. Thus, a software session originated from one B4 to the software concentrator, is assigned to one member PIC and all packets (IPv4-in-IPv6 and inner IPv4) in both directions (originated from B4 and destined to B4) related to that software session are serviced in the same PIC.

For prefix-based DS-Lite subscribers you need to configure the IPv6-prefix for DS-Lite traffic. When a prefix-based subscriber is active, the configured prefix length is taken from the B4 address and is completed with trailing zeros to form a 128-bit IPv6 NAT subscriber. This means that all B4 entities with a matching prefix and all IPv4 networks behind those matching B4 entities, are all identified as a single subscriber. An option is provided to configure the subscriber prefix length per service-set under the **[edit software-options dslite-ipv6-prefix-length dslite-ipv6-prefix-length]** hierarchy.

NOTE: For CGNAT Next Gen Services on the MX-SPC3 security services card, when you configure prefix-based DS-Lite subscribers always specify the IPv6 prefix address for the software concentrator.

With the prefix-based subscriber feature enabled, only one subscriber context is maintained per-prefix. Hence, the Port Block Allocation (NAT PBA) function would account for port blocks per each subscriber, instead of every single B4 address. Session limits configured under the software concentrator, limit the number of IPv4 sessions per subscriber, instead of per software/B4 address. Enabling the **address-pooling-paired** option in prefix-based subscriber configurations results in one public IP address for the subscriber instead of per B4 address.

Release History Table

Release	Description
20.2R1	Starting in Junos OS release 20.2R1, DS-Lite is supported Next Gen Services on MX240, MX480 and MX960 routers with the MX-SPC3.

RELATED DOCUMENTATION

<i>Junos Address Aware Network Addressing Overview</i>
Configuring Next Gen Services DS-Lite Softwires 174
DS-Lite Subnet Limitation 180
DS-Lite Per Subnet Limitation Overview 180

Configuring Next Gen Services DS-Lite Softwires

IN THIS SECTION

- [Configuring Next Gen Services Software Rules | 175](#)
- [Configuring Service Sets for Next Gen Services Softwires | 176](#)
- [Configuring the DS-Lite Software | 178](#)

Configuring Next Gen Services Software Rules

You configure software rules to instruct the router how to direct traffic to the addresses specified for 6rd, DS-Lite, or MAP-E software concentrators. Software rules do not perform any filtration of the traffic. They do not include a **from** statement, and the only option in the **then** statement is to specify the address of the software concentrator.

Starting in Junos OS release 19.3R2 6rd softwires are supported. Starting in Junos OS release 20.2, DS-Lite and Mapping of Address and Port with Encapsulation (MAP-E).

You can create a software rule consisting of one or more terms and associate a particular 6rd, DS-Lite, or MAP-E software concentrator with each term. You can include the software rule in service sets along with other services rules.

To configure a software rule set:

1. Assign a name to the rule set.

```
[edit services softwires]
user@host# edit rule-set rule-set-name
```

For example:

```
[edit services softwires]
user@host# edit rule-set swrs1
```

2. Configure the input and output match directions for the rule set.

```
[edit services softwires rule-set swrs1]
user@host# set match-direction input
```

3. Specify the name of the rule to apply if the match in this direction is met.

```
[edit services softwires rule-set swrs1]
user@host# edit rule rule-name
```

For example:

```
[edit services softwires rule-set swrs1]
user@host# edit rule swr1
```

4. Associate a 6rd, DS-Lite or MAP-E software concentrator with this term.

```
[edit services softwires rule-set swrs1 rule swr1]
user@host# set then ds-lite | map- | v6rd
```

For example, to associate a DS-Lite software specify the name of the DS-Lite software.

```
[edit services softwires rule-set swrs1 rule swr1]
user@host# set then ds-lite dslsw1
```

5. Repeat steps 2 and 3, and 4 for the output direction.

SEE ALSO

[DS-Lite Softwires—IPv4 over IPv6 for Next Gen Services | 171](#)

[DS-Lite Subnet Limitation | 180](#)

[DS-Lite Per Subnet Limitation Overview | 180](#)

Configuring Service Sets for Next Gen Services Softwires

You must include previously defined NAT or stateful firewall software rules or a software rule set in a service set to enable software processing.

Starting in Junos OS release 20.2R1, DS-Lite, MAP-E and 6rd softwires are supported in MX240, MX480, and MX960 routers. MAP-E and 6rd softwires are supported inline on an MPC by specifying the si-1/0/0 interface naming convention. DS-Lite is softwires run on the MX-SPC3 security services card.

To configure service sets for softwires:

1. Specify a name for the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

For example:

```
[edit services]
user@host# edit service-set vms-sw-ss
```

2. Specify the IPv6 prefix length for the subscriber addresses.

```
[edit services service-set vms-sw-ss]
user@host# set software-options dslite-ipv6-prefix-length dslite-ipv6-prefix-length
```

We support four prefix lengths: 56, 64, 96 and 128, which is the default.

3. For NAT, you can include a NAT rule for flows originated by DS-Lite softwires.

NOTE:

Currently a NAT rule configuration is required with a DS-Lite software configuration when you use interface service set configurations; NAT is not required when using next-hop service set configurations. NAT processing from IPv4 to IPv6 address pools and vice versa is not currently supported. FTP, HTTP, and RSTP are supported.

NOTE: With a DS-Lite software, if you configure stateful firewall rules without configuring NAT rules, using an interface service set causes the ICMP echo reply messages to not be sent correctly to DS-Lite. This behavior occurs if you apply a service set to both inet and inet6 families. In such a scenario, the traffic that is not destined to the DS-Lite software concentrator is also processed by the service set and the packets might be dropped, although the service set must not process such packets.

To prevent the problem to incorrect processing of traffic applicable for DS-Lite, you must configure a next-hop style service set and not an interface style service set. This problem does not occur when you configure NAT rules with interface service sets for DS-Lite.

Specify the name of the NAT rule set.

```
[edit services service-set vms-sw-ss]
user@host# edit nat-rule-sets nat-rule-set-name
```

4. Specify the service interface to be used.

```
[edit services service-set vms-sw-ss]
user@host# set interface-service service-interface vms-interface-name
```

5. Specify the name of the previously defined softwires rule set that you want to apply to this service set.

```
[edit services service-set vms-sw-ss]
user@host# set softwires-rule-set rule-set-name
```

SEE ALSO

Configuring the DS-Lite Software

Starting in Junos OS release 20.2R1, you can configure DS-Lite softwires for Next Gen Services on the MX-SPC3 services card.

1. Specify a name for the DS-Lite software.

```
[edit]
user@host# edit services softwires software-types ds-lite name
```

2. Specify a name for the DS-Lite software.

```
[edit]
user@host# edit services softwires software-types ds-lite name
```

For example:

```
user@host# edit services softwires software-types ds-lite dslsw1
```

3. Specify the IPv6 address of the software concentrator.

NOTE: For CGNAT Next Gen Services on the MX-SPC3 security services card, when you configure DS-Lite concentrator, use the following rules:

- For non-prefix based DS-Lite subscribers, specify the B4 IPv6 address
- For prefix-based DS-Lite subscribers, specify the IPv6 prefix address

For example:

```
[edit services softwires software-types ds-lite dslsw1]
user@host# set software-concentrator B4-IPv6-address or IPv6-prefix-address
```

4. You can specify the maximum transmission unit (MTU) for the software tunnel automatically or manually.
 - a. To manually specify the MTUs for the software tunnel:

```
[edit services softwires software-types ds-lite dslsw1]
user@host# set mtu-v4 bytes
user@host# set mtu-v6 bytes
```

NOTE: This MTU-v6 option sets the maximum transmission unit when encapsulating IPv4 packets into IPv6. If the final length is greater than the MTU-v4 value, the IPv6 packet is fragmented. This option is mandatory because it depends on other network parameters under administrator control.

5. Specify the maximum number of flows for the software.

```
[edit services softwires software-types ds-lite dslsw1]
user@host# set flow-limit 1000
```

6. (Optional) For prefix-based DS-Lite subscriber softwires, configure the maximum number of subscriber sessions allowed per prefix. You can configure from 0 through 16,384 sessions.

```
[edit services softwires software-types ds-lite dslsw1]
user@host# set session-limit-per-prefix 12
```

NOTE: You cannot use **flow-limit** and **session-limit-per-prefix** in the same DS-Lite configuration.

7. Configure the size of the IPv4 subnet prefix to which limiting is applied. ipv4prefix=6rd customer edge ipv4

```
[edit services softwires software-types ds-lite dslsw1]
user@host# set ipv4-prefix
```

8. Configure the size of the IPv6 subnet prefix to which limiting is applied. Specify a prefix length of 56, 64, 96, or 128.

```
[edit services softwires software-types ds-lite dslsw1]
user@host# set v6rd-prefix
```

NOTE: Ensure that all mappings are cleared before changing the prefix length.

DS-Lite Subnet Limitation

IN THIS SECTION

- [DS-Lite Per Subnet Limitation Overview | 180](#)
- [Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks | 182](#)

DS-Lite Per Subnet Limitation Overview

Junos OS enables you to limit the number of software flows from a subscriber's basic bridging broadband (B4) device at a given point in time, preventing subscribers from excessive use of addresses within the subnet. This limitation reduces the risk of denial-of-service (DoS) attacks. This limitation is supported on MX Series routers equipped with MS-DPCs. Starting in Junos OS Release 18.2R1, MS-MPCs and MS-MICs also support the subnet limitation feature. Starting in Junos OS Release 19.2R1, MX Virtual Chassis and MX Broadband Network Gateway (BNG) routers also support the subnet limitation feature. Starting in Junos OS release 20.2R1, DS-Lite is supported for CGNAT Next Gen Services on MX240, MX480 and MX960 routers.

A household using IPv6 with DS-Lite is a subnet, not just an individual IP address. The subnet limitation feature associates a subscriber and mapping with an IPv6 prefix instead of an IPv6 address. A subscriber can use any IPv6 addresses in that prefix as a DS-Lite B4 address and potentially exhaust carrier-grade NAT resources. The subnet limitation feature enables greater control of resource utilization by identifying a subscriber with a prefix instead of a specific address.

The subnet limit provides the following features:

- Flows utilize the complete B4 address.
- Prefix length can be configured per service set under software-options for the individual service-set.

- Port blocks are allocated per prefix of the subscriber B4 device, and not on each B4 address (if the prefix length is less than 128). If the prefix length is 128, then each IPv6 address is treated as a B4. Port blocks are allocated per 128-bit IPv6 address.
- Session limit, defined under the DS-Lite software concentrator configuration, limits the number of IPv4 sessions for the prefix.
- EIM, EIF, and PCP mappings are created per software tunnel (full 128 bit IPv6 address). Stale mappings time out based on timeout values.
- If prefix length is configured, then PCP **max-mappings-per-subscriber** (configurable under **pcp-server**) is based on the prefix only, and not the full B4 address.
- SYSLOGS for PBA allocation and release contain the prefix portion of the address completed with all zeros. SYSLOGS for PCP allocate and release, flow creation and deletion will still contain the complete IPv6 address.

The **show services nat mappings address-pooling-paired** operational command output now shows the mapping for the prefix. The mapping shows the address of the active B4.

The **show services software statistics ds-lite** output includes a new field that displays the number of times the session limit was exceeded for the MPC.

For Next Gen Services on MX240, MX480, and MX960 routers, the subnet limit statistic is displayed in the **Software session limit exceeded** field.

show services software statistics (MX-SPC3)

```
user@host> show services software statistics
vms-2/0/0
  Total Session Interest events      :3
  Total Session Destroy events      :2
  Total Session Public Request events :0
  Total Session Accepts             :1
  Total Session Discards             :0
  Total Session Ignores              :0
  Total Session extension alloc failures :0
  Total Session extension set failures :0
Software statistics
  Total Software sessions created    :1
  Total Software sessions deleted    :2
  Total Software sessions created for reverse packets :1
  Total Software session create failed for reverse pkts :0
  Total Software rule match success  :1
  Total Software rule match failed   :0
  Software session limit exceeded     :0
Software packet statistics
  Total Packets processed            :1
```

```

Total packets encapsulated           :1
Total packets decapsulated           :1
Encapsulation errors                 :0
Decapsulation errors                 :0
Encapsulated pkts re-inject failures :0
Decapsulated pkts re-inject failures :0
DS-Lite ICMPv4 Echo replies sent     :0
DS-Lite ICMPv4 TTL exceeded messages sent :0
ICMPv6 ECHO request messages received destined to AFTR :0
ICMPv6 ECHO reply messages sent from AFTR :0
ICMPv6 ECHO requests to AFTR process failures :0
V6 untunnelled packets destined to AFTR dropped :1
Software policy add errors            :0
Software policy delete errors         :0
Software policy memory alloc failures :0
Software Untunnelled packets ignored  :0
Software Misc errors
    DS-Lite ICMPv4 TTL exceed message process errors :0

```

SEE ALSO

[show services nat source mappings address-pooling-paired](#) | 859

[show services software statistics](#) | 1013

Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks

You can configure the DS-Lite per subnet limitation on MX Series routers equipped with MS-DPCs. Starting in Junos OS Release 18.2R1, MS-MPCs and MS-MICs also support the subnet limitation feature. Starting in Junos OS Release 20.2R1, the Next Gen Services MX-SPC3 security services card supports the subnet limitation feature.

Starting in Junos OS Release 19.2R1, MX Virtual Chassis and MX Broadband Network Gateway (BNG) routers also support the subnet limitation feature.

To configure DS-Lite per subnet session limitation:

1. Configure the size of the subnet prefix to which limiting is applied. Specify a prefix length of 56, 64, 96, or 128.

[edit]


```
user@host# set services service-set service-set-name software-options dslite-ipv6-prefix-length
dslite-ipv6-prefix-length
```

NOTE: Ensure that all mappings are cleared before changing the prefix length.

2. If you are using a next-hop service set on an AMS interface for DS-Lite, set the AMS inside interface's IPv6 source prefix length to the same value you use for the subnet prefix in Step 1.

```
[edit interfaces interface-name unit interface-unit-number load-balancing-options hash-keys]
user@host# set ipv6-source-prefix-length ipv6-source-prefix-length
```

3. Configure the maximum number of subscriber sessions allowed per prefix. You can configure from 0 through 16,384 sessions.

```
[edit]
user@host# set services software software-concentrator dslite dslite-concentrator-name session-limit-per-prefix
12
```

For Next Gen Services DS-Lite, MAP-E and V6rd softwires, configure the maximum number of subscriber sessions allowed per prefix:

```
[edit]
user@host# set services softwires software-types ds-lite | map-e | v6rd session-limit-per-prefix limit
```

NOTE: You cannot use **flow-limit** and **session-limit-per-prefix** in the same **dslite** configuration.

SEE ALSO

[DS-Lite Per Subnet Limitation Overview](#) | 180

[clear services nat mappings](#)

[software-options](#) | 697

[ds-lite](#) | 523

Protecting CGN Devices Against Denial of Service (DOS) Attacks

IN THIS SECTION

- [Mapping Refresh Behavior | 184](#)
- [EIF Inbound Flow Limit | 184](#)

You can now choose configuration options that help prevent or minimize the effect of attempted denial of service (DOS) attacks.

Mapping Refresh Behavior

Prior to the implementation of the new options for configuring NAT mapping refresh behavior, described in this topic, a conversation was kept alive when either inbound or outbound flows were active. This remains the default behavior. You can now also specify mapping refresh for only inbound flows or only outbound flows. To configure mapping refresh behavior, include the **mapping-refresh (inbound | outbound | inbound-outbound)** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.

EIF Inbound Flow Limit

Previously, the number of inbound connections on an EIF mapping was limited only by the maximum flows allowed on the system. You can now configure the number of inbound flows allowed for an EIF. To limit the number of inbound connections on an EIF mapping, include the **EIF-flow-limit *number-of-flows*** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.

Reducing Traffic and Bandwidth Requirements Using Port Control Protocol

IN THIS CHAPTER

- [Port Control Protocol Overview | 185](#)
- [Configuring Port Control Protocol | 188](#)

Port Control Protocol Overview

Port Control Protocol (PCP) provides a way to control the forwarding of incoming packets by upstream devices, such as NAT44 and firewall devices, and a way to reduce application keepalive traffic. PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICs. Starting in Junos OS Release 17.4R1, PCP for NAPT44 is also supported on the MS-MPC and MS-MIC. Starting in Junos 20.2R1, PCP for CGNAT DS-Lite services are supported for Next Gen Services. Starting in Junos OS Release 18.2R1, PCP on the MS-MPC and MS-MIC supports DS-Lite. In Junos OS Release 18.1 and earlier releases, PCP on the MS-MPC and MS-MIC does not support DS-Lite.

PCP is designed to be implemented in the context of both Carrier-Grade NATs (CGNs) and small NATs (for example, residential NATs). PCP enables hosts to operate servers for a long time (as in the case of a webcam) or a short time (for example, while playing a game or on a phone call) when behind a NAT device, including when behind a CGN operated by their ISP. PCP enables applications to create mappings from an external IP address and port to an internal IP address and port. These mappings are required for successful inbound communications destined to machines located behind a NAT or a firewall. After a mapping for incoming connections is created, remote computers must be informed about the IP address and port for the incoming connection. This is usually done in an application-specific manner.

Junos OS supports PCP version 2 and version 1.


PCP consists of the following components:

- PCP client—A host or gateway that issues PCP requests to a PCP server in order to obtain and control resources.
- PCP server—Typically a CGN gateway or co-located server that receives and processes PCP requests

Junos OS enables configuring PCP servers for mapping flows using NAPT44 capabilities such as port forwarding and port block allocation. Flows can be processed from these sources:

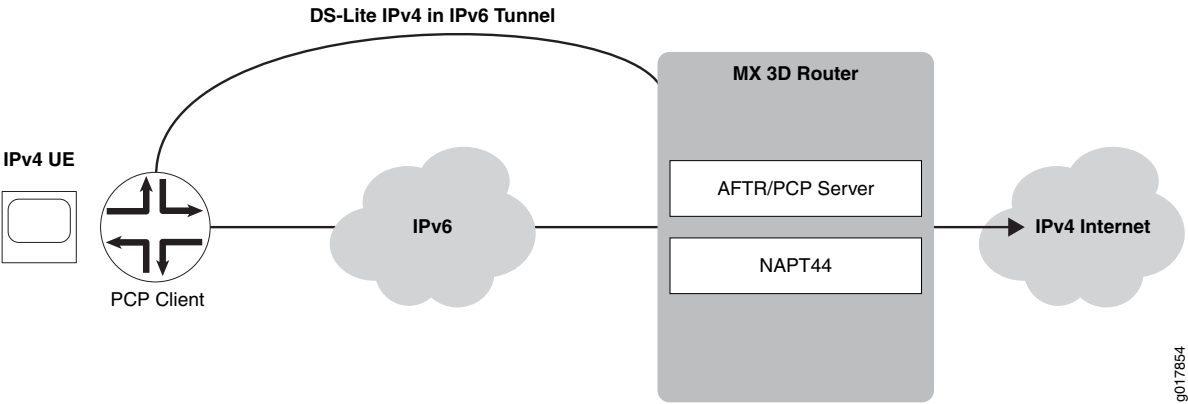
- Traffic containing PCP requests received directly from user equipment, as shown in [Figure 4 on page 186](#).

Figure 4: Basic PCP NAPT44 Topology



- Mapping of traffic containing PCP requests added by a router functioning as a DS-Lite software initiator (B4). This mode, known as *DS-Lite plain mode*, is shown in [Figure 5 on page 186](#).

Figure 5: PCP with DS-Lite Plain Mode



NOTE: Junos OS does not support deterministic port block allocation for PCP-originated traffic.

Benefits of Port Control Protocol

Many NAT-friendly applications send frequent application-level messages to ensure their sessions are not being timed out by a NAT device. PCP is used to:

- Reduce the frequency of these NAT keepalive messages
- Reduce bandwidth on the subscriber's access network
- Reduce traffic to the server
- Reduce battery consumption on mobile devices

Port Control Protocol Version 2

Starting with Junos OS Release 15.1, Port Control Protocol (PCP) version 2 is supported, which is in compliance with RFC 6887. PCP provides a way to control the forwarding of incoming packets by upstream devices, such as NAT44, and firewall devices, and a way to reduce application keep-alive traffic. PCP version 2 supports nonce authentication. PCP allows applications to create mappings from an external IP address and port to an internal IP address and port. A nonce payload prevents a replay attack and it is sent by default unless it is explicitly disabled.

Client nonce verification for version 2 map requests (for refresh or delete) requires that the nonce received in the original map request that causes the PCP mapping to be created is preserved. The version of the initial request that enables the mapping to be created is also preserved. This behavior of saving the nonce and version parameters denotes that 13 bytes per PCP mapping are used. This slight increase in storage space is not significant when matched with the current memory usage of a system for a single requested mapping (taking into account the endpoint-independent mapping (EIM) and endpoint-independent filtering (EIF) that are created along with it). In a customer deployment, PCP causes EIM and EIF mappings to represent a fraction of all such mappings.

Until Junos Release 15.1, services PICs support PCP servers on Juniper Networks routers in accordance with PCP draft version 22 with version 1 message encoding. With PCP being refined from the draft version as defined in *Port Control Protocol (PCP) draft-ietf-pcp-base-22 (July 2012 expiration)* to a finalized, standard version as defined in RFC 6887 -- Port Control Protocol (PCP), the message encoding changed to version 2 with the addition of a random nonce payload to authenticate peer and map requests as necessary. Version 1 does not decode messages compliant with version 2 format and nonce authentication is not supported. In a real-world network environment, with customer premises equipment (CPE) devices increasingly supporting version 2 only, it is required to parse and send version 2 messages. Backward compatibility with version 1-supporting CPE devices is maintained (version negotiation is part of the standard) and authenticates request nonce payload packets when v2 messages are in use.

The output of the **show services pcp statistics** command contains the PCP unsupported version field, which is incremented to indicate whenever the version is not 1 or 2. A new field, PCP request nonce does not match existing mapping, is introduced to indicate the number of PCP version 2 requests that were ignored because the nonce payload did not match the one recorded in the mapping (authentication failed). If version 2 is in use, the client nonce is used for authentication.

Release History Table

Release	Description
20.2R1	Starting in Junos 20.2R1, PCP for CGNAT DS-Lite services are supported for Next Gen Services.
18.2R1	Starting in Junos OS Release 18.2R1, PCP on the MS-MPC and MS-MIC supports DS-Lite.
17.4R1	Starting in Junos OS Release 17.4R1, PCP for NAPT44 is also supported on the MS-MPC and MS-MIC.
15.1	Starting with Junos OS Release 15.1, Port Control Protocol (PCP) version 2 is supported, which is in compliance with RFC 6887.

RELATED DOCUMENTATION

[Configuring Port Control Protocol](#) | [188](#)

Configuring Port Control Protocol

IN THIS SECTION

- [Configuring PCP Server Options](#) | [189](#)
- [Configuring a PCP Rule](#) | [191](#)
- [Configuring a NAT Rule](#) | [192](#)
- [Configuring a Service Set to Apply PCP](#) | [193](#)
- [SYSLOG Message Configuration](#) | [193](#)

This topic describes how to configure port control protocol (PCP). PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICs. Starting in Junos OS Release 17.4R1, PCP for NAPT44 is also supported on the MS-MPC and MS-MIC. Starting in Junos OS Release 18.2R1, PCP on the MS-MPC and MS-MIC supports DS-Lite. In Junos OS Release 18.1 and earlier releases, PCP on the MS-MPC and MS-MIC does not support DS-Lite. Starting in Junos OS release 20.2R1 PCP is supported on the MX-SPC3 security services card for CGNAT services.

Perform the following configuration tasks:

Configuring PCP Server Options

1. Specify a PCP server name.

```
user @host# edit services pcp server server-name
```

2. Set the IPv4 or IPv6 addresses of the server. For PCP DS-Lite, the **ipv6-address** must match the address of the AFTR (Address Family Transition Router or software concentrator).

NOTE: Starting in Junos OS Release 18.2R1, PCP on the MS-MPC and MS-MIC supports DS-Lite. In Junos OS Release 18.1 and earlier releases, PCP on the MS-MPC and MS-MIC does not support DS-Lite.

```
[edit services pcp server server-name]  
user @host# set ipv6-address ipv6-address
```

or

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

3. For PCP DS-Lite, provide the name of the DS-Lite software concentrator configuration.

```
[edit services pcp server server-name]  
user @host# set software-concentrator software-concentrator-name
```

4. Specify the minimum and maximum mapping lifetimes for the server.

```
[edit services pcp server server-name]  
user @host# set mapping-lifetime-minimum mapping-lifetime-min  
user @host# set mapping-lifetime-maximum mapping-lifetime-max
```

5. Specify the time limits for generating short lifetime or long lifetime errors.

```
[edit services pcp server server-name]
```

```
user @host# set short-lifetime-error short-lifetime-error
user @host# set long-lifetime-error long-lifetime-error
```

6. (Optional)—Enable PCP options on the specified PCP server. The following options are available—**third-party** and **prefer-failure**. The third-party option is required to enable third-party requests by the PCP client. DS-Lite requires the **third-party** option. The **prefer-failure** option requests generation of an error message when the PCP client requests a specific IP address/port that is not available, rather than assigning another available address from the NAT pool. If **prefer-failure** is not specified NAPT44 assigns an available address/port from the NAT pool based on the configured NAT options.

```
[edit services pcp server server-name]
user @host# set pcp-options third-party
user @host# set pcp-options prefer-failure
```

7. (Optional)—Specify which NAT pool to use for mapping.

```
[edit services pcp server server-name]
user @host# set nat-options pool-name1 <poolname2...>
```

NOTE: When you do not explicitly specify a NAT pool for mapping, the Junos OS performs a partial rule match based on source IP, source port, and protocol, and the Junos OS uses the NAT pool configured for the first matching rule to allocate mappings for PCP.

You *must* use explicit configuration in order to use multiple NAT pools.

For the MX-SPC3 security services card and Next Gen Services, the **nat-options** statement supports only one pool name to attach to a PCP server.

8. (Optional)—Configure the maximum number of mappings per client. The default is 32 and maximum is 128.

```
[edit services pcp server server-name]
user @host# set max-mappings-per-client max-mappings-per-client
```


Configuring a PCP Rule

A PCP rule has the same basic options as all service set rules:

- A **term** option that allows a single rule to have multiple applications.

A term is not required when running the MX-SPC3 security services card for Next Gen Services.

- A **from** option that identifies the traffic that is subject to the rule.
- A **then** option that identifies what action is to be taken. In the case of a PCP rule, this option identifies the pcsp server that handles selected traffic

1. Go to the **[edit services pcsp rule rule-name]** hierarchy level and specify **match-direction** input.

```
user @host# edit services pcsp rule rule-name
user @host# set match-direction input
```

2. Go to the **[edit services pcsp rule rule-name term term-name]** hierarchy level and provide a term name.

```
user @host# edit term term-name
```

This step is not required when running the MX-SPC3 security services card for Next Gen Services.

3. (Optional)—Provide a **from** option to filter the traffic to be selected for processing by the rule. When you omit the **from** option, all traffic handled by the service set's service interface is subject to the rule. The following options are available at the **[edit services pcsp rule rule-name term term-name from]** hierarchy level:

application-sets set-name —Traffic for the application set is processed by the PCP rule.

This step is not required when running the MX-SPC3 security services card for Next Gen Services.

applications [application-name]—Traffic for the application is processed by the PCP rule.

This option is not required when running the MX-SPC3 security services card for Next Gen Services.

destination-address address <except>—Traffic for the destination address or prefix is processed by the PCP rule. If you include the **except** option, traffic for the destination address or prefix is *not* processed by the PCP rule.

destination-address-range high maximum-value low minimum-value <except>—Traffic for the destination address range is processed by the PCP rule. If you include the **except** option, traffic for the destination address range is *not* processed by the PCP rule.

destination-port high maximum-value low minimum-value—Traffic for the destination port range is processed by the PCP rule.

destination-prefix-list *list-name* <except>—Traffic for a destination address in the prefix list is processed by the PCP rule. If you include the **except** option, traffic for a destination address in the prefix list is *not* processed by the PCP rule.

source-address *address* <except>—Traffic from the source address or prefix is processed by the PCP rule. If you include the **except** option, traffic from the source address or prefix is *not* processed by the PCP rule.

source-address-range high *maximum-value* low *minimum-value* <except>—Traffic from the source address range is processed by the PCP rule. If you include the **except** option, traffic from the source address range is *not* processed by the PCP rule.

source-prefix-list *list-name* <except>—Traffic from a source address in the prefix list is processed by the PCP rule. If you include the **except** option, traffic from a source address in the prefix list is *not* processed by the PCP rule.

4. Set the **then** option to identify the target PCP server.

```
[edit services pcp rule rule-name term term-name]
user@host# set then pcp-server server-name
```

Configuring a NAT Rule

To configure a NAT rule:

1. Configure the NAT rule name and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

2. Specify the NAT pool to use:

```
[edit services nat rule-name term term-name then translated]
user@host# set source-pool nat-pool-name
```

3. Configure the translation type.

```
[edit services nat rule-name term term-name then translated]
user@host# set translation-type translation-type
```

4. If you are using PCP with IPv4-to-IPv4 NAT or with DS-Lite, configure endpoint-independent mapping (EIM) and endpoint-independent filtering (EIF).

```
[edit services nat rule-name term term-name then translated]
user@host# set mapping-type endpoint-independent
user@host# set filtering-type endpoint-independent
```

NOTE: The PCP mappings are not created if you do not configure EIM and EIF with PCP for IPv4-to-IPv4 NAT or for DS-Lite.

Configuring a Service Set to Apply PCP

To use PCP, you must provide the rule name (or name of a list of rule names) in the **pcp-rule rule-name** option.

1. Go to the **[edit services service-set service-set-name]** hierarchy level.

```
user @host# edit services service-set service-set-name
```

2. If this is a new service set, provide basic service set information, including interface information and any other rules that may apply.
3. Specify the name of the PCP rule or rule list used to send traffic to the specified PCP server.

```
[edit services service-set service-set-name ]
user @host# set pcp-rule rule-name | rule-listname
```

NOTE: Your service set must also identify any required **nat-rule** and **softwire-rule**.

SYSLOG Message Configuration

A new syslog class, configuration option, **pcp-logs**, has been provided to control PCP log generation. It provides the following levels of logging:

- **protocol**—All logs related to mapping creation, deletion are included at this level of logging.

- **protocol-error**—All protocol error related logs (such as mapping refresh failed, PCP look up failed, mapping creation failed). are included in this level of logging.
- **system-error**—Memory and infrastructure errors are included in this level of logging.

SEE ALSO

| [Port Control Protocol Overview](#) | 185

Transitioning to IPv6 Using Mapping of Address and Port with Encapsulation (MAP-E)

IN THIS CHAPTER

- [Mapping of Address and Port with Encapsulation \(MAP-E\) for Next Gen Services | 195](#)
- [Equal Cost Multiple Path \(ECMP\) support for Mapping of Address and Port with Encapsulation \(MAP-E\) | 202](#)

Mapping of Address and Port with Encapsulation (MAP-E) for Next Gen Services

IN THIS SECTION

- [Understanding Mapping of Address and Port with Encapsulation \(MAP-E\) | 195](#)
- [Configuring Mapping of Address and Port with Encapsulation \(MAP-E\) for Next Gen Services | 198](#)

Understanding Mapping of Address and Port with Encapsulation (MAP-E)

IN THIS SECTION

- [Benefits of Mapping of Address and Port with Encapsulation \(MAP-E\) | 196](#)
- [Mapping of Address and Port with Encapsulation \(MAP-E\) Terminology | 196](#)
- [Mapping of Address and Port with Encapsulation \(MAP-E\) Functionality | 196](#)
- [Mapping of Address and Port with Encapsulation \(MAP-E\) Supported and Unsupported Features | 197](#)

This topic provides an overview of Mapping of Address and Port with Encapsulation (MAP-E) feature and its benefit to service providers when used as an inline service on MX Series routers with MPC and MIC interfaces. Starting in Junos OS release 20.2R1, MAP-E softwires are supported under Next Gen Services on either an MPC or MIC by specifying the inline services **si-1/1/0** naming convention. Starting in Junos OS release 20.3R1, MPC10E and MX2K-MPC11E support MAP-E.

Benefits of Mapping of Address and Port with Encapsulation (MAP-E)

Reduces administrative overhead and creates a scalable network infrastructure that easily supports connectivity to a large number of IPv4 subscribers over the ISP's IPv6 access network.

Mapping of Address and Port with Encapsulation (MAP-E) Terminology

Border Relay (BR)—MAP-E-enabled provider edge device in a MAP domain. A BR device has at least an IPv6-enabled interface and an IPv4 interface connected to the native IPv4 network.

MAP-E Customer Edge (CE)—MAP-E-enabled customer edge device in a MAP deployment.

MAP domain—One or more MAP-E CE devices and BR devices connected to the same virtual link.

Port Set ID (PSID)—Separate part of the transport layer port space that is denoted as port set ID.

Embedded Address (EA) Bits—EA-bits in the IPv6 address identify an IPv4 prefix or address or a shared IPv4 address and a port-set identifier.

Softwire—Tunnel between two IPv6 end-points to carry IPv4 packets or two IPv4 end-points to carry IPv6 packets.


Softwire Initiator (SI)—Softwire at the customer end that encapsulates native packets and tunnels them to a softwire concentrator at the service provider.

Softwire Concentrator (SC)—Softwire that decapsulates the packets received from a softwire initiator and sends them to their destination.

Mapping of Address and Port with Encapsulation (MAP-E) Functionality

Figure 6 on page 196 illustrates a simple MAP-E deployment scenario.

Figure 6: Sample MAP-E Deployment



In the MAP-E network topology, there are two MAP-E customer edge (CE) devices, each connected to a private IPv4 host. The MAP-E CE devices are dual stack and are capable of Network Address Port Translation (NAPT). The MAP-E CE devices connect to a MAP-E Border Relay (BR) device through an IPv6-only MAP-E network domain. The MAP-E BR device is dual stack and is connected to both a public IPv4 network and an IPv6 MAP-E network.

The MAP-E functionality is as follows:

1. The MAP-E CE devices are capable of NAPT. On receiving an IPv4 packet from the host, the MAP-E CE device performs NAT translation on the incoming IPv4 packets.
2. The NAT translated IPv4 packets are then encapsulated into IPv6 packets by the MAP-E CE device, and sent to the MAP-E BR device.
3. The IPv6 packet gets transported through the IPv6-only service provider network and reaches the MAP-E BR device.
4. On receiving the IPv6 packets, the incoming IPv6 packets are decapsulated by the MAP-E CE device and routed to the IPv4 public network.

In the reverse path, the incoming IPv4 packet is encapsulated into an IPv6 packet by the MAP-E BR device, and routed to the MAP-E CE devices.

Mapping of Address and Port with Encapsulation (MAP-E) Supported and Unsupported Features

Junos OS supports the following MAP-E features and functionality:

- MAP-E implementation supports line card throughput of 100 Gigabits.
- support for Inline MAP-E Border Relay (BR) solution that adheres to draft version 03 of RFC 7597
Fully compliant with draft version 03 of RFC 7597, *Mapping of Address and Port with Encapsulation (MAP)*, when the **version-3** option is disabled at the **services softwires software-types map-e map-e-concentrator-name**
- Support chassis-wide scale of 250 shared MAP-E rules.
- Support the feature on all MPCs using service interfaces with 100 Gigabits.
- Ability to ping MAP-E BR IPv6 address.
- Support only next-hop style of configuration for MAP-E.
- Support reassembly of fragmented IPv4 traffic arriving from IPv4 network before encapsulating it into an IPv6 packet.
- Support fragmentation of inner IPv4 packet if the packet size after encapsulation exceeds the MAP-E maximum transmission unit (MTU).

- Packets having Internet Control Message Protocol (ICMP) payload with the following message types are accepted for MAP-E encapsulation and decapsulation:
 - Echo or Echo Reply Message of type 0 and 8
 - Timestamp or Timestamp Reply Message of type 13 and 14
 - Information Request or Information Reply Message of type 15 and 16
 - Source quench, destination_unreachable, time_exceeded, icmp_redirect, icmp_address_mask_reply and parameter_problem errors
- Border Relay (BR) anycast is supported.

The following features and functionality are not supported with the MAP-E feature:

- Anti-spoof check is not supported for fragmented IPv4 packets coming from a customer edge (CE) device.
- Section 8.2 of the Internet draft draft-ietf-softwire-map-03 (expires on July 28, 2013), *Mapping of Address and Port with Encapsulation (MAP)* is not supported. Instead of responding with an ICMPv6 Destination Unreachable, Source address failed ingress/egress policy (Type 1, Code 5) message, spoof packets are silently dropped and the counter is incremented.
- IPv6 reassembly is not supported.
- ICMP v6-to-v4 translation at the BR is not supported.
- Inline MAP-E with virtual routing and forwarding (VRF) is not supported.
- Inline MAP-E with inline Network Address Translation (NAT) or dual stack (DS)-Lite is not supported.
- Interface-style MAP-E configuration is not supported.

Configuring Mapping of Address and Port with Encapsulation (MAP-E) for Next Gen Services

This example shows you how to configure the MAP-E Border Relay (BR) solution using a next hop-based style of configuration.

To configure MAP-E:

1. Create service interface on the device with 100g bandwidth support.

```
[edit chassis]
user@host# set fpc 0 pic 0 inline-services bandwidth 100g
```

2. Configure the dual stack service interface unit 0.

```
[edit interfaces]
user@host# set si-0/0/0 unit 0 family inet
```



```
user@host# set si-0/0/0 unit 0 family inet6
```

3. Configure service interface inside the dual stack domain.

```
[edit interfaces]
user@host# set si-0/0/0 unit 1 family inet
user@host# set si-0/0/0 unit 1 family inet family inet6
user@host# set si-0/0/0 unit 1 service-domain inside
```

4. Configure service interface outside the dual stack domain.

```
[edit interfaces]
user@host# set si-0/0/0 unit 2 family inet
user@host# set si-0/0/0 unit 2 family inet family inet6
user@host# set si-0/0/0 unit 2 service-domain outside
```

5. Configure the IPv4-facing interface on BR.

```
[edit interfaces]
user@host# set ge-0/2/7 unit 0 family inet address 10.10.10.1/16
```

6. Configure the CPE-facing interface on BR.

```
[edit interfaces]
user@host# set ge-0/2/8 unit 0 family inet6 address 3abc::1/16
```

7. Configure the MAP-E software concentrator and associated parameters.

a. (Optional) Configure MAPE version 3.

NOTE: For full RFC 7597 compliance do not configure MAP-E version 3.

b. Specify a name for MAP-E concentrator.

```
[edit]
user@host# edit services softwires software-types map-e mape-tun1
```

c. Specify the IPv6 address of the BR.

```
user@host# set br-address 2001:db8:ffff::1/128
```

d. Specify the rules for the MAP-E concentrator.

NOTE:

When configuring the MAP-E software concentrator, take the following into consideration:

- Possible values for **ea-bits-len** is 0 through 48.
- Possible values for **v4-prefix-len** is 0 through 32.
- If **v4-prefix-len** is 0 then **ea-bits-len** must be non-zero, and vice versa.
- It is possible that **ea-bits-len** is equal to 0, but **psid-len** is non-zero.
- If the sum of **v4-prefix-len** and **ea-bits-len** is less than 32, then the **psid-len** must be equal to the difference between 32 and the sum total of **v4-prefix-len** and **ea-bits-len**.
- The MAP-E IPv4 and IPv6 prefix must be unique per software concentrator.
- MAP-E PSID offset has a default value of 4, and MAP-E tunnel maximum transmission unit (MTU) has a default value of 9192.

i. Specify the rule length for the IPv4 and IPv6 prefixes.

```
user@host# edit services softwires software-types map-e mape-tun1
user@host# edit rule r1
[edit services softwires software-types map-e mape-tun1]
user@host# set rule r1 ipv4-prefix 192.0.2.0/24
```

```
user@host# set rule r1 ipv6-prefix 2001:db8:0000::/40
```

- ii. Configure the rule length for embedded addresses.

```
[edit services softwires software-types map-e mape-tun1]
user@host# set ea-bits-length 16
```

- iii. Configure the rule for the PSID offset.

```
[edit services softwires software-types map-e mape-tun1]
user@host# set psid-offset 4
```

- iv. Configure the rule for the PSID length.

```
[edit services softwires software-types map-e mape-tun1]
user@host# set psid-len 8
```

- v. Specify the MAP-E IPv6 tunnel MTU values.

```
[edit services softwires software-types map-e mape-tun1]
user@host# set mtu-v6 9192
user@host# set v4-reassembly
user@host# set v6-reassembly
```

- vi. Configure the software rule, which specifies the direction of the traffic to be tunneled through the MAP-E software.

```
[edit services softwires]
user@host# set rule-set domain-1 rule r1 then map-e map-e-dom-1
```

8. Configure the service-set for MAP-E.

```
[edit]
user@host# edit services service-set sset1
[edit services service-set sset1]
user@host# set softwires-rule-set domain-1
user@host# set next-hop-service inside-service-interface si-4/2/0.1
```

```
user@host# set next-hop-service outside-service-interface si-4/2/0.2
```

Release History Table

Release	Description
20.3R1	Starting in Junos OS release 20.3R1, MPC10E and MX2K-MPC11E support MAP-E.
20.2R1	Starting in Junos OS release 20.2R1, MAP-E softwires are supported under Next Gen Services on either an MPC or MIC by specifying the inline services si-1/1/0 naming convention.

Equal Cost Multiple Path (ECMP) support for Mapping of Address and Port with Encapsulation (MAP-E)

IN THIS SECTION

- [Equal Cost Multiple Path \(ECMP\) support for Mapping of Address and Port with Encapsulation \(MAP-E\) | 202](#)
- [Disabling auto-routes to support ECMP with Mapping of Address and Port with Encapsulation \(MAP-E\) | 203](#)

Equal Cost Multiple Path (ECMP) support for Mapping of Address and Port with Encapsulation (MAP-E)

This topic provides an overview of Equal Cost Multiple Path (ECMP) support for Mapping of Address and Port with Encapsulation (MAP-E) feature and its benefit to service providers when used as an inline service on MX Series routers with MPC and MIC interfaces.

In a MAP-E network topology, in the reverse path, the border relay router receives IPv4 traffic and encapsulates it in a IPv6 packet. Longer routes are used for faster matching. However, they do not facilitate EMCP load balancing on the PIC, as the routes point to a single PIC. Starting in 19.3R1, you can disable auto-routes by configuring the **disable-auto-route** statement at the **[edit services software software-concentrator map-e <domain-name>]** hierarchy, and direct the static routes to an ECMP load balancer. Hence, the packets can be distributed among different inline service interfaces.

Benefits

Enable load-balancing by distributing packets among different inline service interfaces.

Disabling auto-routes to support ECMP with Mapping of Address and Port with Encapsulation (MAP-E)

This example shows you how to disable auto-routes on a MAP-E Border Relay (BR) solution to support ECMP.

1. Create service interface on the device with 100g bandwidth support.

```
[edit chassis]
user@host# set fpc 0 pic 0 inline-services bandwidth 100g
```

2. Configure the dual stack service interface unit 0.

```
[edit interfaces]
user@host# set si-0/0/0 unit 0 family inet
user@host# set si-0/0/0 unit 0 family inet6
```

3. Configure service interface inside the dual stack domain.

```
[edit interfaces]
user@host# set si-0/0/0 unit 1 family inet
user@host# set si-0/0/0 unit 1 family inet family inet6
user@host# set si-0/0/0 unit 1 service-domain inside
```

4. Configure service interface outside the dual stack domain.

```
[edit interfaces]
user@host# set si-0/0/0 unit 2 family inet
user@host# set si-0/0/0 unit 2 family inet family inet6
user@host# set si-0/0/0 unit 2 service-domain outside
```

5. Configure the IPv4-facing interface on BR.

```
[edit interfaces]
user@host# set ge-0/2/7 unit 0 family inet address 10.10.10.1/16
```

6. Configure the CPE-facing interface on BR.

```
[edit interfaces]
```

```
user@host# set ge-0/2/8 unit 0 family inet6 address 3abc::1/16
```

7. Configure MAP-E domain 1 and associated parameters.

```
[edit services software software-concentrator]
user@host# set map-e mape-domain-1 version03
user@host# set map-e mape-domain-1 software-address 2001:db8:ffff::1
user@host# set map-e mape-domain-1 ipv4-prefix 192.0.2.0/24 mape-prefix 2001:db8::/32
user@host# set map-e mape-domain-1 ea-bits-len 16
user@host# set map-e mape-domain-1 psid-offset 4
user@host# set map-e mape-domain-1 psid-length 8
user@host# set map-e mape-domain-1 mtu-ipv6 9192
user@host# set map-e mape-domain-1 disable-auto-route
```

8. Configure MAP-E domain 2 and associated parameters.

```
[edit services software software-concentrator]
user@host# set map-e mape-domain-2 version03
user@host# set map-e mape-domain-2 software-address 2001:db8:ffff::1
user@host# set map-e mape-domain-2 ipv4-prefix 192.0.3.0/24 mape-prefix 2002:db8::/32
user@host# set map-e mape-domain-2 ea-bits-len 16
user@host# set map-e mape-domain-2 psid-offset 4
user@host# set map-e mape-domain-2 psid-length 8
user@host# set map-e mape-domain-2 mtu-ipv6 9192
user@host# set map-e mape-domain-2 disable-auto-route
```

9. Configure a software rule for MAP-E domain-1 to specify the direction of traffic to be tunneled.

```
[edit services software]
user@host# set rule sw-rule1 match-direction input term t1 then map-e mape-domain-1
```

10. Configure a software rule for MAP-E domain-2 to specify the direction of traffic to be tunneled.

```
[edit services software]
user@host# set rule sw-rule2 match-direction input term t1 then map-e mape-domain-2
```

11. Configure a single rule-set to combine both the rules.

```
[edit services software]
```

```

user@host# set rule-set ecmp-rules rule sw-rule1
user@host# set rule-set ecmp-rules rule sw-rule2

```

12. Configure the service set for MAP-E.

```

[edit services service-set]
user@host# set sset1 software-rule-sets ecmp-rules
user@host# set sset1 next-hop-service inside-service-interface si-0/0/0.1
user@host# set sset1 next-hop-service outside-service-interface si-0/0/0.2
user@host# set sset2 software-rule-sets ecmp-rules
user@host# set sset2 next-hop-service inside-service-interface si-0/1/0.1
user@host# set sset2 next-hop-service outside-service-interface si-0/1/0.2

```

13. Configure static routes for MAP-E BR IPv6 address.

```

[edit routing-options]
user@host# set rib inet6.0 static route 2001:db8:ffff::1/128 next-hop si-0/0/0.1
user@host# set rib inet6.0 static route 2001:db8:ffff::1/128 next-hop si-0/1/0.1
user@host# set rib inet.0 static route 192.0.2.0/24 next-hop si-0/0/0.2
user@host# set rib inet.0 static route 192.0.2.0/24 next-hop si-0/1/0.2
user@host# set rib inet.0 static route 192.0.3.0/24 next-hop si-0/0/0.2
user@host# set rib inet.0 static route 192.0.3.0/24 next-hop si-0/1/0.2

```

14. Enable load balancing.

```

[edit ]
user@host# set policy-options policy-statement LB then load-balance per-packet
user@host# set routing-options forwarding-table export LB

```

15. Verify the status of the routes.

```

[edit ]
user@host# run show route 2001:db8:ffff::1
inet6.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:ffff::1/128
    *[Static/5] 00:00:12
    > via si-1/0/0.1
    via si-1/1/0.1

```

The service sets of the PICs have *ecmp-rules* configured and they carry the MAP-E rules of domain-1 and domain-2. From the output, you can understand that when the **disable-auto-route** is enabled and *ecmp -rules* configured, instead of the longer auto routes, static routes are created.

RELATED DOCUMENTATION

| [map-e](#) | 583

Monitoring and Troubleshooting Softwires

IN THIS CHAPTER

- Ping and Traceroute for DS-Lite | 207
- Monitoring Software Statistics | 207
- Monitoring CGN, Stateful Firewall, and Software Flows | 209

Ping and Traceroute for DS-Lite

With Junos OS Release 11.4, you can use the **ping** and **traceroute** commands to determine the status of the DS-Lite software tunnels:

- IPv6 ping—The software address endpoint on the DS-Lite software terminator (AFTR) is usually configured only at the **[edit services software]** hierarchy level; it need not be hosted on any interface. Previous releases of the Junos OS software did not provide replies to pings to the IPv6 software address when the AFTR was not configured on a specific interface or loopback. An IPv6 ping enables the software initiator (B4) to verify the software address of the AFTR before creating a tunnel.
- IPv4 ping—A special IPv4 address, 192.0.0.1, is reserved for the AFTR. Previous releases of the Junos OS did not respond to any pings sent to this address. A B4 and other IPv4 nodes can now ping to this address to determine whether the DS-Lite tunnel is working.
- Traceroute—The AFTR now generates and forwards traceroute packets over the DS-Lite tunnel.

NOTE: No additional CLI configuration is necessary to use the new functionality.

Monitoring Software Statistics

Purpose

You can review software global statistics by using the **show services software** or **show services software statistics** command.

Action

user@host# **show services software**

```
Interface: sp-0/0/0, Service set: sset
Software Direction Flow count
2001:0:0:1::1 -> 1001::1 I 3
```

user@host# **show services software statistics**

```
DS-Lite Statistics:
Service PIC Name: :sp-0/0/0
Statistics
-----
Softwires Created :2
Softwires Deleted :1
Softwires Flows Created :2
Softwires Flows Deleted :1
Slow Path Packets Processed :2
Fast Path Packets Processed :274240
Fast Path Packets Encapsulated :583337
Rule Match Failed :0
Rule Match Succeeded :2
IPv6 Packets Fragmented :0
Transient Errors
-----
Flow Creation Failed - Retry :0
Slow Path Failed - Retry :0
Errors
-----
Software Creation Failed :0
Flow Creation Failed :0
Slow Path Failed :0
Packet not IPv4-in-IPv6 :0
IPv6 Fragmentation Error :0
Slow Path Failed - IPv6 Next Header Offset :0
Decapsulated Packet not IPv4 :0
Fast Path Failed - IPv6 Next Header Offset :0
No Software ID :0
No Flow Extension :0
Flow Limit Exceeded :0
6rd Statistics:
Service PIC Name :sp-0/0/0
```

```

Statistics
-----
Softwires Created :0
Softwires Deleted :0
Softwires Flows Created :0
Softwires Flows Deleted :0
Slow Path Packets Processed :0
Fast Path Packets Processed :0
Fast Path Packets Encapsulated :0
Rule Match Failed :0
Rule Match Succeeded :0
Transient Errors
-----
Flow Creation Failed - Retry :0
Slow Path Failed - Retry :0
Errors
-----
Softwire Creation Failed :0
Flow Creation Failed :0
Slow Path Failed :0
Packet not IPv6-in-IPv4 :0
Slow Path Failed - IPv6 Next Header Offset :0
Decapsulated Packet not IPv6 :0
Encapsulation Failed - No packet memory :0
No Softwire ID :0
No Flow Extension :0
ICMPv4 Dropped Packets :0

```

Monitoring CGN, Stateful Firewall, and Software Flows

Purpose

Use the following commands to check the creation of the softwires, pre-NAT flows, and post-NAT flows. Output can be filtered using more specific fields such as AFTR or B4 address or both for DS-Lite, and software-concentrator or software-initiator or both for 6rd.

- [show services stateful-firewall flows](#)
- [show services software flows](#)

Action

```
user@host# show services stateful-firewall flows
```

Interface: sp-0/1/0, Service set: dslite-svc-set2

Flow				State	Dir	Frm count
TCP	200.200.200.2:80	->	44.44.44.1:1025	Forward	O	219942
	NAT dest	44.44.44.1:1025	->	20.20.1.4:1025		
	Softwire	2001::2	->	1001::1		
TCP	20.20.1.2:1025	->	200.200.200.2:80	Forward	I	110244
	NAT source	20.20.1.2:1025	->	44.44.44.1:1024		
	Softwire	2001::2	->	1001::1		
TCP	200.200.200.2:80	->	44.44.44.1:1024	Forward	O	219140
	NAT dest	44.44.44.1:1024	->	20.20.1.2:1025		
	Softwire	2001::2	->	1001::1		
DS-LITE	2001::2	->	1001::1	Forward	I	988729
TCP	200.200.200.2:80	->	44.44.44.1:1026	Forward	O	218906
	NAT dest	44.44.44.1:1026	->	20.20.1.3:1025		
	Softwire	2001::2	->	1001::1		
TCP	20.20.1.3:1025	->	200.200.200.2:80	Forward	I	110303
	NAT source	20.20.1.3:1025	->	44.44.44.1:1026		
	Softwire	2001::2	->	1001::1		
TCP	20.20.1.4:1025	->	200.200.200.2:80	Forward	I	110944
	NAT source	20.20.1.4:1025	->	44.44.44.1:1025		
	Softwire	2001::2	->	1001::1		

RELATED DOCUMENTATION

Tunneling Services for IPv4-to-IPv6 Transition Overview

Port Forwarding Overview and Configuration

IN THIS CHAPTER

- [Port Forwarding for Next Gen Services | 211](#)

Port Forwarding for Next Gen Services

IN THIS SECTION

- [Port Forwarding Overview | 212](#)
- [Configuring Port Forwarding with Static Destination Address Translation for Next Gen Services | 212](#)
- [Configuring Port Forwarding without Static Destination Address Translation for Next Gen Services | 215](#)

Port Forwarding Overview

Port forwarding allows the public destination address and port of a packet to be translated to an IP address and port in a private network. This translation is a static, one-to-one mapping.

Port forwarding allows a packet to reach a host within a masqueraded, typically private, network, based on the port number on which the packet was received from the originating host. An example of this type of destination is the host of a public HTTP server within a private network.

If you only need to change the destination port, you can also configure port forwarding without translating the destination address.

Port forwarding is supported for destination NAT and twice NAT 44. Port forwarding works only with the FTP application-level gateway (ALG), and has no support for technologies that offer IPv6 services over IPv4 infrastructure, such as IPv6 rapid deployment (6rd) and dual-stack lite (DS-Lite).

Benefits

- Allows remote computers, such as public machines on the Internet, to connect to a non-standard port of a specific computer that is hidden within a private network.

Configuring Port Forwarding with Static Destination Address Translation for Next Gen Services

You can configure port forwarding with static destination address translation, which changes the destination address and port of a packet so it can reach the correct host and port within a masqueraded, typically private, network.

1. [Configuring the Destination Pool for Destination Address Translation | 212](#)
2. [Configuring the Mappings for Port Forwarding | 213](#)
3. [Configuring the NAT Rule for Port Forwarding with Destination Address Translation | 213](#)
4. [Configuring the Service Set for Port Forwarding with Destination Address Translation | 215](#)

Configuring the Destination Pool for Destination Address Translation

To configure the destination pool for the static destination address translation:

1. Create a destination pool.

```
user@host# edit services nat destination pool nat-pool-name
```

2. Define the addresses or subnets to which destination addresses are translated.

```
[edit services nat destination pool nat-pool-name]
user@host# set address address-prefix
```

3. To allow the IP addresses of a NAT destination pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the Mappings for Port Forwarding

1. Configure the port forwarding map name.

```
[edit services nat destination]
user@host# set port-forwarding map-name
```

2. Specify the original destination port number that needs to be translated and the port number to which the original port is mapped. You can configure a maximum of 32 destination port mappings in a port forwarding map.

```
[edit services nat destination port-forwarding map-name]
user@host# set destined-port port-id translated-port port-id
```

In the following example, the destination port number that needs to be translated is 23 and the port to which traffic is mapped is 45.

```
[edit services nat destination port-forwarding map1]
user@host# set destined-port 32 translated-port 45
```

Configuring the NAT Rule for Port Forwarding with Destination Address Translation

To configure the NAT rule for port forwarding with destination address translation:

1. Configure the NAT rule name.

```
[edit services destination source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the destination addresses that the NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

- Specify the destination port range that the NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-port low-port to high-port
```

- Specify the NAT pool that contains the destination addresses for translated traffic.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat pool nat-pool-name
```

- Specify the name of the mapping for port forwarding. You can only configure one mapping within a NAT rule term.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then port-forwarding-mappings map-name
```

- Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```


Configuring the Service Set for Port Forwarding with Destination Address Translation

To configure the service set for static destination NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

NOTE: You cannot use an AMS interface in a port forwarding service set.

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

Configuring Port Forwarding without Static Destination Address Translation for Next Gen Services

You can configure port forwarding without static destination address translation, which changes the destination port of a packet so it can reach the correct port on the destination host.

1. [Configuring the Mappings for Port Forwarding | 216](#)
2. [Configuring the NAT Rule for Port Forwarding without Destination Address Translation | 216](#)
3. [Configuring the Service Set for Port Forwarding without Destination Address Translation | 217](#)

Configuring the Mappings for Port Forwarding

1. Configure the port forwarding map name.

```
[edit services destination source]
user@host# set port-forwarding map-name
```

2. Specify the original destination port number that needs to be translated and the port number to which the original port is mapped. You can configure a maximum of 32 destination port mappings in a port forwarding map.

```
[edit services nat destination port-forwarding map-name]
user@host# set destined-port port-id translated-port port-id
```

In the following example, the destination port number that needs to be translated is 23 and the port to which traffic is mapped is 45.

```
[edit services nat destination port-forwarding map1]
user@host# set destined-port 32 translated-port 45
```

Configuring the NAT Rule for Port Forwarding without Destination Address Translation

To configure the NAT rule for port forwarding without destination address translation:

1. Configure the NAT rule name.

```
[edit services destination source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the destination addresses that the NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

4. Specify that there is no address translation for the rule.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat off
```

5. Specify the name of the mapping for port forwarding. You can only configure one mapping within a NAT rule term.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then port-forwarding-mappings map-name
```

6. Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Port Forwarding without Destination Address Translation

To configure the service set for static destination NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
```

```
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]  
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface  
                  interface-name
```

NOTE: You cannot use an AMS interface in a port forwarding service set.

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]  
user@host# set nat-rule-sets rule-set-name
```

Port Translation Features Overview and Configuration

IN THIS CHAPTER

- [Address Pooling and Endpoint Independent Mapping for Port Translation | 219](#)
- [Round-Robin Port Allocation | 221](#)
- [Secured Port Block Allocation for Port Translation | 222](#)

Address Pooling and Endpoint Independent Mapping for Port Translation

IN THIS SECTION

- [Address Pooling | 219](#)
- [Endpoint Independent Mapping and Endpoint Independent Filtering | 220](#)

Address Pooling

Address pooling, or address pooling paired (APP) ensures assignment of the same external IP address for all sessions originating from the same internal host. You can use this feature when assigning external IP addresses from a pool. This option does not affect port utilization.

Address pooling solves the problems of an application opening multiple connections. For example, when Session Initiation Protocol (SIP) client sends Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) packets, the SIP generally server requires that they come from the same IP address, even if they have been subject to NAT. If RTP and RTCP IP addresses are different, the receiving endpoint might drop packets. Any point-to-point (P2P) protocol that negotiates ports (assuming address stability) benefits from address pooling paired.

The following are use cases for address pooling:

- A site that offers instant messaging services requires that chat and their control sessions come from the same public source address. When the user signs on to chat, a control session authenticates the user.

A different session begins when the user starts a chat session. If the chat session originates from a source address that is different from the authentication session, the instant messaging server rejects the chat session, because it originates from an unauthorized address.

- Certain websites such as online banking sites require that all connections from a given host come from the same IP address.

NOTE: When you deactivate a service set that contains address pooling paired (APP) for that service set, messages are displayed on the PIC console and the mappings are cleared for that service set. These messages are triggered when the deletion of a service-set commences and again generated when the deletion of the service set is completed. The following sample messages are displayed when deletion starts and ends:

- **Nov 15 08:33:13.974 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion initiated**
- **Nov 15 08:33:14.674 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion completed**

In a scaled environment that contains a large number of APP in a service set, a heavy volume of messages is generated and this process takes some amount of time. We recommend that you wait until the console messages indicating the completion of deletion of the service set are completed before you reactivate the service-set again.

Endpoint Independent Mapping and Endpoint Independent Filtering

Endpoint independent mapping (EIM) ensures the assignment of the same external address *and* port for all connections from a given host if they use the same internal port. This means if they come from a different source port, you are free to assign a different external address.

EIM and APP differ as follows:

- APP ensures assigning the same external IP address.
- EIM provides a stable external IP address and port (for a period of time) to which external hosts can connect. Endpoint independent filtering (EIF) controls which external hosts can connect to an internal host.

NOTE: When you deactivate a service set that contains endpoint independent mapping (EIM) mapping for that service set, messages are displayed on the PIC console and the mappings are cleared for that service set. These messages are triggered when the deletion of a service set commences and again generated when the deletion of the service set is completed. The following sample messages are displayed when deletion starts and ends:

- Nov 15 08:33:13.974 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion initiated
- Nov 15 08:33:14.674 LOG: Critical] SVC-SET ss1 (iid 5) deactivate/delete: NAT Mappings and flows deletion completed

In a scaled environment that contains a large number of EIM mappings in a service set, a heavy volume of messages is generated and this process takes some amount of time. We recommend that you wait until the console messages indicating the completion of deletion of the service set are completed before you reactivate the service-set again.

Round-Robin Port Allocation

Round-robin allocation is one method you can configure to allocate private addresses to external addresses and ports. Round-robin allocation assigns one port from each external address in a range before repeating the process for each address in the next range. After ports have been allocated for all addresses in the last range, the allocation process wraps around and allocates the next unused port for addresses in the first range. For example, if you have a NAT pool range of 100.0.0.1 through 100.0.0.12 and the first port is 3333:

- The first connection is allocated to the address:port 100.0.0.1:3333.
- The second connection is allocated to the address:port 100.0.0.2:3333.
- The third connection is allocated to the address:port 100.0.0.3:3333.
- The fourth connection is allocated to the address:port 100.0.0.4:3333.
- The fifth connection is allocated to the address:port 100.0.0.5:3333.
- The sixth connection is allocated to the address:port 100.0.0.6:3333.
- The seventh connection is allocated to the address:port 100.0.0.7:3333.
- The eighth connection is allocated to the address:port 100.0.0.8:3333.
- The ninth connection is allocated to the address:port 100.0.0.9:3333.
- The tenth connection is allocated to the address:port 100.0.0.10:3333.

- The eleventh connection is allocated to the address:port 100.0.0.11:3333.
- The twelfth connection is allocated to the address:port 100.0.0.12:3333.
- Wraparound occurs and the thirteenth connection is allocated to the address:port 100.0.0.1:3334.

Secured Port Block Allocation for Port Translation

You can configure secured port block allocation, which allocates blocks of ports to a subscriber for source NAT port translation. The most recently allocated block is the current active block. New requests for NAT ports for the subscriber are served from the active block. Ports are allocated randomly from the current active block.

Carriers track subscribers using the IP address (RADIUS or DHCP) log. If they use port translation without port block allocation, an IP address is shared by multiple subscribers, and the carrier must track the IP address and port, which are part of the NAT log. Because ports are used and reused at a very high rate, tracking subscribers using the log becomes difficult because of the large number of messages, which are difficult to archive and correlate. By using port block allocation, you can significantly reduce the number of logs, making it easier to track subscribers.

With port block allocation, we generate one syslog log per set of ports allocated for a subscriber. These logs are UDP based and can be lost in the network, particularly for long-running flows. You can configure an interim logging interval to re-send logs for active blocks that have traffic on at least one of the ports.

Static Source NAT Overview and Configuration

IN THIS CHAPTER

- [Static Source NAT Overview | 223](#)
- [Configuring Static Source NAT44 or NAT66 for Next Gen Services | 224](#)

Static Source NAT Overview

IN THIS SECTION

- [Benefits | 224](#)

Static source NAT performs a one-to-one static mapping of the original private domain host source address to a public source address. A block of external addresses is set aside for this mapping, and source addresses are translated as hosts in a private domain originate sessions to the external domain. Static source NAT does not perform port mapping. For packets outbound from the private network, static source NAT translates source IP addresses and related fields such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, static source NAT translates the destination IP address and the checksums.

Benefits

- Allows hosts in the private network to connect with the external domain, while hiding the private network.

Configuring Static Source NAT44 or NAT66 for Next Gen Services

1. [Configuring the Source Pool for Static Source NAT44 or NAT66 | 224](#)
2. [Configuring the NAT Rule for Static Source NAT44 or NAT66 | 225](#)
3. [Configuring the Service Set for Static Source NAT44 or NAT66 | 226](#)

Configuring the Source Pool for Static Source NAT44 or NAT66

To configure the source pool for static source NAT44 or NAT66:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. Configure a one-to-one static shifting of a range of original source addresses to the range of addresses in the source pool by specifying the base address of the original source address range.

```
[edit services nat source pool nat-pool-name]  
user@host# set host-address-base ip-address
```

For example, if the host address base is 198.51.100.30 and the NAT pool uses the range 203.0.113.10 to 203.0.113.20, then 198.51.100.30 translates to 203.0.113.10, 198.51.100.31 translates to 203.0.113.11, and so on.

4. To allow the IP addresses of a NAT source pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Rule for Static Source NAT44 or NAT66

To configure the NAT source rule for static source NAT44 or NAT66 :

1. Configure the NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address any-unicast
```

- Specify one or more application protocols to which the NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

- Specify the NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

- Configure the address-pooling paired feature if you want to ensure assignment of the same external IP address for all sessions originating from the same internal host.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat mapping-type]
user@host# set address-pooling-paired
```

- Specify the timeout period for **address-pooling-paired** mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Static Source NAT44 or NAT66

To configure the service set for static source NAT44 or NAT66:

- Define the service set.

```
[edit services]
```

```
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]  
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]  
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface  
interface-name
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]  
user@host# set nat-rule-sets rule-set-name
```

SEE ALSO

| [Static Source NAT Overview](#) | 223

Static Destination NAT Overview and Configuration

IN THIS CHAPTER

- [Static Destination NAT Overview | 228](#)
- [Configuring Static Destination NAT for Next Gen Services | 229](#)

Static Destination NAT Overview

Static destination NAT translates the IPv4 destination address of an incoming packet to the IPv4 address of a private server. This redirects traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

Static destination NAT uses a one-to-one mapping between the original address and the translated address; the mapping is configured statically.

You can also statically translate the destination port by using port forwarding. See [“Port Forwarding for Next Gen Services” on page 211](#).

Benefits of Static Destination NAT

- Allows external traffic to communicate with a private host without revealing the host's private IP address
- Does not require port mapping

RELATED DOCUMENTATION

| [Configuring Static Destination NAT for Next Gen Services | 229](#)

Configuring Static Destination NAT for Next Gen Services

1. [Configuring the Destination Pool for Static Destination NAT | 229](#)
2. [Configuring the NAT Rule for Static Destination NAT | 229](#)
3. [Configuring the Service Set for Static Destination NAT | 231](#)

Configuring the Destination Pool for Static Destination NAT

To configure the destination pool for static destination NAT:

1. Create a destination pool.

```
user@host# edit services nat destination pool nat-pool-name
```

2. Define the addresses or subnets to which destination addresses are translated.

```
[edit services nat destination pool nat-pool-name]  
user@host# set address address-prefix
```

3. To allow the IP addresses of a NAT destination pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]  
user@host# set allow-overlapping-pools
```

Configuring the NAT Rule for Static Destination NAT

To configure the NAT rule for static destination NAT:

1. Configure the NAT rule name.

```
[edit services nat destination]  
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the destination NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]  
user@host# set match-direction (in | out | in-out)
```

3. Specify the source addresses of traffic that the NAT rule applies to.

To specify one address or prefix value:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match source-address any-unicast
```

4. Specify the destination addresses that the NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

5. Specify one or more application protocols to which the destination NAT rule applies. The number of applications listed in the rule must not exceed 3072.


```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

- Specify the NAT pool that contains the destination addresses for translated traffic.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat pool nat-pool-name
```

- Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Static Destination NAT

To configure the service set for static destination NAT:

- Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

- Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

- Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

RELATED DOCUMENTATION

| [Static Destination NAT Overview](#) | 228

Twice NAT Overview and Configuration

IN THIS CHAPTER

- Twice NAT Overview | 233
- Configuring Twice NAT for Next Gen Services | 234

Twice NAT Overview

Twice NAT translates both the source and destination IP addresses.

The private source address is translated by dynamically assigning a public address from a pool and a port number. Multiple private IP addresses can be mapped to the same external address because each private address is mapped to a different port of the external address.

The destination address is translated to the IPv4 address of a private server. This redirects traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address). The destination addresses is translated with a one-to-one static mapping to an address in a pool. Port mapping is not performed for the destination address.

You can also statically translate the destination port by using port forwarding. See [“Port Forwarding for Next Gen Services” on page 211](#).

Benefits

- Allows hosts in the private network to connect with the external domain, while hiding the private network.
- Minimizes the number of public IP addresses that are allocated for NAT.
- Allows external traffic to communicate with a private host without revealing the host’s private IP address

Configuring Twice NAT for Next Gen Services

IN THIS SECTION

- [Configuring the Source and Destination Pools for Twice NAT | 234](#)
- [Configuring the NAT Rules for Twice NAT | 238](#)
- [Configuring the Service Set for Twice NAT | 241](#)

Configuring the Source and Destination Pools for Twice NAT

To configure the source and destination pools for twice NAT:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. To configure automatic port assignment, specify either random allocation or round-robin allocation.

```
[edit services nat source pool nat-pool-name port]  
user@host# set automatic (random-allocation | round-robin)
```

Random allocation randomly assigns a port from the range 1024 through 65535 for each port translation. Round robin allocation first assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

4. To disable round-robin port allocation for all NAT pools that do not specify an **automatic** (**random-allocation** | **round-robin**) setting, configure the global setting.

```
[edit services nat source]
user@host# set port-round-robin disable
```

5. To configure a range of ports to assign to a pool, perform the following:

NOTE: If you specify a range of ports to assign, the **automatic** statement is ignored.

- a. Specify the low and high values for the port. If you do not configure automatic port assignment, you must configure a range of ports.

```
[edit services nat source pool nat-pool-name port]
user@host# set range port-low to port-high
```

- b. Specify either random allocation or round-robin allocation. Round-robin allocation is the default.

```
[edit services nat source pool nat-pool-name port range]
user@host# set (random-allocation | round-robin)
```

6. Assign a port within the same range as the incoming port—either 0 through 1023 or 1024 through 65,535. This feature is not available if you configure port-block allocation.

```
[edit services nat source pool nat-pool-name port]
user@host# set preserve-range
```

7. Assign a port with the same parity (even or odd) as the incoming port. This feature is not available if you configure port-block allocation.

```
[edit services nat source pool nat-pool-name port]
user@host# set preserve-parity
```

8. Configure a global default port range for NAT pools that use port translation. This port range is used when a NAT pool does not specify a port range and does not specify automatic port assignment. The global port range can be from 1024 through 65,535.

```
[edit services nat source]
user@host# set pool-default-port-range port-low to port-high
```

9. If you want to allocate a block of ports for each subscriber to use for NAT, configure port-block allocation:

- a. Configure the number of ports in a block. The range is 1 through 64,512 and the default is 128.

```
[edit services nat source pool nat-pool-name port]
user@host# set block-allocation block-size block-size
```

- b. Configure the interval, in seconds, for which the block is active. After the timeout, a new block is allocated, even if ports are available in the active block. If you set the timeout to 0, port blocks are filled completely before a new port block is allocated, and the last port block remains active indefinitely. The range is 0 through 86,400, and the default is 0.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set active-block-timeout timeout-interval
```

- c. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

- d. Configure the maximum number of blocks that can be allocated to a user address. The range is 1 through 512, and the default is 8.

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set maximum-blocks-per-host maximum-block-number
```

- e. Specify how often to send interim system logs for active port blocks and for inactive port blocks with live sessions. This increases the reliability of system logs, which are UDP-based and can get lost in the network. The range is 1800 through 86,400 seconds, and the default is 0 (interim logs are disabled).

```
[edit services nat source pool nat-pool-name port block-allocation]
user@host# set interim-logging-interval timeout-interval
```

10. Specify the timeout period for endpoint independent translations that use the specified NAT pool. Mappings that are inactive for this amount of time are dropped. The range is 120 through 86,400 seconds. If you do not configure **ei-mapping-timeout**, then the **mapping-timeout** value is used for endpoint independent translations.

```
[edit services nat source pool nat-pool-name]
user@host# set ei-mapping-timeout ei-mapping-timeout
```

11. Specify the timeout period for address-pooling paired mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

12. Define the NAT pool utilization levels that trigger SNMP traps. The **raise-threshold** is the pool utilization percentage that triggers the trap, and the range is 50 through 100. The **clear-threshold** is the pool utilization percentage that clears the trap, and the range is 40 through 100. For pools that use port-block allocation, the utilization is based on the number of ports that are used; for pools that do not use port-block allocation, the utilization is based on the number of addresses that are used.

```
[edit services nat source pool nat-pool-name]
user@host# set pool-utilization-alarm raise-threshold value
user@host# set pool-utilization-alarm clear-threshold value
```

If you do not configure **pool-utilization-alarm**, traps are not created.

13. Create a destination pool. Do not use the same name that you used for the source pool.

```
user@host# edit services nat destination pool nat-pool-name
```

14. Define the addresses or subnets to which destination addresses are translated.

```
[edit services nat destination pool nat-pool-name]
user@host# set address address-prefix
```

15. To allow the IP addresses of a NAT source pool or destination pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**. However, pools that configure port-block allocation must not overlap with other pools.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Rules for Twice NAPT

To configure the source and destination NAT rules for twice NAPT:

1. Configure the source NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:


```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address any-unicast
```

4. Specify one or more application protocols to which the NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

5. Specify the NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

6. If you want to ensure that the same external address and port are assigned to all connections from a given host, configure endpoint-independent mapping:

- a. Configure the mapping type as endpoint independent.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set mapping-type endpoint-independent
```

- b. Specify prefix lists that contain the hosts that are allowed to establish inbound connections using the endpoint-independent mapping. (Prefix lists are configured at the **[edit policy-options]** hierarchy level.)

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set filtering-type endpoint-independent prefix-list [allowed-host] except [denied-host]
```

- c. Specify the maximum number of inbound flows allowed simultaneously on an endpoint-independent mapping.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set secure-nat-mapping eif-flow-limit number-of-flows
```

- d. Specify the direction in which active endpoint-independent mapping is refreshed. By default, mapping is refreshed for both inbound and outbound active flows.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat]
user@host# set secure-nat-mapping mapping-refresh (inbound | inbound-outbound | outbound)
```

7. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

8. Configure the destination NAT rule name.

```
[edit services nat destination]
user@host# set rule-set rule-set-name rule rule-name
```

9. Specify the traffic direction to which the destination NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

10. Specify the destination addresses of traffic that the destination NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

11. Specify one or more application protocols to which the destination NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

12. Specify the destination NAT pool that contains the destination addresses for translated traffic.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat pool nat-pool-name
```

13. Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Twice NAT

To configure the service set for twice NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

3. Specify the NAT rule sets to be used with the service set. Include the source NAT rule set and the destination NAT rule set.

```
[edit services service-set service-set-name]
```

```
user@host# set nat-rule-sets rule-set-name
```

Twice NAT Overview and Configuration

IN THIS CHAPTER

- [Twice Static NAT Overview | 243](#)
- [Configuring Twice Static NAT44 for Next Gen Services | 244](#)
- [Twice Dynamic NAT Overview | 248](#)
- [Configuring Twice Dynamic NAT for Next Gen Services | 249](#)

Twice Static NAT Overview

IN THIS SECTION

- [Benefits | 243](#)

Twice static NAT translates both the source and destination IP addresses. An addresses is translated with a one-to-one static mapping to an address in a pool. Port mapping is not performed.

The original private domain host source address is translated to a public source address.

The destination address is translated to the IPv4 address of a private server. This redirects traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

Benefits

- Allows hosts in the private network to connect with the external domain, while hiding the private network.
- Hides a private network
- Allows external traffic to communicate with a private host without revealing the host's private IP address
- Does not require port mapping

Configuring Twice Static NAT44 for Next Gen Services

IN THIS SECTION

- [Configuring the Source and Destination Pools for Twice Static NAT44 | 244](#)
- [Configuring the NAT Rules for Twice Static NAT44 | 245](#)
- [Configuring the Service Set for Twice Static NAT44 | 247](#)

Configuring the Source and Destination Pools for Twice Static NAT44

To configure the source and destination pools for twice static NAT44:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. Configure a one-to-one static shifting of a range of original source addresses to the range of addresses in the source pool by specifying the base address of the original source address range.

```
[edit services nat source pool nat-pool-name]  
user@host# set host-address-base ip-address
```

For example, if the host address base is 198.51.100.30 and the NAT pool uses the range 203.0.113.10 to 203.0.113.20, then 198.51.100.30 translates to 203.0.113.10, 198.51.100.31 translates to 203.0.113.11, and so on.

4. Create a destination pool. Do not use the same name that you used for the source pool.

```
user@host# edit services nat destination pool nat-pool-name
```

5. Define the addresses or subnets to which destination addresses are translated.

```
[edit services nat destination pool nat-pool-name]
user@host# set address address-prefix
```

6. To allow the IP addresses of a NAT pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Rules for Twice Static NAT44

To configure the source and destination NAT rules for twice static NAT44:

1. Configure the source NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
```

```
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address any-unicast
```

4. Specify one or more application protocols to which the source NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

5. Specify the source NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

6. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

7. Configure the destination NAT rule name.

```
[edit services nat destination]
user@host# set rule-set rule-set-name rule rule-name
```

8. Specify the traffic direction to which the destination NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

9. Specify the destination addresses of traffic that the destination NAT rule applies to.


```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

10. Specify one or more application protocols to which the destination NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

11. Specify the destination NAT pool that contains the destination addresses for translated traffic.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat pool nat-pool-name
```

12. Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Twice Static NAT44

To configure the service set for twice static NAT44:

1. Define the service set.

```
[edit services]
```

```
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]  
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]  
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface  
interface-name
```

3. Specify the NAT rule sets to be used with the service set. Include the source NAT rule set and the destination NAT rule set.

```
[edit services service-set service-set-name]  
user@host# set nat-rule-sets rule-set-name
```

Twice Dynamic NAT Overview

Twice dynamic NAT translates both the source and destination IP addresses. Port mapping is not performed.

The private source address is translated by dynamically assigning a public address from a pool, and the mapping from the original source address to the translated source address is maintained as long as there is at least one active flow that uses this mapping.

The destination address is translated to the IPv4 address of a private server. This redirects traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address). The destination addresses is translated with a one-to-one static mapping to an address in a pool.

Benefits

- Allows hosts in the private network to connect with the external domain, while hiding the private network.
- Allows a few public IP addresses to be used by several private hosts

- Allows external traffic to communicate with a private host without revealing the host's private IP address
- Does not require port mapping

Configuring Twice Dynamic NAT for Next Gen Services

IN THIS SECTION

- [Configuring the Source and Destination Pools for Twice Dynamic NAT | 249](#)
- [Configuring the NAT Rules for Twice Dynamic NAT | 250](#)
- [Configuring the Service Set for Twice Dynamic NAT | 253](#)

Configuring the Source and Destination Pools for Twice Dynamic NAT

To configure the source and destination pools for twice dynamic NAT:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. Disable port translation.

```
[edit services nat destination pool nat-pool-name]  
user@host# set port no-translation
```

4. Define the NAT pool utilization levels that trigger SNMP traps. The **raise-threshold** is the pool utilization percentage that triggers the trap, and the range is 50 through 100. The **clear-threshold** is the pool utilization percentage that clears the trap, and the range is 40 through 100. The utilization is based on the number of addresses that are used.

```
[edit services nat source pool nat-pool-name]
user@host# set pool-utilization-alarm raise-threshold value
user@host# set pool-utilization-alarm clear-threshold value
```

If you do not configure **pool-utilization-alarm**, traps are not created.

5. Create a destination pool. Do not use the same name that you used for the source pool.

```
user@host# edit services nat destination pool nat-pool-name
```

6. Define the addresses or subnets to which destination addresses are translated.

```
[edit services nat destination pool nat-pool-name]
user@host# set address address-prefix
```

7. To allow the IP addresses of a NAT source pool or destination pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Rules for Twice Dynamic NAT

To configure the source and destination NAT rules for twice dynamic NAT:

1. Configure the source NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

To specify any unicast address:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address any-unicast
```

4. Specify one or more application protocols to which the source NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

5. Configure the address-pooling paired feature if you want to ensure assignment of the same external IP address for all sessions originating from the same internal host.

```
[edit services nat source rule-set rule-set-name rule rule-name then source-nat mapping-type]
user@host# set address-pooling-paired
```

6. Specify the timeout period for **address-pooling-paired** mappings that use the NAT pool. The range is 120 through 86,400 seconds, and the default is 300. Mappings that are inactive for this amount of time are dropped.

```
[edit services nat source pool nat-pool-name]
user@host# set mapping-timeout mapping-timeout
```

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

7. Specify the source NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

8. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

9. Configure the destination NAT rule name.

```
[edit services nat destination]
user@host# set rule-set rule-set-name rule rule-name
```

10. Specify the traffic direction to which the destination NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

11. Specify the destination addresses of traffic that the destination NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

12. Specify one or more application protocols to which the destination NAT rule applies. The number of applications listed in the rule must not exceed 3072.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match application [application-name]
```

13. Specify the destination NAT pool that contains the destination addresses for translated traffic.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat pool nat-pool-name
```

14. Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Twice Dynamic NAT

To configure the service set for twice dynamic NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service, which requires a single service interface, or a next-hop service, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

3. Specify the NAT rule sets to be used with the service set. Include the source NAT rule set and the destination NAT rule set.

```
[edit services service-set service-set-name]  
user@host# set nat-rule-sets rule-set-name
```


Class of Service Overview and Configuration

IN THIS CHAPTER

- [Class of Service for Services PICs \(Next Gen Services\) | 255](#)

Class of Service for Services PICs (Next Gen Services)

IN THIS SECTION

- [Class of Service Overview for Services PICs \(Next Gen Services\) | 255](#)
- [Configuring CoS for Traffic Processed by a Services PIC \(Next Gen Services\) | 256](#)

Class of Service Overview for Services PICs (Next Gen Services)

You can configure CoS Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignment for packets transiting a services PIC while being processed by a service set.

Configure services CoS rules, which identify the matching conditions for packet source and destination addresses and for packet applications, and the actions to take on those packets. You must apply CoS rules to a service set before the rules can be applied to traffic. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

You can also configure specific CoS actions for FTP and for SIP traffic by creating an application profile. The application profile can then be referenced in the CoS rule actions.

The services CoS rules do not support scheduling. You must configure scheduling at the **[edit class-of-service]** hierarchy level on the output interface or fabric.

NOTE: When configuring Next Gen Services with the MX-SPC3 services card, the **service set** must include at least one stateful firewall (SFW) rule or NAT rule, or services CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a **service set**. CoS works without NAT and SFW rules also.

Benefits

CoS for traffic on a services PIC lets you classify traffic flows based on stateful firewall and NAT configurations.

SEE ALSO

[Configuring CoS for Traffic Processed by a Services PIC \(Next Gen Services\) | 256](#)

Configuring CoS for Traffic Processed by a Services PIC (Next Gen Services)

IN THIS SECTION

- [Configuring CoS Rules | 256](#)
- [Configuring Application Profiles for CoS Rules | 259](#)
- [Configuring CoS Rule Sets | 260](#)
- [Configuring the Service Set for CoS | 261](#)

Configuring CoS Rules

1. Configure a name for the CoS rule.

```
user@host# edit services cos rule rule-name
```

2. Specify the traffic flow direction for the CoS rule.

```
[edit services cos rule rule-name]
user@host# set match-direction (input | input-output | output)
```

If this CoS rule is applied to an interface-type service set, the direction is determined by whether a packet is entering or leaving the interface on which the service set is applied. If this CoS rule is applied to a next-hop service set, the direction is input if the inside interface is used to route the packet, and the direction is output if the outside interface is used to route the package.

If you configure **input-output**, the rule is applied to sessions initiated from either direction.

3. Configure a name for a CoS rule policy.

```
[edit services cos rule rule-name]
user@host# set policy policy-name
```

You can configure multiple policies for a CoS rule. Each policy identifies the matching conditions for packet source and destination addresses and for packet applications, and the CoS actions to take on those packets. Once a policy in the rule matches a packet, that policy is applied and no other policies in the rule are processed.

4. Specify one or more port-based applications that match the policy.

```
[edit services cos rule rule-name policy policy-name]
user@host# set match application [application-names]
```

5. Specify the destination address that matches the policy.

```
[edit services cos rule rule-name policy policy-name]
user@host# set match destination-address address
```

6. Specify a range of destination addresses that match the policy.

```
[edit services cos rule rule-name policy policy-name]
user@host# set match destination-address-range low minimum-value high maximum-value
```

7. Specify the destination port number that matches the policy.

```
[edit services cos rule rule-name policy policy-name]
user@host# set match destination-port port-number
```

8. Specify the source address that matches the policy.

```
[edit services cos rule rule-name policy policy-name]
user@host# set match source-address address
```

9. Specify a range of source addresses that match the policy.

```
[edit services cos rule rule-name policy policy-name]
user@host# set match source-address-range low minimum-value high maximum-value
```

10. Specify a prefix list of source address prefixes that match the policy.

```
[edit services cos rule rule-name policy policy-name]
user@host# set match source-prefix-list list-name
```

You configure a prefix list by using the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

11. Specify the application profile that defines the CoS policy actions for FTP and SIP traffic.

```
[edit services cos rule rule-name policy policy-name]
user@host# set then application-profile profile-name
```

12. Specify the DSCP value to apply to the packet.

```
[edit services cos rule rule-name policy policy-name]
user@host# set then dscp (alias | bits)
```

The DSCP can be either a code point alias or a DSCP bit value.

13. Specify the forwarding class name to apply to the packet.

```
[edit services cos rule rule-name policy policy-name]
user@host# set then forwarding-class class-name
```

The choices are:

- assured-forwarding
- best-effort
- expedited-forwarding
- network-control
- user-defined classifiers.

You can define classifiers under [edit class-of-service classifiers dscp] hierarchy.

14. Configure system logging for the CoS rule policy.

15. Specify the treatment of flows in the reverse direction of the matching direction. Perform only one of the following:

- Configure unique values for the reverse direction:

```
[edit services cos rule rule-name policy policy-name]
user@host# set then reverse application-profile profile-name
user@host# set then reverse dscp (alias | bits)
user@host# set then reverse forwarding-class class-name
```

- Apply the CoS rule policy actions to flows in the reverse direction as well as to flows in the matching direction.

```
[edit services cos rule rule-name policy policy-name]
user@host# set then reflexive
```

- Store the DSCP and forwarding class of a packet that is received in the match direction of the rule and then apply that DSCP and forwarding class to packets that are received in the reverse direction of the same session.

```
[edit services cos rule rule-name policy policy-name]
user@host# set then revert
```

Configuring Application Profiles for CoS Rules

Configure CoS actions for FTP and SIP traffic. The application profile can then be used in CoS rule actions.

1. Configure a name for the application profile.

```
user@host# edit services cos application-profile profile-name
```

2. Specify the DSCP value to apply to the FTP or SIP (voice or video) packets.

For FTP traffic:

```
[edit services cos application-profile profile-name]
user@host# set ftp data dscp (alias | bits)
```

For SIP voice or video traffic:

```
[edit services cos application-profile profile-name]
user@host# set sip video | voice dscp dscp
```

The DSCP can be either a code point alias or a DSCP bit value.

3. Specify the forwarding class to apply to FTP or SIP packets.

For FTP traffic:

```
[edit services cos application-profile profile-name]
user@host# set ftp data forwarding-class class-name
```

For SIP voice or video traffic:

```
[edit services cos application-profile profile-name]
user@host# set sip video | voice forwarding-class forwarding-class dscp
```

The choices are:

- assured-forwarding
- best-effort
- expedited-forwarding
- network-control

Configuring CoS Rule Sets

A CoS rule set lets you specify a set of services CoS rules. You can then assign the rule set to a service set, which processes the rules in the order they appear. Once a rule matches the packet, the router performs the corresponding action, and no further rules in the rule set are applied.

1. Configure a name for the CoS rule set.

```
user@host# edit services cos rule-set rule-set-name
```

2. Specify the CoS rules that belong to the rule set.

```
[edit services cos rule-set rule-set-name]
user@host# set rule [rule-name]
```

Configuring the Service Set for CoS

You must apply CoS rules to a service set before the rules can be applied to traffic. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

To configure a service set with CoS rules:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service set, which requires a single service interface, or a next-hop service set, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

3. Specify the CoS rules to be used with the service set. You can either specify individual rules or rule sets.

To apply individual CoS rules:

```
[edit services service-set service-set-name]
user@host# set cos-rules [cos-rule-name]
```

To apply CoS rule sets:

```
[edit services service-set service-set-name]
user@host# set cos-rule-sets [cos-rule-set-name]
```

The service set processes the CoS rules or rule sets in the order in which they appear in the service set configuration.

4. (Optional) Assign at least one stateful firewall rule or NAT rule to the service set.

5. (Optional) Configure the service set to create a CoS session even if a packet is first received in the reverse direction of the matching direction of the CoS rule. The CoS rule values are then applied as soon as a packet in the correct match direction is received.

```
[edit services service-set service-set-name]  
user@host# set cos-options match-rules-on-reverse-flow
```

SEE ALSO

[Class of Service Overview for Services PICs \(Next Gen Services\)](#) | 255

3

PART

Stateful Firewall Services

Stateful Firewall Services Overview and Configuration | **264**

Stateful Firewall Services Overview and Configuration

IN THIS CHAPTER

- [Stateful Firewall Overview for Next Gen Services | 264](#)
- [Configuring Stateful Firewalls for Next Gen Services | 267](#)

Stateful Firewall Overview for Next Gen Services

Services PICs employ a type of firewall called a stateful firewall. Contrasted with a stateless firewall, which inspects packets in isolation, a stateful firewall provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions for new communication attempts.

Stateful firewalls group relevant flows into conversations, and decide whether the conversation is allowed to be established. If a conversation is allowed, all flows within the conversation are permitted, including flows that are created during the life cycle of the conversation.

Benefits

By inspecting the application protocol data of a flow, the stateful firewall intelligently enforces security policies and permits only the minimally required packet traffic.

Flows and Conversations

A typical Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) conversation consists of two flows: the initiation flow and the responder flow. However, some conversations, such as an FTP conversation, might consist of two control flows and many data flows.

A flow is identified by the following five properties:

- Source address
- Source port
- Destination address

- Destination port
- Protocol

Stateful Firewall Rules

Stateful firewall rules govern whether the conversation is allowed to be established. A rule consists of matching conditions and actions to take.

Matching conditions include direction, source address, destination address, and application protocol or service. In addition to the specific values you configure, you can assign the value **any**, **any-ipv4**, **any-ipv6**, or you can use an **address-book** under **services** to define address lists and ranges for use within stateful firewall rules. Finally, you can specify matches that result in the rule *not* being applied.

Actions in a stateful firewall rule include allowing the traffic or dropping the traffic.

Stateful firewall rules are directional. For each new conversation, the router software determines whether the initiation flow direction matches the rule direction.

Stateful firewall rules are ordered. The software checks the rules in the order in which you include them in the configuration. The first time the software finds a matching rule for a flow, the router implements the action specified by that rule, and ignores subsequent rules.

The stateful firewall rules are configured in relation to an interface. By default, the stateful firewall allows all sessions initiated from the hosts behind the interface to pass through the router.

Stateful Firewall Anomaly Checking

The stateful firewall recognizes the following events as anomalies and sends them to the IDS software for processing:

- IP anomalies:
 - IP version is not correct.
 - IP header length field is too small.
 - IP header length is set larger than the entire packet.
 - Bad header checksum.
 - IP total length field is shorter than header length.
 - Packet has incorrect IP options.
 - Internet Control Message Protocol (ICMP) packet length error.
 - Time-to-live (TTL) equals 0.
- IP address anomalies:

- IP packet source is broadcast or multicast.
- Land attack (source IP equals destination IP).
- IP fragmentation anomalies:
 - IP fragment overlap.
 - IP fragment missed.
 - IP fragment length error.
 - IP packet length is more than 64 kilobytes (KB).
 - Tiny fragment attack.
- TCP anomalies:
 - TCP port 0.
 - TCP sequence number 0 and flags 0.
 - TCP sequence number 0 and FIN/PSH/RST flags set.
 - TCP flags with wrong combination (TCP FIN/RST or SYN/(URG|FIN|RST)).
 - Bad TCP checksum.
- UDP anomalies:
 - UDP source or destination port 0.
 - UDP header length check failed.
 - Bad UDP checksum.
- Anomalies found through stateful TCP or UDP checks:
 - SYN followed by SYN-ACK packets without ACK from initiator.
 - SYN followed by RST packets.
 - SYN without SYN-ACK.
 - Non-SYN first flow packet.
 - ICMP unreachable errors for SYN packets.
 - ICMP unreachable errors for UDP packets.
- Packets dropped by stateful firewall rules.

Configuring Stateful Firewalls for Next Gen Services

IN THIS SECTION

- [Configuring Stateful Firewall Rules for Next Gen Services | 267](#)
- [Configuring Stateful Firewall Rule Sets for Next Gen Services | 269](#)
- [Configuring the Service Set for Stateful Firewalls for Next Gen Services | 269](#)

To configure stateful firewalls, you configure stateful firewall rules, and apply those rules to a service set. You can also configure stateful firewall rule sets, which contain a set of stateful firewall rules.

Configuring Stateful Firewall Rules for Next Gen Services

A stateful firewall rule specifies which traffic is processed and what action to apply to the traffic.

To configure a stateful firewall rule:

1. Configure a name for the stateful firewall rule.

```
user@host# edit services policies stateful-firewall-rule rule-name
```

2. Specify the traffic flow direction to which the stateful firewall rule applies.

```
[edit services policies stateful-firewall-rule rule-name]  
user@host# set match-direction (input | input-output | output)
```

If you configure **input-output**, the rule is applied to sessions initiated from either direction.

If this stateful firewall rule is applied to an interface-type service set, the direction is determined by whether a packet is entering or leaving the interface on which the service set is applied. If this stateful firewall rule is applied to a next-hop service set, the direction is input if the inside interface is used to route the packet, and the direction is output if the outside interface is used to route the package.

3. Configure a name for a policy.

```
[edit services policies stateful-firewall-rule rule-name]  
user@host# set policy policy-name
```

You can configure multiple policies for a stateful firewall rule. Each policy identifies the matching conditions for a flow, and whether or not to allow the flow. Once a policy in the rule matches a packet, that policy is applied and no other policies in the rule are processed.

4. Specify the destination address of the flows to which the policy applies.

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
user@host# set match destination-address (address | any | any-ipv4 | any-ipv6)
```

Alternatively, you can specify an **address-book** under the **services** configuration hierarchy to use in this step.

The destination address can be IPv4 or IPv6.

5. Specify the destination address of the flows to which the policy does not apply.

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
user@host# set match destination-address-excluded address
```

The destination address can be IPv4 or IPv6.

6. Specify the source address of the flows to which the policy applies.

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
user@host# set match source-address (address | any | any-ipv4 | any-ipv6)
```

Alternatively, you can specify an **address-book** under the **services** configuration hierarchy to use in this step.

The source address can be IPv4 or IPv6.

7. Specify the source address of the flows to which the policy does not apply.

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
user@host# set match source-address-excluded address
```

The source address can be IPv4 or IPv6.

8. Specify one or more application protocols to which the policy applies.

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
user@host# set match application [application-name]
```

Use an application protocol definition you have configured at the **[edit applications]** hierarchy level.

9. Specify an action that the policy takes.

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
user@host# set then (count | deny | reject | permit)
```

where:

count— Enables a count, in bytes or kilobytes, of all network traffic the policy allows to pass.

deny— Drop the packets.

permit— Accept the packets and send them to their destination.

reject— Drop the packets. For TCP traffic, send a TCP reset (RST) segment to the source host. For UDP traffic, send an ICMP **destination unreachable, port unreachable** message (type 3, code 3) to the source host.

Configuring Stateful Firewall Rule Sets for Next Gen Services

A stateful firewall rule set lets you specify a set of stateful firewall rules, which are processed in the order in which they appear in the rule set configuration. Once a stateful firewall rule in the rule set matches a packet, that rule is applied and no other rules in the rule set are processed`.

To configure a stateful firewall rule set:

1. Configure a name for the stateful firewall rule set.

```
user@host# edit services policies stateful-firewall-rule-set rule-set-name
```

2. Specify the stateful firewall rules that belong to the rule set.

```
[edit services policies stateful-firewall-rule-set rule-set-name]
user@host# set stateful-firewall-rule [rule-name]
```

Configuring the Service Set for Stateful Firewalls for Next Gen Services

Stateful firewall rules must be assigned to a service set before they can be applied to traffic.

To configure a service set to apply stateful firewall rules:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service set, which requires a single service interface, or a next-hop service set, which requires an inside and outside service interface.

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface interface-name
```

or

```
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface interface-name outside-service-interface
interface-name
```

3. Specify the stateful firewall rules to be used with the service set. You can specify either individual rules or rule sets but not both.

To apply individual stateful firewall rules:

```
[edit services service-set service-set-name]
user@host# set stateful-firewall-rules [rule-name]
```

To apply stateful firewall rule sets:

```
[edit services service-set service-set-name]
user@host# set stateful-firewall-rule-sets [rule-set-name]
```

The service set processes the stateful firewall rules or rule sets in the order in which they appear in the service set configuration.

4

PART

Intrusion Detection Services

[IDS Screens for Network Attack Protection Overview and Configuration](#) | 272

IDS Screens for Network Attack Protection Overview and Configuration

IN THIS CHAPTER

- [Understanding IDS Screens for Network Attack Protection | 272](#)
- [Configuring Network Attack Protection With IDS Screens for Next Gen Services | 276](#)

Understanding IDS Screens for Network Attack Protection

IN THIS SECTION

- [Intrusion Detection Services | 272](#)
- [Benefits | 273](#)
- [Session Limits | 273](#)
- [Suspicious Packet Patterns | 274](#)

Intrusion Detection Services

Intrusion detection services (IDS) screens give you a way to identify and drop traffic that is part of a network attack.

In an IDS screen, you can specify:

- The limits on the number of sessions that originate from individual sources or that terminate at individual destinations
- The types of suspicious packets

You can also choose to log an alarm when an IDS screen identifies a packet, rather than drop the packet.

In addition to IDS screens, you can use firewall filters and policers to stop illegal TCP flags and other bad flag combinations, and to specify general rate limiting (see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*). IDS screens add a more granular level of filtering.

Use firewall filters and stateful firewall filters to filter out traffic that does not need to be processed by an IDS screen.

Benefits

Provides protection against several types of network attacks.

Session Limits

You can use IDS screens to set session limits for traffic from an individual source or to an individual destination. This protects against network probing and flooding attacks. Traffic that exceeds the session limits is dropped. You can specify session limits either for traffic with a particular IP protocol, such as ICMP, or for traffic in general.

You decide whether the limits apply to individual addresses or to an aggregation of traffic from individual subnets of a particular prefix length. For example, if you aggregate limits for IPv4 subnets with a prefix length of 24, traffic from 192.0.2.2 and 192.0.2.3 is counted against the limits for the 192.0.2.0/24 subnet.

Some common network probing and flooding attacks that session limits protect against include:

ICMP Address Sweep—The attacker sends ICMP request probes (pings) to multiple targets. If a target machine replies, the attacker receives the IP address of the target.

ICMP Flood—The attacker floods a target machine by sending a large number of ICMP packets from one or more source IP addresses. The target machine uses up its resources as it attempts to process those ICMP packets, and then it can no longer process valid traffic.

TCP Port Scan—The attacker sends TCP SYN packets from one source to multiple destination ports of the target machine. If the target replies with a SYN-ACK from one or more destination ports, the attacker learns which ports are open on the target.

TCP SYN Flood—The attacker floods a target machine by sending a large number of TCP SYN packets from one or more source IP addresses. The attacker might use real source IP addresses, which results in a completed TCP connection, or might use fake source IP addresses, resulting in the TCP connection not being completed. The target creates states for all the completed and incomplete TCP connections. The target uses up its resources as it attempts to manage the connection states, and then it can no longer process valid traffic.

UDP Flood—The attacker floods a target machine by sending a large number of UDP packets from one or more source IP addresses. The target machine uses up its resources as it attempts to process those UDP packets, and then it can no longer process valid traffic.

Session limits for traffic from a source or to a destination include:

- maximum number of concurrent sessions
- maximum number of packets per second
- maximum number of connections per second

IDS screens also install a dynamic filter on the PFEs of line cards for suspicious activity when the following conditions occur:

- Either the packets per second or the number of connections per second for an individual source or destination address exceeds four times the session limit in the IDS screen. (Dynamic filters are not created from IDS screens that use subnet aggregation.)
- The services card CPU utilization percentage exceeds a configured value (default value is 90 percent).

The dynamic filter drops the suspicious traffic at the PFE, without the traffic being processed by the IDS screen. When the packet or connection rate no longer exceeds four times the limit in the IDS screen, the dynamic filter is removed.

Suspicious Packet Patterns

You can use IDS screens to identify and drop traffic with a suspicious packet pattern. This protects against attackers that craft unusual packets to launch denial-of-service attacks.

Suspicious packet patterns and attacks that you can specify in an IDS screen are:

ICMP fragmentation attack—The attacker sends the target ICMP packets that are IP fragments. These are considered suspicious packets because ICMP packets are usually short. When the target receives these packets, the results can range from processing packets incorrectly to crashing the entire system.

Malformed ICMPv6 packets—Malformed ICMPv6 packets can cause damage to the device and network. Examples of malformed IPv6 packets are packets that are too big (message type 2), that have the next header set to routing (43), or that have a routing header set to hop-by hop.

ICMP large packet attack—The attacker sends the target ICMP frames with an IP length greater than 1024 bytes. These are considered suspicious packets because most ICMP messages are small.

Ping of death attack—The attacker sends the target ICMP ping packets whose IP datagram length (ip_len) exceeds the maximum legal length (65,535 bytes) for IP packets, and the packet is fragmented. When the target attempts to reassemble the IP packets, a buffer overflow might occur, resulting in a system crashing, freezing, and restarting.

Bad option attack—The attacker sends the target packets with incorrectly formatted IPv4 options or IPv6 extension headers. This can cause unpredictable issues, depending on the IP stack implementation of routers and the target.

Fragmented IP packets—IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the target receives these packets, the results can range from processing the packets incorrectly to crashing the entire system.

IPv6 extension headers—Attackers can maliciously use extension headers for denial-of-service attacks or to bypass filters.

IPv4 options—Attackers can maliciously use IPv4 options for denial-of-service attacks.

IP teardrop attack—The attacker sends the target fragmented IP packets that overlap. The target machine uses up its resources as it attempts to reassemble the packets, and then it can no longer process valid traffic.

IP unknown protocol attack—The attacker sends the target packets with protocol numbers greater than 137 for IPv4 and 139 for IPv6. An unknown protocol might be malicious.

TCP FIN No ACK attack—The attacker sends the target TCP packets that have the FIN bit set but have the ACK bit unset. This can allow the attacker to identify the operating system of the target or to identify open ports on the target.

Land attack—The attacker sends the target spoofed SYN packets that contain the target's IP address as both the destination and the source IP address. The target uses up its resources as it repeatedly replies to itself. In another variation of the land attack, the SYN packets also contain the same source and destination ports.

TCP SYN ACK ACK attack—The attacker initiates Telnet or FTP connections with the target without completing the connections. The target's session table can fill up, resulting in the device rejecting legitimate connection requests.

TCP SYN FIN attack—The attacker sends the target TCP packets that have both the SYN and the FIN bits set. This can cause unpredictable behavior on the target, depending on its TCP stack implementation.

SYN fragment attack—The attacker sends the target SYN packet fragments. The target caches SYN fragments, waiting for the remaining fragments to arrive so it can reassemble them and complete the connection. A flood of SYN fragments eventually fills the host's memory buffer, preventing valid traffic connections.

TCP no flag attack—The attacker sends the target TCP packets containing no flags. This can cause unpredictable behavior on the target, depending on its TCP stack implementation.

TCP WinNuke attack—The attacker sends a TCP segment with the urgent (URG) flag set and destined for port 139 of a target running Windows. This might cause the target machine to crash.

RELATED DOCUMENTATION

Configuring Network Attack Protection With IDS Screens for Next Gen Services | 276

Configuring Network Attack Protection With IDS Screens for Next Gen Services

IN THIS SECTION

- [Configuring the IDS Screen Name, Direction, and Alarm Option | 276](#)
- [Configuring Session Limits in the IDS Screen | 277](#)
- [Configuring Suspicious Packet Pattern Detection in the IDS Screen | 281](#)
- [Configuring the Service Set for IDS | 284](#)

Configuring the IDS Screen Name, Direction, and Alarm Option

Configure the IDS screen name, traffic direction, and optional alarm.

1. Specify a name for the IDS screen.

```
[edit services screen]  
user@host# set ids-option screen-name
```

2. Specify whether the IDS screen is applied to input traffic, output traffic, or both.

```
[edit services screen ids-option screen-name]  
user@host# set match-direction (input | input-output | output)
```

3. If you want the IDS screen to log an alarm when packets exceed the session limit, rather than drop packets, configure **alarm-without-drop**.

```
[edit services screen ids-option screen-name]  
user@host# set alarm-without-drop
```

Configuring Session Limits in the IDS Screen

You can use IDS screens to set session limits for traffic from individual addresses or subnets and to individual addresses or subnets. This protects against network probing and flooding attacks. [Table 36 on page 277](#) shows the session limit options that protect against some common network probing and flooding attacks.

Table 36: IDS Screen Options for Network Attacks Type

Network Attack Type	[edit services screen ids-options screen-name limit-sessions] Options to Set
ICMP Address Sweep	by-source by-protocol icmp { maximum-sessions <i>number</i> ; packet-rate <i>number</i> ; session-rate <i>number</i> ; }
ICMP Flood	by-destination by-protocol icmp { maximum-sessions <i>number</i> ; packet-rate <i>number</i> ; session-rate <i>number</i> ; }
TCP Port Scan	(by-destination by-source) by-protocol tcp { maximum-sessions <i>number</i> ; packet-rate <i>number</i> ; }
TCP SYN Flood	(by-destination by-source) by-protocol tcp { maximum-sessions <i>number</i> ; packet-rate <i>number</i> ; session-rate <i>number</i> ; }
UDP Flood	by-destination by-protocol udp { maximum-sessions <i>number</i> ; packet-rate <i>number</i> ; session-rate <i>number</i> ; }

To configure the session limits in an IDS screen:

1. If you want to apply session limits to an aggregation of all sessions to individual destination subnets or from individual source subnets rather than individual addresses, configure aggregation.
 - a. To apply session limits to an aggregation of all sessions from within an individual IPv4 subnet, specify the subnet prefix length. The range is from 1 through 32.

```
[edit services screen ids-option screen-name aggregations]
user@host# set source-prefix-mask prefix-value
```

For example, the following statement configures an IPv4 prefix length of 24, and sessions from 192.0.2.2 and 192.0.2.3 are counted as sessions from the 192.0.2.0/24/24 subnet.

```
[edit services screen ids-option screen1 aggregations]
user@host# set source-prefix-mask 24
```

- b. To apply session limits to an aggregation of all sessions from within an individual IPv6 subnet, specify the subnet prefix length. The range is from 1 through 128.

```
[edit services screen ids-option screen-name aggregations]
user@host# set source-prefix-ipv6-mask prefix-value
```

For example, the following statement configures an IPv6 prefix length of 64, and sessions from 2001:db8:1234:72a2::2 and 2001:db8:1234:72a2::3 are counted as sessions from the 2001:db8:1234:72a2::/64 subnet.

```
[edit services screen ids-option screen1 aggregations]
user@host# set source-prefix-ipv6-mask 64
```

- c. To apply session limits to an aggregation of all sessions to an individual IPv4 subnet, specify the subnet prefix length. The range is from 1 through 32.

```
[edit services screen ids-option screen-name aggregations]
user@host# set destination-prefix-mask prefix-value
```

- d. To apply session limits to an aggregation of all sessions to an individual IPv6 subnet, specify the subnet prefix length. The range is from 1 through 128.

```
[edit services screen ids-option screen-name aggregations]
user@host# set destination-prefix-ipv6-mask prefix-value
```

2. If you want to apply session limits from a source for a particular IP protocol:
 - a. Configure the maximum number of concurrent sessions allowed from an individual source IP address or subnet for a particular IP protocol.


```
[edit services screen ids-option screen-name limit-session by-source ]
user@host# set by-protocol (icmp | tcp | udp) maximum-sessions number
```

- b. Configure the maximum number of packets per second allowed from an individual source IP address or subnet for a particular protocol.

```
[edit services screen ids-option screen-name limit-session by-source ]
user@host# set by-protocol (icmp | tcp | udp) packet-rate number
```

- c. Configure the maximum number of connections per second allowed from an individual source IP address or subnet for a particular protocol.

```
[edit services screen ids-option screen-name limit-session by-source ]
user@host# set by-protocol (icmp | tcp | udp) session-rate number
```

3. If you want to apply session limits to a destination for a particular IP protocol:

- a. Configure the maximum number of concurrent sessions allowed to an individual destination IP address or subnet for a particular IP protocol.

```
[edit services screen ids-option screen-name limit-session by-destination]
user@host# set by-protocol (icmp | tcp | udp) maximum-sessions number
```

- b. Configure the maximum number of packets per second allowed to an individual destination IP address or subnet for a particular protocol.

```
[edit services screen ids-option screen-name limit-session by-destination ]
user@host# set by-protocol (icmp | tcp | udp) packet-rate number
```

- c. Configure the maximum number of connections per second allowed to an individual destination IP address or subnet for a particular protocol.

```
[edit services screen ids-option screen-name limit-session by-destination ]
user@host# set by-protocol (icmp | tcp | udp) session-rate number
```

4. If you want to apply session limits from a source regardless of the IP protocol:

- a. Configure the maximum number of concurrent sessions allowed from an individual source IP address or subnet.

```
[edit services screen ids-option screen-name limit-session by-source ]
user@host# set maximum-sessions number
```

- b. Configure the maximum number of packets per second allowed from an individual source IP address or subnet

```
[edit services screen ids-option screen-name limit-session by-source ]
user@host# set packets-rate number
```

- c. Configure the maximum number of connections per second allowed from an individual source IP address or subnet.

```
[edit services screen ids-option screen-name limit-session by-source ]
user@host# set session-rate number
```

5. If you want to apply session limits to a destination regardless of the IP protocol:

- a. Configure the maximum number of concurrent sessions allowed to an individual destination IP address or subnet.

```
[edit services screen ids-option screen-name limit-session by-destination ]
user@host# set maximum-sessions number
```

- b. Configure the maximum number of packets per second allowed to an individual destination IP address or subnet

```
[edit services screen ids-option screen-name limit-session by-destination ]
user@host# set packets-rate number
```

- c. Configure the maximum number of connections per second allowed to an individual destination IP address or subnet.

```
[edit services screen ids-option screen-name limit-session by-destination]
user@host# set session-rate number
```

6. Specify the services card CPU utilization percentage that triggers the installation of a dynamic filter on the PFEs of the line cards for suspicious traffic. The default value is 90.

```
[edit services screen]
user@host# set cpu-throttle percentage percent
```

In addition to the CPU utilization percentage threshold, the packet rate or connection rate for an individual source or destination address must exceed four times the session limit in the IDS screen before the dynamic filter is installed. Dynamic filters are not created from IDS screens that use subnet aggregation.

The dynamic filter drops the suspicious traffic at the PFE, without the traffic being processed by the IDS screen. When the packet or connection rate no longer exceeds four times the limit in the IDS screen, the dynamic filter is removed.

Configuring Suspicious Packet Pattern Detection in the IDS Screen

You can use IDS screens to identify and drop suspicious packets. This protects against attackers that craft unusual packets to launch denial-of-service attacks.

To configure suspicious pattern detection:

1. To protect against ICMP fragmentation attacks, identify and drop ICMP packets that are IP fragments.

```
[edit services screen ids-option screen-name icmp]
user@host# set fragment
```

2. To identify and drop malformed ICMPv6 packets, configure **icmpv6-malformed**.

```
[edit services screen ids-option screen-name icmp]
user@host# set icmpv6-malformed
```

3. To protect against ICMP large packet attacks, identify and drop ICMP packets that are larger than 1024 bytes.

```
[edit services screen ids-option screen-name icmp]
user@host# set large
```

4. To protect against ping of death attacks, identify and drop oversized and irregular ICMP packets.

```
[edit services screen ids-option screen-name icmp]
user@host# set ping-death
```

5. To protect against bad option attacks, identify and drop packets with incorrectly formatted IPv4 options or IPv6 extension headers.

```
[edit services screen ids-option screen-name ip]
user@host# set bad-option
```

6. To identify and drop fragmented IP packets, configure **block-frag**.

```
[edit services screen ids-option screen-name ip]
user@host# set block-frag
```

7. To drop IPv6 packets with particular extension header values, specify the values.

```
[edit services screen ids-option screen-name ip]
user@host# set ipv6-extension-header header
```

The following header values can be configured:

ah-header—Authentication Header extension header

esp-header—Encapsulating Security Payload extension header

fragment-header—Fragment Header extension header

hop-by-hop-header—Hop-by-Hop option with the specified option:

CALIPSO-option—Common Architecture Label IPv6 Security Option

jumbo-payload-option—IPv6 jumbo payload option

quick-start-option—IPv6 quick start option

router-alert-option—IPv6 router alert option

RPL-option—Routing Protocol for Low-Power and Lossy Networks option

SFM-DPD-option—Simplified Multicast Forwarding IPv6 Duplicate Packet Detection option

user-defined-option-type *type-low* to *type-high*—A range of header types

Range: 1 through 255.

mobility-header—Mobility Header extension header.

routing-header—Routing Header extension header.

8. To drop IPv4 packets with particular IPv4 option values, specify the values.

```
[edit services screen ids-option screen-name ip]
user@host# set option
```

The following IPv4 option values can be configured:

loose-source-route-option— IP option of 3 (Loose Source Routing)

record-route-option— IP option of 7 (Record Route)

security-option— IP option of 2 (Security)

source-route-option—IP option of 3 (Loose Source Routing) or the IP option of 9 (Strict Source Routing)

stream-option—IP option of 8 (Stream ID)

strict-source-route-option—IP option of 9 (Strict Source Routing)

timestamp-option—IP option of 4 (Internet timestamp)

9. To protect against IP teardrop attacks, identify and drop fragmented IP packets that overlap.

```
[edit services screen ids-option screen-name ip]
user@host# set tear-drop
```

10. To protect against IP unknown protocol attacks, identify and drop IP frames with protocol numbers greater than 137 for IPv4 and 139 for IPv6.

```
[edit services screen ids-option screen-name ip]
user@host# set unknown-protocol
```

11. To protect against TCP FIN No ACK Attacks, identify and drop any packet with the FIN flag set and without the ACK flag set.

```
[edit services screen ids-option screen-name tcp]
user@host# set fin-no-ack
```

12. To protect against land attacks, identify and drop SYN packets that have the same source and destination address or port.

```
[edit services screen ids-option screen-name tcp]
user@host# set land
```

13. To protect against TCP SYN ACK ACK attacks, configure the maximum number of connections from an IP address that can be opened without being completed.

```
[edit services screen ids-option screen-name tcp]
user@host# set syn-ack-ack-proxy number
```

14. To protect against TCP SYN FIN attacks, identify and drop packets that have both the SYN and FIN flags set.

```
[edit services screen ids-option screen-name tcp]
user@host# set syn-fin
```

15. To protect against SYN fragment attacks, identify and drop SYN packet fragments.

```
[edit services screen ids-option screen-name tcp]
user@host# set syn-frag
```

16. To protect against TCP no flag attacks, identify and drop TCP packets that have no flag fields set.

```
[edit services screen ids-option screen-name tcp]
user@host# set tcp-no-flag
```

17. To protect against TCP WinNuke attacks, identify and drop TCP segments that are destined for port 139 and have the urgent (URG) flag set.

```
[edit services screen ids-option screen-name tcp]
user@host# set winnuke
```

Configuring the Service Set for IDS

Configure a service set to apply the IDS screen.

1. Assign the IDS screen to a service set.

```
[edit services]
user@host# set service-set service-set-name ids-option screen-name
```

If the service set is associated with an AMS interface, then the session limits you configure are applicable to each member interface.

2. Limit the packets that the IDS screen processes by configuring a stateful firewall rule . The stateful firewall rule can identify either the traffic that should undergo IDS processing or the traffic that should skip IDS processing:
 - To allow IDS processing on the traffic that matches the stateful firewall rule, include **accept** at the **[edit services stateful-firewall rule *rule-name* term *term-name* then]** hierarchy level.
 - To skip IDS processing on the traffic that matches the stateful firewall rule, include **accept skip-ids** at the **[edit services stateful-firewall rule *rule-name* term *term-name* then]** hierarchy level.
3. Assign the stateful firewall rule to the service set.

```
[edit services]
user@host# set service-set service-set-name stateful-firewall-rules rule-name
```

4. To protect against header anomaly attacks, configure a header integrity check for the service set.

```
[edit services]
user@host# set service-set service-set-name service-set-options header-integrity-check enable-all
```

RELATED DOCUMENTATION

Understanding IDS Screens for Network Attack Protection | 272

5

PART

Traffic Load Balancing

Traffic Load Balancing Overview and Configuration | **287**

Traffic Load Balancing Overview and Configuration

IN THIS CHAPTER

- Traffic Load Balancer Overview | 287
- Configuring TLB | 296

Traffic Load Balancer Overview

IN THIS SECTION

- Traffic Load Balancing Support Summary | 287
- Traffic Load Balancer Application Description | 288
- Traffic Load Balancer Modes of Operation | 289
- Traffic Load Balancer Functions | 292
- Traffic Load Balancer Application Components | 293
- Traffic Load Balancer Configuration Limits | 295

Traffic Load Balancing Support Summary

Table 37 on page 287 provides a summary of the traffic load balancing support on the MS-MPC and MS-MIC cards for Adaptive Services versus support on the MX-SPC3 security services card for Next Gen Services.

Table 37: Traffic Load Balancing Support Summary

	MS-MPC		MX-SPC3
Junos Release	< 16.1R6 & 18.2.R1	≥ 16.1R6 & 18.2R1	19.3R2
Max # of Instances per Chassis	32	2,000 / 32 in L2 DSR mode	2,000

Table 37: Traffic Load Balancing Support Summary (continued)

	MS-MPC		MX-SPC3
Max # of Virtual Services per Instance	32	32	32
Max # of virtual IP address per virtual service		1	1
Max # of Groups per Instances	32	32	32
Max # of Real-Services (Servers) per Group	255	255	255
Max # of groups per virtual service		1	1
Max # of Network Monitor Profiles per Group		2	2
Max # of HC's per security services per PIC/NPU in 5-sec's		4,000	1,250 – 19.3R2 10,000 – 20.1R1
Supported Health Check Protocols	ICMP, TCP, UDP, HTTP, SSL, Custom		ICMP, TCP, UDP, HTTP, SSL, Custom

Traffic Load Balancer Application Description

Traffic Load Balancer (TLB) is supported on MX Series routers with either of the Multiservices Modular Port Concentrator (MS-MPC), Multiservices Modular Interface Card (MS-MIC), or the MX Security Services Processing Card (MX-SPC3) and in conjunction with the Modular Port Concentrator (MPC) line cards supported on the MX Series routers as described in [Table 38 on page 288](#).

NOTE: You cannot run Deterministic NAT and TLB simultaneously.

Table 38: TLB MX Series Router Platform Support Summary

TLB Mode	MX Platform Coverage
Multiservices Modular Port Concentrator (MS-MPC)	MX240, MX2480, MX960, MX2008, MX2010, MX2020

Table 38: TLB MX Series Router Platform Support Summary (*continued*)

TLB Mode	MX Platform Coverage
Multiservices Modular Interface Card (MS-MIC)	MX5, MX10, MX40, MX80, MX104, MX240, MX2480, MX960, MX2008, MX2010, MX2020
MX Security Services Processing Card (MX-SPC3)	MX240, MX480, MX960

- TLB enables you to distribute traffic among multiple servers.
- TLB employs an MS-MPC-based control plane and a data plane using the MX Series router forwarding engine.
- TLB uses an enhanced version of equal-cost multipath (ECMP). Enhanced ECMP facilitates the distribution of flows across groups of servers. Enhancements to native ECMP ensure that when servers fail, only flows associated with those servers are impacted, minimizing the overall network churn on services and sessions.
- TLB provides application-based health monitoring for up to 255 servers per group, providing Intelligent traffic steering based on health checking of server availability information. You can configure an aggregated multiservices (AMS) interface to provide one-to-one redundancy for MS-MPCs or Next Gen Services MX-SPC3 card used for server health monitoring.
- TLB applies its flow distribution processing to ingress traffic.
- TLB supports multiple virtual routing instances to provide improved support for large scale load balancing requirements.
- TLB supports static virtual-IP-address-to-real-IP-address translation, and static destination port translation during load balancing.

Traffic Load Balancer Modes of Operation

IN THIS SECTION

- [Transparent Mode Layer 2 Direct Server Return | 290](#)
- [Translated Mode | 291](#)
- [Transparent Mode Layer 3 Direct Server Return | 292](#)

Traffic Load Balancer provides three modes of operation for the distribution of outgoing traffic and for handling the processing of return traffic.

Table 39 on page 290 summarizes the TLB support and which cards it's supported on.

Table 39: TLB Versus Security Service Cards Summary

Security Service Card	MS-MPC/MS-MIC	MX-SPC3
Translate	Yes	Yes
Transparent Layer 3 Direct Server Return	Yes	Yes
Transparent Layer 2 Direct Server Return	Yes	Not Supported


Transparent Mode Layer 2 Direct Server Return

When you use transparent mode Layer 2 direct server return (DSR):

- The PFE processes data.
- Load balancing works by changing the Layer 2 MAC of packets.
- An MS-MPC performs the network-monitoring probes.
- Real servers must be directly (Layer 2) reachable from the MX Series router.
- TLB installs a route and all the traffic over that route is load-balanced.
- TLB never modifies Layer 3 and higher level headers.

Figure 7 on page 290 shows the TLB topology for transparent mode Layer 2 DSR.

Figure 7: TLB Topology for Transparent Mode



Translated Mode

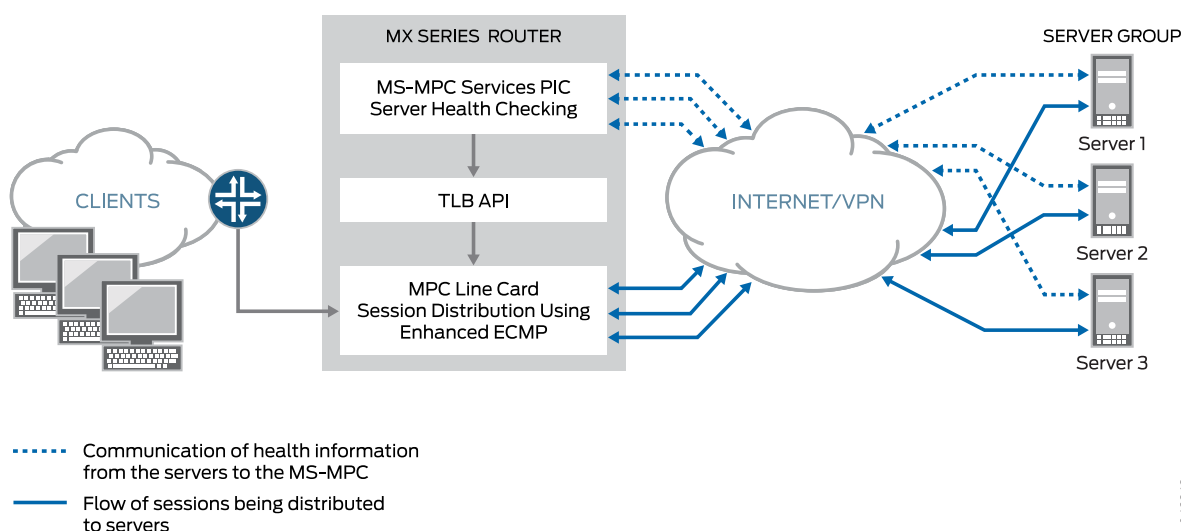
Translated mode provides greater flexibility than transparent mode Layer 2 DSR. When you choose translated mode:

- An MS-MPC performs the network-monitoring probes.
- The PFE performs stateless load balancing:
 - Data traffic directed to a virtual IP address undergoes translation of the virtual IP address to a real server IP address and translates the virtual port to a server listening port. Return traffic undergoes the reverse translation.
 - Client to virtual IP traffic is translated; the traffic is routed to reach its destination.
 - Server-to-client traffic is captured using implicit filters and directed to an appropriate load-balancing next hop for reverse processing. After translation, traffic is routed back to the client.
 - Two load balancing methods are available: random and hash. The random method is only for UDP traffic and provides quavms-random distribution. While not literally random, this mode provides fair distribution of traffic to an available set of servers. The hash method provides a hash key based on any combination of the source IP address, destination IP address, and protocol.

NOTE: Translated mode processing is only available for IPv4-to-IPv4 and IPv6-to-IPv6 traffic.

Figure 8 on page 291 shows the TLB topology for translated mode.

Figure 8: TLB Topology for Translated Mode



Transparent Mode Layer 3 Direct Server Return

Transparent mode Layer 3 DSR load balancing distributes sessions to servers that can be a Layer 3 hop away. Traffic is returned directly to the client from the real-server.

Traffic Load Balancer Functions

TLB provides the following functions:

- TLB always distributes the *requests* for any flow. When you specify DSR mode, the response returns directly to the source. When you specify translated mode, reverse traffic is steered through implicit filters on server-facing interfaces.
- TLB supports hash-based load balancing or random load balancing.
- TLB enables you to configure servers offline to prevent a performance impact that might be caused by a rehashing for all existing flows. You can add a server in the administrative down state and use it later for traffic distribution by disabling the administrative down state. Configuring servers offline helps prevent traffic impact to other servers.
- When health checking determines a server to be down, only the affected flows are rehashed.
- When a previously down server is returned to service, all flows belonging to that server based on hashing return to it, impacting performance for the returned flows. For this reason, you can disable the automatic rejoining of a server to an active group. You can return servers to service by issuing the **request services traffic-load-balance real-service rejoin** operational command.

NOTE: NAT is not applied to the distributed flows.

- Health check monitoring application runs on an MS-MPC/NPU. This network processor unit (NPU) is not used for handling data traffic.
- TLB supports static virtual-IP-address-to-real-IP-address translation, and static destination port translation during load balancing.
- TLB provides multiple VRF support.

RELATED DOCUMENTATION

[Interchassis High-Availability](#)
[Understanding AMS Interfaces](#)

Traffic Load Balancer Application Components

Servers and Server Groups

TLB enables configuration of groups of up to 255 servers (referred to in configuration statements as *real services*) for use as alternate destinations for stateless session distribution. All servers used in server groups must be individually configured before assignment to groups. Load balancing uses hashing or randomization for session distribution. Users can add and delete servers to and from the TLB server distribution table and can also change the administrative status of a server.

NOTE: TLB uses the session distribution next-hop API to update the server distribution table and retrieve statistics. *Applications do not have direct control on the server distribution table management. They can only influence changes indirectly through the add and delete services of the TLB API.*

Server Health Monitoring – Single Health Check and Dual Health Check

TLB supports TCP, HTTP, SSL Hello, and custom health check probes to monitor the health of servers in a group. You can use a single probe type for a server group, or a dual health check configuration that includes two probe types. The configurable health monitoring function resides on either an MX-SPC3 or an MS-MPC. By default, probe requests are sent every 5 seconds. Also by default, a real server is declared down only after five consecutive probe failures and declared up only after five consecutive probe successes.

Use a custom health check probe to specify the following:

- Expected string in the probe response
- String that is sent with the probe
- Server status to assign when the probe times out (up or down)
- Server status to assign when the expected response to the probe is received (up or down)
- Protocol – UDP or TCP

TLB provides *application stickiness*, meaning that server failures or changes do not affect traffic flows to other active servers. Changing a server's administrative state from up to down does not impact any active flows to remaining servers in the server distribution table. Adding a server or deleting a server from a group has some traffic impact for a length of time that depends on your configuration of the interval and retry parameters in the monitoring profile.

TLB provides two levels of server health monitoring:

- **Single Health Check**—One probe type is attached to a server group by means of the **network-monitoring-profile** configuration statement.

- **TLB Dual Health Check (TLB-DHC)**—Two probe types are associated with a server group by means of the **network-monitoring-profile** configuration statement. A server's status is declared based on the result of two health check probes. Users can configure up to two health check profiles per server group. If a server group is configured for dual health check, a real-service is declared to be UP only when both health-check probes are simultaneously UP; otherwise, a real-service is declared to be DOWN.

NOTE:

The following restrictions apply to AMS interfaces used for server health monitoring:

- An AMS interface configured under a TLB instance uses its configured member interfaces exclusively for health checking of configured multiple real servers.
- The member interfaces use unit 0 for single VRF cases, but can use units other than 1 for multiple VRF cases.
- TLB uses the IP address that is configured for AMS member interfaces as the source IP address for health checks.
- The member interfaces must be in the same routing instance as the interface used to reach real servers. This is mandatory for TLB server health-check procedures.

Virtual Services

The virtual service provides a virtual IP address (VIP) that is associated with the group of servers to which traffic is directed as determined by hash-based or random session distribution and server health monitoring. In the case of Layer2 DSR and Layer3 DSR, the special address 0.0.0.0 causes all traffic flowing to the forwarding instance to be load balanced.

The virtual service configuration includes:

- **Mode**—indicating how traffic is handled (translated or transparent).
- The group of servers to which sessions are distributed.
- The load balancing method.
- Routing instance and route metric.

BEST PRACTICE: Although you can assign a virtual address of 0.0.0.0 in order to use default routing, we recommend using a virtual address that can be assigned to a routing instance set up specifically for TLB.

Traffic Load Balancer Configuration Limits

Traffic Load Balancer configuration limits are described in [Table 40 on page 295](#).

Table 40: TLB Configuration Limits

Configuration Component	Configuration Limit
Maximum number of instances.	<p>Starting in Junos OS Release 16.1R6 and Junos OS Release 18.2R1, the TLB application supports 2000 TLB instances for virtual services that use the direct-server-return or the translated mode. In earlier releases, the maximum number of instances is 32.</p> <p>If multiple virtual services are using the same server group, then all of those virtual services must use the same load balancing method to support 2000 TLB instances.</p> <p>For virtual services that use the layer2-direct-server-return mode, TLB supports only 32 TLB instances. To perform the same function as the layer2-direct-server-return mode and have support for 2000 TLB instances, you can use the direct-server-return mode and use a service filter with the skip action.</p>
Maximum number of servers per group	255
Maximum number of virtual services per services PIC	32
Maximum number of health checks per services PIC in a 5-second interval	<p>For MS-MPC services cards: 2000</p> <p>For Next Gen Services mode and the MX-SPC3 services cards: 1250</p>
Maximum number of groups per virtual service	1
Maximum number of virtual IP addresses per virtual service	1
Supported health checking protocols	<p>ICMP, TCP, HTTP, SSL, Custom</p> <p>NOTE: ICMP health checking is supported only on MS-MPC services cards.</p>

Release History Table

Release	Description
16.1R6	Starting in Junos OS Release 16.1R6 and Junos OS Release 18.2R1, the TLB application supports 2000 TLB instances for virtual services that use the direct-server-return or the translated mode.

RELATED DOCUMENTATION

| [Configuring TLB](#) | [296](#)

Configuring TLB

IN THIS SECTION

- [Loading the TLB Service Package](#) | [296](#)
- [Configuring a TLB Instance Name](#) | [297](#)
- [Configuring Interface and Routing Information](#) | [297](#)
- [Configuring Servers](#) | [300](#)
- [Configuring Network Monitoring Profiles](#) | [301](#)
- [Configuring Server Groups](#) | [302](#)
- [Configuring Virtual Services](#) | [304](#)
- [Configuring Tracing for the Health Check Monitoring Function](#) | [307](#)

The following topics describe how to configure TLB. To create a complete application, you must also define interfaces and routing information. You can optionally define firewall filters and policy options in order to differentiate TLB traffic.

Loading the TLB Service Package

Load the TLB service package on each service PIC on which you want to run TLB.

NOTE: For Next Gen Services and the MX-SPC3 services card, you do not need to load this package.

To load the TLB service package on a service PIC:

- Load the **jservices-traffic-dird** package.

```
[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]
user@host# set package jservices-traffic-dird
```

For example:

```
[edit chassis fpc 3 pic 0 adaptive-services service-package extension-provider]
user@host# set package jservices-traffic-dird
```

Configuring a TLB Instance Name

To configure a name for the TLB instance:

- At the **[edit services traffic-load-balance]** hierarchy level, identify the TLB instance name.

```
[edit services traffic-load-balance]
user@host# set instance instance-name
```

For example:

```
[edit services traffic-load-balance]
user@host# set instance tlb-instance1
```

Configuring Interface and Routing Information

To configure interface and routing information:

1. At the **[edit services traffic-load-balance instance *instance-name*]** hierarchy level, identify the service interface associated with this instance.

```
user@host# [edit services traffic-load-balance instance instance-name]
user@host# set interface interface-name
```

For example, on an MS-MPC:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set interface ms-1/0/0
```

For example, for Next Gen Services on an MX-SPC3:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set interface vms-1/0/0
```

2. Enable the routing of health-check packet responses from real servers to the service interface that you identified in Step 1.

```
[edit interfaces]
user@host# set interface-name unit 0 ip-address-owner service-plane
```

For example, on an MS-MPC:

```
[edit interfaces]
user@host# set ms-1/0/0 unit 0 ip-address-owner service-plane
```

For example, on an MX-SPC3:

```
[edit interfaces]
user@host# set vms-1/0/0 unit 0 ip-address-owner service-plane
```

3. Specify the client interface for which an implicit filter is defined to direct traffic in the forward direction. This is required only for translated mode.

```
user@host# [edit services traffic-load-balance instance instance-name]
user@host# set client-interface interface-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set client-interface ge-5/2/0.0
```

4. Specify the virtual routing instance used to route data traffic in the forward direction to servers. This is required for SLT and Layer 3 DSR; it is optional for Layer 2 DSR.

```
user@host# [edit services traffic-load-balance instance instance-name]
```

```
user@host# set server-vrf server-vrf
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set server-vrf server-vrf
```

5. Specify the server interface for which implicit filters are defined to direct return traffic to the client.

NOTE: Implicit filters for return traffic are not used for DSR.

```
user@host# [edit services traffic-load-balance instance instance-name]
user@host# set server-interface server-interface
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set server-interface ge-5/2/1.0
```

6. (Optional) Specify the filter used to bypass health checking for return traffic.

```
user@host# [edit services traffic-load-balance instance instance-name]
user@host# set server-inet-bypass-filter server-inet-bypass-filter
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set server-inet-bypass-filter tlb-ipv4-bypass
```

7. Specify the virtual routing instance in which you want the data in the reverse direction to be routed to the clients.

```
user@host# [edit services traffic-load-balance instance instance-name]
user@host# set client-vrf client-vrf
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set client-vrf client-vrf
```

NOTE: Virtual routing instances for routing data in the reverse direction are not used with DSR.

Configuring Servers

To configure servers for the TLB instance:

- Configure a logical name and IP address for each server to be made available for next-hop distribution.

```
[edit services traffic-load-balance instance instance-name]
user@host# set real-service real-service-name address server-ip-address
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1]
user@host# set real-service rs138 address 172.26.99.138
user@host# set real-service rs139 address 172.26.99.139
user@host# set real-service rs140 address 172.26.99.140
```

Configuring Network Monitoring Profiles

A network monitoring profile configures a health check probe, which you assign to a server group to which session traffic is distributed.

To configure a network monitoring profile:

1. Configure the type of probe to use for health monitoring — **icmp**, **tcp**, **http**, **ssl-hello**, or **custom**.

NOTE: **icmp** probes are supported only on MS-MPC cards.

Next Gen Services and the MX-SPC3 do not support ICMP probes in this release.

- For an ICMP probe:

```
[edit services network-monitoring profile profile-name]
user@host.com# set icmp
```

- For a TCP probe:

```
[edit services network-monitoring profile profile-name]
user@host.com# set tcp port tcp-port-number
```

- For an HTTP probe:

```
[edit services network-monitoring profile profile-name]
user@host.com# set http host hostname url url port http-port-number method (get | option)
```

- For an SSL probe:

```
[edit services network-monitoring profile profile-name]
user@host.com# set ssl-hello port port ssl-version
```

- For a custom probe:

```
[edit services network-monitoring profile profile-name]
user@host.com# set custom cmd priority default-real-service-status (down | up) expect (ascii | binary)
receive-string port port real-service-action (down | up) send (ascii | binary) send-string
```

2. Configure the interval for probe attempts, in seconds (1 through 180).

```
[edit services network-monitoring profile profile-name]
```

```
user@host.com# set probe-interval interval
```

For example:

```
[edit services network-monitoring profile profile1-icmp]
user@host.com# set probe-interval 2
```

3. Configure the number of failure retries, after which the real server is tagged as down.

```
[edit services network-monitoring profile profile-name]
user@host.com# set failure-retries number-of-retries
```

For example:

```
[edit services network-monitoring profile profile1-icmp]
user@host.com# set failure-retries 3
```

4. Configure the number of recovery retries, which is the number of successful probe attempts after which the server is declared up.

```
[edit services network-monitoring profile profile-name]
user@host.com# set recovery-retries number-of-retries
```

For example:

```
[edit services network-monitoring profile profile1-icmp]
user@host.com# set recovery-retries 1
```

Configuring Server Groups

Server groups consist of servers to which traffic is distributed by means of stateless, hash-based session distribution and server health monitoring.

To configure a server group:

1. Specify the names of one or more configured real servers.

```
[edit services traffic-load-balance instance instance-name groups group-name]
user@host.com# set real-services real-service-name, ...
```


For example:

```
[edit services traffic-load-balance instance tlb-instance1 groups tlb-group1]
user@host.com# set real-services [ rs138 rs139 rs140 ]
```

2. Configure the routing instance for the group when you do not want to use the default instance, **inet.0**.

```
[edit services traffic-load-balance instance instance-name groups group-name]
user@host.com# set routing-instance routing-instance-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 groups tlb-group1]
user@host.com# set routing-instance tlb-routing-instance1
```

3. (Optional) Disable the default option that allows a server to rejoin the group automatically when it comes up.

```
[edit services traffic-load-balance instance instance-name group group-name]
user@host.com# set real-service-rejoin-options no-auto-rejoin
```

4. (Optional) Configure the logical unit of the instance's service interface to use for health checking.
 - a. Specify the logical unit.

```
[edit services traffic-load-balance instance instance-name group group-name]
user@host.com# set health-check-interface-subunit health-check-interface-subunit
```

- b. Enable the routing of health-check packet responses from real servers to the interface.

```
[edit interfaces]
user@host.com# set interface-name unit subunit ip-address-owner service-plane
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 group tlb-group1]
user@host.com# set health-check-interface-subunit 30
[edit interfaces]
user@host.com# set ms-1/0/0 unit 30 ip-address-owner service-plane
```

5. Configure one or two network monitoring profiles to be used to monitor the health of servers in this group.

```
[edit services traffic-load-balance instance instance-name groups group-name]
user@host.com# set network-monitoring-profile profile-name1 profile-name2
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 groups tlb-group1]
user@host.com# set network-monitoring-profile profile1-icmp profile2-http
```

Configuring Virtual Services

A virtual service provides an address that is associated with a the group of servers to which traffic is directed as determined by hash-based or random session distribution and server health monitoring. You may optionally specify filters and routing instances to steer traffic for TLB.

To configure a virtual service:

1. At the **[edit services traffic-load-balance instance *instance-name*]** hierarchy level, specify a non-zero address for the virtual service.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# set address virtual-ip-address
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set address 192.0.2.11
```

2. Specify the server group used for this virtual service.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# set group group-name
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set group tlb-group1
```

3. (Optional) Specify a routing instance for the virtual service. If you do not specify a routing instance, the default routing instance is used.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# set routing-instance routing-instance
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set routing-instance msp-tproxy-server-vrf31
```

4. Specify the processing mode for the virtual service.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# set mode (layer2-direct-server-return | direct-server-return | translated)
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set mode translated
```

5. (Optional) For a translated mode virtual service, enable the addition of the IP addresses for all the real servers in the group under the virtual service to the server-side filters. Doing this allows you to configure two virtual services with the same listening port and protocol on the same interface and VRF.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# set include-real-server-ips-in-server-filter
```

6. (Optional) Specify a routing metric for the virtual service.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# set routing-metric routing-metric
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set routing-metric 128
```

- Specify the method used for load balancing. You can specify a hash method that provides a hash key based on any combination of the source IP address, destination IP address, and protocol, or you can specify **random**.

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set load-balancing-method (hash hash-key (source-ip | destination-ip | proto) | random)
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set load-balancing-method hash hash-key source-ip
```

or

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set load-balancing-method random
```

NOTE: If you switch between the hash method and the random method for a virtual service, the statistics for the virtual service are lost.

- For a translated mode virtual service, specify a service for translation, including a virtual-port, server-listening-port, and protocol.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# set service service-name virtual-port virtual-port server-listening-port server-listening-port protocol
(udp | tcp)
```

For example:

```
[edit services traffic-load-balance instance tlb-instance1 virtual-service virtual-service1]
user@host# set service fast-track-service virtual-port 1111 server-listening-port 22 protocol tcp
```

- Commit the configuration.

```
[edit services traffic-load-balance instance instance-name virtual-service virtual-service-name]
user@host# commit
```

NOTE: In the absence of a client-interface configuration under the TLB instance, the implicit client filter (for VIP) is attached to the client-vrf configured under the TLB instance. In this case, the routing-instance under a translate mode virtual service cannot be the same as the client-vrf configured under the TLB instance. if it is, the commit fails.

Configuring Tracing for the Health Check Monitoring Function

To configure tracing options for the health check monitoring function:

- 1. Specify that you want to configure tracing options for the health check monitoring function.

```
[edit services network-monitoring]
user@host# edit traceoptions
```

- 2. (Optional) Configure the name of the file used for the trace output.

```
[edit services network-monitoring traceoptions]
user@host# set file file-name
```

- 3. (Optional) Disable remote tracing capabilities.

```
[edit services network-monitoring traceoptions]
user@host# set no-remote-trace
```

- 4. (Optional) Configure flags to filter the operations to be logged.

```
[edit services network-monitoring traceoptions]
user@host# set flag flag
```

[Table 41 on page 307](#) describes the flags that you can include.

Table 41: Trace Flags

Flag	Support on MS-MPC and MX-SPC3 Cards	Description
all	MS-MPC and MX-SPC3	Trace all operations.
all-real-services	MX-SPC3	Trace all real services.

Table 41: Trace Flags (*continued*)

Flag	Support on MS-MPC and MX-SPC3 Cards	Description
config	MS-MPC and MX-SPC3	Trace traffic load balancer configuration events.
connect	MS-MPC and MX-SPC3	Trace traffic load balancer ipc events.
database	MS-MPC and MX-SPC3	Trace database events.
file-descriptor-queue	MS-MPC	Trace file descriptor queue events.
inter-thread	MS-MPC	Trace inter-thread communication events.
filter	MS-MPC and MX-SPC3	Trace traffic load balancer filter programming events.
health	MS-MPC and MX-SPC3	Trace traffic load balancer health events.
messages	MS-MPC and MX-SPC3	Trace normal events.
normal	MS-MPC and MX-SPC3	Trace normal events.
operational-commands	MS-MPC and MX-SPC3	Trace traffic load balancer show events.
parse	MS-MPC and MX-SPC3	Trace traffic load balancer parse events.
probe	MS-MPC and MX-SPC3	Trace probe events.
probe-infra	MS-MPC and MX-SPC3	Trace probe infra events.
route	MS-MPC and MX-SPC3	Trace traffic load balancer route events.
snmp	MS-MPC and MX-SPC3	Trace traffic load balancer SNMP events.
statistics	MS-MPC and MX-SPC3	Trace traffic load balancer statistics events.
system	MS-MPC and MX-SPC3	Trace traffic load balancer system events.

5. (Optional) Configure the level of tracing.

```
[edit services network-monitoring traceoptions]
user@host# set level (all | error | info | notice | verbose | warning)
```

6. (Optional) Configure tracing for a particular real server within a particular server group.

```
[edit services network-monitoring traceoptions]
user@host# set monitor monitor-object-name group-name group-name real-services-name real-service-name
```

7. (Optional) Starting in Junos OS Release 16.1R6 and 18.2R1, configure tracing for a particular virtual service and instance.

```
[edit services traffic-load-balance traceoptions]
user@host# set monitor monitor-object-name instance-name instance-name virtual-svc-name
virtual-service-name
```

Release History Table

Release	Description
16.1R6	Starting in Junos OS Release 16.1R6 and 18.2R1, configure tracing for a particular virtual service and instance.

RELATED DOCUMENTATION

| [Traffic Load Balancer Overview](#) | [287](#)



DNS Request Filtering

DNS Request Filtering Overview and Configuration | **311**

DNS Request Filtering Overview and Configuration

IN THIS CHAPTER

- [DNS Request Filtering for Disallowed Website Domains | 311](#)
- [DNS Request Filtering System Logging Error Messages | 320](#)

DNS Request Filtering for Disallowed Website Domains

IN THIS SECTION

- [Overview of DNS Request Filtering | 311](#)
- [How to Configure DNS Request Filtering | 313](#)

Overview of DNS Request Filtering

IN THIS SECTION

- [Benefits | 312](#)
- [Disallowed Domain Filter Database File | 313](#)
- [DNS Filter Profile | 313](#)


Starting in Junos OS Release 18.3R1, you can configure DNS filtering to identify DNS requests for disallowed website domains. Starting in Junos OS Release 19.3R2, you can configure DNS filtering if you are running Next Gen Services with the MX-SPC3 services card. Next Gen Services are supported on MX240, MX480 and MX960 routers. For DNS request types A, AAAA, MX, CNAME, TXT, SRV, and ANY, you configure the action to take for a DNS request for a disallowed domain. You can either:

- Block access to the website by sending a DNS response that contains the IP address or fully qualified domain name (FQDN) of a DNS sinkhole server. This ensures that when the client attempts to send traffic to the disallowed domain, the traffic instead goes to the sinkhole server (see [Figure 9 on page 312](#)).
- Log the request and allow access.

For other DNS request types for a disallowed domain, the request is logged and access is allowed.

The actions that the sinkhole server takes are not controlled by the DNS request filtering feature; you are responsible for configuring the sinkhole server actions. For example, the sinkhole server could send a message to the requestor that the domain is not reachable and prevent access to the disallowed domain.

Figure 9: DNS Request for Disallowed Domain



Benefits

DNS filtering redirects DNS requests for disallowed website domains to sinkhole servers, while preventing anyone operating the system from seeing the list of disallowed domains. This is because the disallowed domain names are in an encrypted format.

Disallowed Domain Filter Database File

DNS request filtering requires a disallowed domain filter database .txt file, which identifies each disallowed domain name, the action to take on a DNS request for the disallowed domain, and the IP address or fully qualified domain name (FQDN) of a DNS sinkhole server.

DNS Filter Profile

You configure a DNS filter profile to specify which disallowed domain filter database file to use. You can also specify the interfaces on which DNS request filtering is performed, limit the filtering to requests for specific DNS servers, and limit the filtering to requests from specific source IP address prefixes.

How to Configure DNS Request Filtering

IN THIS SECTION

- [How to Configure a Domain Filter Database | 313](#)
- [How to Configure a DNS Filter Profile | 314](#)
- [How to Configure a Service Set for DNS Filtering | 319](#)

To filter DNS requests for disallowed website domains, perform the following:

How to Configure a Domain Filter Database

Create one or more domain filter database files that include an entry for each disallowed domain. Each entry specifies what to do with a DNS request for a disallowed website domain.

To configure a domain filter database file:

1. Create the name for the file. The database file name can have a maximum length of 64 characters and must have a **.txt** extension.
2. Add a file header with a format such as
20170314_01:domain,sinkhole_ip,v6_sinkhole,sinkhole_fqdn,id,action.
3. Add an entry in the file for each disallowed domain. You can include a maximum of 10,000 domain entries. Each entry in the database file has the following items:

hashed-domain-name,IPv4 sinkhole address, IPv6 sinkhole address, sinkhole FQDN, ID, action

where:

- **hashed-domain-name** is a hashed value of the disallowed domain name (64 hexadecimal characters). The hash method and hash key that you use to produce the hashed domain value are needed when you configure DNS filtering with the Junos OS CLI.
 - **IPv4 sinkhole address** is the address of the DNS sinkhole server for IPv4 DNS requests.
 - **IPv6 sinkhole address** is the address of the DNS sinkhole server for IPv6 DNS requests.
 - **sinkhole FQDN** is the fully qualified domain name of the DNS sinkhole server.
 - **ID** is a 32-bit number that uniquely associates the entry with the hashed domain name.
 - **action** is the action to apply to a DNS request that matches the disallowed domain name. If you enter **replace**, the MX Series router sends the client a DNS response with the IP address or FQDN of the DNS sinkhole server. If you enter **report**, the DNS request is logged and then sent to the DNS server.
4. In the last line of the file, include the file hash, which you calculate by using the same key and hash method that you used to produce the hashed domain names.
 5. Save the database files on the Routing Engine in the `/var/db/url-filterd` directory.
 6. Validate the domain filter database file.

```
user@host> request services web-filter validate dns-filter-file-name filename hash-key key-string hash-method
hash-method-name
```

7. If you make any changes to the database file, apply the changes.

```
user@host> request services web-filter update dns-filter-database filename
```

How to Configure a DNS Filter Profile

A DNS filter profile includes general settings for filtering DNS requests for disallowed website domains, and includes up to 32 templates. The template settings apply to DNS requests on specific uplink and downlink logical interfaces or routing instances, or to DNS requests from specific source IP address prefixes, and override the corresponding settings at the DNS profile level. You can configure up to eight DNS filter profiles.

To configure a DNS filter profile:

1. Configure the name for a DNS filter profile:

```
[edit]
user@host# edit services web-filter profile profile-name
```

The maximum number of profiles is 8.

2. Configure the interval for logging per-client statistics for DNS filtering. The range is 0 through 60 minutes and the default is 5 minutes.

```
[edit services web-filter profile profile-name]
user@host# set global-dns-stats-log-timer minutes
```

3. Configure general DNS filtering settings for the profile. These values are used if a DNS request does not match a specific template.

- a. Specify the name of the domain filter database to use when filtering DNS requests.

```
[edit services web-filter profile profile-name dns-filter]
user@host# set database-file filename
```

- b. (Optional) To limit DNS filtering to DNS requests that are destined for specific DNS servers, specify up to three IP addresses (IPv4 or IPv6).

```
[edit services web-filter profile profile-name dns-filter]
user@host# set dns-server [ ip-address ]
```

- c. Specify the format for the hash key.

```
[edit services web-filter profile profile-name dns-filter]
user@host# set hash-key ascii-text
```

- d. Specify the hash key that you used to create the hashed domain name in the domain filter database file.

```
[edit services web-filter profile profile-name dns-filter]
user@host# set hash-key key-string
```

- e. Specify the hash method that was used to create the hashed domain name in the domain filter database file.

```
[edit services web-filter profile profile-name dns-filter]
user@host# set hash-method hash-method-name
```

The only supported hash method is **hmac-sha2-256**.

- f. Configure the interval for logging statistics for DNS requests and for sinkhole actions performed for each customer IP address. The range is 1 through 60 minutes and the default is 5 minutes.

```
[edit services web-filter profile profile-name dns-filter]
user@host# set statistics-log-timer minutes
```

- g. Configure the time to live while sending the DNS response after taking the DNS sinkhole action. The range is 0 through 86,400 seconds and the default is 1800.

```
[edit services web-filter profile profile-name dns-filter]
user@host# set dns-resp-ttl seconds
```

- h. Configure the level of subdomains that are searched for a match. The range is 0 through 10. A value of 0 indicates that subdomains are not searched.

```
[edit services web-filter profile profile-name dns-filter]
user@host# set wildcarding-level level
```

For example, if you set the **wildcarding-level** to 4 and the database file includes an entry for **example.com**, the following comparisons are made for a DNS request that arrives with the domain **198.51.100.0.example.com**:

- **198.51.100.0.example.com**: no match
- **51.100.0.example.com**: no match for one level down
- **100.0.example.com**: no match for two levels down
- **0.example.com**: no match for three levels down
- **example.com**: match for four levels down

4. Configure a template. You can configure a maximum of 8 templates in a profile. Each template identifies filter settings for DNS requests on specific uplink and downlink logical interfaces or routing instances, or for DNS requests from specific source IP address prefixes.

- a. Configure the name for the template.

```
[edit services web-filter profile profile-name]
user@host# set dns-filter-template template-name
```

- b. (Optional) Specify the client-facing logical interfaces (uplink) to which the DNS filtering is applied.

```
[edit services web-filter profile profile-name dns-filter-template template-name]
```

```
user@host# set client-interfaces client-interface-name
```

- c. (Optional) Specify the server-facing logical interfaces (downlink) to which the DNS filtering is applied.

```
[edit services web-filter profile profile-name dns-filter-template template-name]
user@host# set server-interfaces server-interface-name
```

- d. (Optional) Specify the routing instance for the client-facing logical interface to which the DNS filtering is applied.

```
[edit services web-filter profile profile-name dns-filter-template template-name]
user@host# set client-routing-instance client-routing-instance-name
```

- e. (Optional) Specify the routing instance for the server-facing logical interface to which DNS filtering is applied.

```
[edit services web-filter profile profile-name dns-filter-template template-name]
user@host# set server-routing-instance server-routing-instance-name
```

NOTE: If you configure the client and server interfaces or the client and server routing instances, implicit filters are installed on the interfaces or routing instances to direct DNS traffic to the services PIC for DNS filtering. If you configure neither the client and server interfaces nor the routing instances, you must provide a way to direct DNS traffic to the services PIC (for example, via routes).

- f. Specify the name of the domain filter database to use when filtering DNS requests.

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set database-file filename
```

- g. (Optional) To limit DNS filtering to DNS requests that are destined for specific DNS servers, specify up to three IP addresses (IPv4 or IPv6).

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set dns-server ip-address
```

- h. Specify the hash method that was used to create the hashed domain name in the domain filter database file.

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set hash-method hash-method-name
```

The only supported hash method is **hmac-sha2-256**.

- i. Specify the hash key that was used to create the hashed domain name in the domain filter database file.

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set hash-key key-string
```

- j. Configure the interval for logging statistics for DNS requests and for sinkhole actions performed for each customer IP address. The range is 1 through 60 minutes and the default is 5 minutes.

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set statistics-log-timer minutes
```

- k. Configure the time to live while sending the DNS response after taking the DNS sinkhole action. The range is 0 through 86,400 seconds and the default is 1800.

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set dns-resp-ttl seconds
```

- l. Configure the level of subdomains that are searched for a match. The range is 0 through 10. A value of 0 indicates that subdomains are not searched.

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set wildcarding-level level
```

For example, if you set the **wildcarding-level** to 4 and the database file includes an entry for **example.com**, the following comparisons are made for a DNS request that arrives with the domain **198.51.100.0.example.com**:

- **198.51.100.0.example.com**: no match
- **51.100.0.example.com**: no match for one level down
- **100.0.example.com**: no match for two levels down

- **0.example.com**: no match for three levels down
- **example.com**: match for four levels down

m. (Optional) Specify the response error code for SRV and TXT query types.

```
[edit services web-filter profile profile-name dns-filter-template template-name dns-filter]
user@host# set txt-resp-err-code (Noerror | Refused)
user@host# set srv-resp-err-code (Noerror | Refused)
```

n. Configure a term for the template. You can configure a maximum of 64 terms in a template.

```
[edit services web-filter profile profile-name dns-filter-template template-name]
user@host# set term term-name
```

o. (Optional) Specify the source IP address prefixes of DNS requests you want to filter. You can configure a maximum of 64 prefixes in a term.

```
[edit services web-filter profile profile-name dns-filter-template template-name term term-name]
user@host# set from src-ip-prefix source-prefix
```

p. Specify that the sinkhole action identified in the domain filter database is performed on disallowed DNS requests.

```
[edit services web-filter profile profile-name dns-filter-template template-name term term-name]
user@host# set then dns-sinkhole
```

How to Configure a Service Set for DNS Filtering

- Associate the DNS filter profile with a next-hop service set and enable logging for DNS filtering. The service interface can be an ms- or vms- interface Next Gen Services with MX-SPC3 services card), or it can be an aggregated multiservices (AMS) interface.

```
[edit services service-set service-set-name]
user@host# set web-filter-profile profile-name
user@host# set syslog host hostname class urlf-logs
user@host# set next-hop-service inside-service-interface interface-name.unit-number
user@host# set next-hop-service outside-service-interface interface-name.unit-number
```

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, you can configure DNS filtering if you are running Next Gen Services with the MX-SPC3 services card. Next Gen Services are supported on MX240, MX480 and MX960 routers.

DNS Request Filtering System Logging Error Messages

IN THIS SECTION

- [System Logging for DNS Request Filtering Overview | 320](#)
- [DNS Match-Event Syslog Format | 321](#)
- [Reason Mask Values & Interpretations for DNS Filtering | 324](#)
- [Per-Term Statistics Syslog Format | 325](#)
- [DNS Filtering Disallow-List File Add/Change Syslog Format | 326](#)
- [DNS Filtering Summary Report Statistics Syslog Format | 328](#)
- [DNS Filtering Per-Client-IP Statistics Syslog Format | 328](#)

The message format for system logs related to DNS request filtering differs slightly for the Next Gen Services MX-SPC3 services card versus early services cards. This topic describes the differences in the DNS request filtering related system log messages and provides a description of all fields in these messages.

System Logging for DNS Request Filtering Overview

Next Gen Services DNS request filtering system logging generates these events:

1. DNS match events (DNS_SR_MATCH_EVENT)
 - a. A single syslog is generated for each DNS match to the list of filtered domains.
2. Per-term statistics (DNS_SR_CUSTOMER_STATS)
 - a. Each term in the template represents a customer, enabling you to collect per-customer statistics.
 - b. You can configure the interval in which you want to collect statistics in each template.

- 3. You can report an event each time a DNS disallow-list file is added or updated (DNS_SR_FILE_UPDATE_NOTICE)
- 4. You can collect per-PIC Summary report statistics (DNS_SR_REPORT_STATS)
 - a. Statistics are generated every 5 minutes. This interval value is not configurable.
 - b. These stats are generated per-PIC basis.

NOTE: To enable these logs you must configure a syslog for each **service-set** for which you've configured dns-filtering.

All system log messages for Next Gen Services are configured at the **service-set** level using the following statement:

```
user@host# edit services service-set service-set-name syslog
```

To collect DNS request filtering system log messages, include **urlf** in the **local-category** statement:

```
[edit services service-set ss1 syslog]  
user@host# set local-category urlf
```

- 5. You can collect per-client IP statistics (DNS_SR_CLIENT_IP_STATS)
 - a. This statistics are generated per-profile.
 - b. The interval for collecting these statistics is configurable per-profile.

DNS Match-Event Syslog Format

NOTE: System system log messages for Next Gen Services DNS request filtering doesn't include the FPC slot/PIC slot and UTC time.

Table 42 on page 321 describes the fields contained in DNS request filtering match events.

Table 42: DNS-Match-Event Syslog Format

Field Name	Description	Example
Time Stamp	Time when log entry was generated	Oct 27 10:04:19

Table 42: DNS-Match-Event Syslog Format (*continued*)

Field Name	Description	Example
Router Name	Host name of the router generating the record	Jnpr-router-01
Log Handle	Log handle to identify the log category	junos-url-filter
Match	Indicates a DNS match was detected.	JSERVICES_URLF_MATCH_EVENT: DNS_SR_MATCH_EVENT
Tag	Log-prefix configured	Tag=<value>
svc-set-name	Service-set name	svc-set-name=<value>
ID	ID assigned to the domain name (Size of ID is assumed to be a 32-bit number)	ID=12345
IP_Src	Source IP	IP_Src=10.1.5.72
IP_Dst	Destination IP (DNS resolver)	IP_Dst=10.1.1.10
Src_Prt	Source Port	Src_Prt=37344
Dst_Prt	Destination Port	Dst_Prt=53
Sinkhole_IP	IP of sinkhole server from Domain Name Input List	Sinkhole_IP=10.1.50.64
Sinkhole_IPv6	IP of IPv6 sinkhole server from Domain Name Input List	Sinkhole_IPv6=8001:1002: 1003:1004:1005:1006:1007:1008
Sinkhole_fqdn	Sinkhole FQDN	Sinkhole_fqdn=NA
Count	Counter for match events to accommodate identical event records	Count=54
Replaced	Designates replacement of response domain (i.e. sinkholing)	Replaced=Y

Table 42: DNS-Match-Event Syslog Format (*continued*)

Field Name	Description	Example
Reason_Mask	Reason for action (if Replaced=N) [See table below for bit position enumeration]	Reason_Mask=0x0
QType	Query Type of the DNS request (A, AAAA, MX, CNAME, SRV, TXT)	QType=A
Profile	Profile Name [The Web filter profile name as configured]	Profile=profile_01
Template	Template Name [The DNS filter template name as configured]	Template=template_01
Term	Term Name [The DNS filter term name as configured]	Term=term_01
Time	UNIX timestamp	Time=Wed Dec 20 12:25:24 2017

Here's an example of MX-SPC3 DNS filtering syslog format:

Feb 20 17:06:36 ce-bras-mx480-o junos-url-filter: JSERVICES_URLF_MATCH_EVENT: DNS_SR_MATCH_EVENT, Tag=tag, svc-set-name= s1, ID=1235, IP_SRC=2.2.2.3, IP_DST=101.10.10.100, SRC_PRT=34342, DST_PRT=53, Sinkhole_IP=1.1.1.1, Sinkhole_IPv6=NA, Sinkhole_fqdn=NA, Count=9, Replaced=Y, Reason_Mask=0x0, QType=A, Profile=webf-prof-1, Template=dnsf-temp-1, Term=dnsf-term-1, Time=Tue Jan 23 13:45:52 2018

Here's an example of MS-MPC DNS filtering syslog format:

Jan 23 13:45:52 cliq (FPC Slot 1, PIC Slot 1) 2018-01-23 21:45:52: {s1}[jservices-urlf]: JSERVICES_URLF_MATCH_EVENT: DNS_SR_MATCH_EVENT ID=1235, IP_SRC=2.2.2.3, IP_DST=101.10.10.100, SRC_PRT=34342, DST_PRT=53, Sinkhole_IP=1.1.1.1, Sinkhole_IPv6=NA, Sinkhole_fqdn=NA, Count=9, Replaced=Y, Reason_Mask=0x0, QType=A, Profile=webf-prof-1, Template=dnsf-temp-1, Term=dnsf-term-1, Time=Tue Jan 23 13:45:52 2018

Reason Mask Values & Interpretations for DNS Filtering

Table 43 on page 324 describes the reason mask value fields and interpretations for MX Next Gen Services DNS filtering.

Table 43: Reason Mask Values & Interpretations for DNS Filtering

Bit Position	Hex Value	Interpretation	Additional Comments
	0x0	Replaced	
0	0x1	Reason Other	<i>Examples:</i> Fragmented packets, malformed packets
1	0x2	Not a supported DNS request type	<i>Examples:</i> SRV, TXT
2	0x4	Indicator action set to "Report-Only"	This is to enable testing of new indicators before putting them into Production.
3	0x8	Replace A/AAAA record error	
4	0x10	Replacement information not available	The domain name entry is marked "replace" but the sinkhole IP is not provided.

Here's an example of MX Next Gen Services syslog format for DNS filtering showing the reason mask and interpretation:

```
Feb 20 17:06:36 ce-bras-mx480-o junos-url-filter: JSERVICES_URLF_MATCH_EVENT:
DNS_SR_MATCH_EVENT, Tag=tag, svc-set-name= s1, ID=1235, IP_SRC=2.2.2.3, IP_DST=101.10.10.100,
SRC_PRT=34342, DST_PRT=53, Sinkhole_IP=1.1.1.1, Sinkhole_IPv6=NA, Sinkhole_fqdn=NA, Count=9,
Replaced=Y, Reason_Mask=0x0, QType=A, Profile=webf-prof-1, Template=dnsf-temp-1,
Term=dnsf-term-1, Time=Tue Jan 23 13:45:52 2018
```

Here's an example of MS-MPC DNS filtering syslog format:

```
Jan 23 13:45:52 cliq (FPC Slot 1, PIC Slot 1) 2018-01-23 21:45:52: {s1}[jservices-urlf]:
JSERVICES_URLF_MATCH_EVENT: DNS_SR_MATCH_EVENT ID=1235, IP_SRC=2.2.2.3,
```

IP_DST=101.10.10.100, SRC_PRT=34342, DST_PRT=53, Sinkhole_IP=1.1.1.1, Sinkhole_IPv6=NA, Sinkhole_fqdn=NA, Count=9, Replaced=Y, Reason_Mask=0x0, QType=A, Profile=webf-prof-1, Template=dnstf-temp-1, Term=dnstf-term-1, Time=Tue Jan 23 13:45:52 2018

Per-Term Statistics Syslog Format

Table 44 on page 325 describes the fields for MX Next Gen Services DNS filtering per-term statistics syslog format.

Table 44: Per-Term Statistics Syslog Format

Field Name	Description	Example
Time Stamp	Time when log entry was generated	Oct 27 10:04:17
Router Name	Host name of the router generating the record	Jnpr-router-01
Log Handle	Log handle to identify the log category	junos-url-filter
Match	A term(customer) statistics record	JSERVICES_URLF_CUSTOMER_STATS: DNS_SR_CUSTOMER_STATS
Tag	Log-prefix configured	Tag=<value>
svc-set-name	Service-set name	svc-set-name=<value>
Profile	Profile Name [The Web filter profile name as configured]	Profile=profile_01
Template	Template Name [The DNS filter template name as configured]	Template=template_01
Term	Term Name [The DNS filter term name as configured]	Term=term_01
Packets_Processed	Total DNS Requests Processed	Requests_Processed=200

Table 44: Per-Term Statistics Syslog Format (*continued*)

Field Name	Description	Example
DNS_UDP_Packets_Processed	DNS UDP Requests Processed	DNS_UDP_Requests_Processed=98
DNS_TCP_Packets_Processed	DNS TCP Requests Processed	DNS_TCP_Requests_Processed=35
DNS_UDP_Requests_sinkholed	DNS UDP Requests sink-holed	DNS_UDP_Requests_Sinkholed =50
DNS_TCP_Requests_sinkholed	DNS TCP Requests sink-holed	DNS_TCP_Requests_Sinkholed =50
DNS_UDP_Requests_reported	DNS UDP Requests reported	DNS_UDP_Requests_Reported =50
DNS_TCP_Requests_reported	DNS TCP Requests reported	DNS_TCP_Requests_Reported =50
Time	UNIX timestamp	Time=Wed Dec 20 12:25:24 2017
Count	Counter to accommodate identical event records	Count=10

Here's an example of MX-SPC3 DNS filtering syslog format for per-term statistics:

Feb 25 14:25:45 curve junos-url-filter: JSERVICES_URLF_CUSTOMER_STATS: DNS_SR_CUSTOMER_STATS, Tag , svc-set-name s1, Profile=DNS_CUSTOMER-A, Template=DNS_CUSTOMER-A, Term=DNS_CUSTOMER-A, Requests_Processed=0, DNS_UDP_Requests_Processed=0, DNS_TCP_Requests_Processed=0, DNS_UDP_Requests_Sinkholed=0, DNS_TCP_Requests_Sinkholed=0, DNS_UDP_Requests_Reported=0, DNS_TCP_Requests_Reported=0, Time=Mon Feb 25 14:25:45 2019, Count=13

Here's an example of MS-MPC DNS filtering syslog format:

Mar 8 12:16:05 iphone3gs (FPC Slot 5, PIC Slot 0) 2019-03-08 20:16:04: {ATT-Zone5}[jsservices-urlf]: JSERVICES_URLF_CUSTOMER_STATS: DNS_SR_CUSTOMER_STATS, Profile=ATT-Profile-5-Zone5, Template=ATT-Profile-5-Zone5-Area1, Term=ATT-Profile-5-Zone5-Area1-Customer3, Requests_Processed=0, DNS_UDP_Requests_Processed=0, DNS_TCP_Requests_Processed=0, DNS_UDP_Requests_Sinkholed=0, DNS_TCP_Requests_Sinkholed=0, DNS_UDP_Requests_Reported=0, DNS_TCP_Requests_Reported=0, Time=Fri Mar 08 12:16:05 2019, Count=111

DNS Filtering Disallow-List File Add/Change Syslog Format

[Table 45 on page 327](#) describes the fields for MX Next Gen Services DNS filtering disallow-list file additions and updates syslog format.

Table 45: Disallow-List File Add/Change Syslog Format

Field Name	Description	Example
Time Stamp	Time when log entry was generated	Oct 27 10:04:17
Router Name	Host name of the router generating the record	Jnpr-router-01
Log Handle	Log handle to identify the log category	junos-url-filter
Match	The domain disallow-list file updated for the template. .	JSERVICES_URLF_FILE_UPDATE_NOTICE: DNS_SR_FILE_UPDATE_NOTICE
Tag	Log-prefix configured	Tag=<value>
svc-set-name	Service-set name	svc-set-name=<value>
File Name	Name of the file	File_Name=shdb.txt
File Version	Version of the file	File_Version=20170314_01
Updated	File Update Time	Domain_Filter_File_Updated=Fri Oct 27 10:56:42 2017
Profile	Profile Name [The Web filter profile name as configured]	Profile=profile_01
Template	Template Name [The DNS filter template name as configured]	Template=template_01
Domains	Number of Domains in the file	Domains=12
Report-Only-Domains	Number of Report-Only domains in the file	Report_Only_Domains=3

Here's an example of the syslog format for MX-SPC3 DNS filtering disallow-list add/change file updates:

**Feb 25 14:36:47 curve junos-url-filter: JSERVICES_URLF_FILE_UPDATE_NOTICE:
DNS_SR_FILE_UPDATE_NOTICE, Tag=, svc-set-name=s1, File_Name=test_dns_sink.txt,**

File_Version=20180911 01, Domain_Filter_File_Updated=Mon Feb 25 14:36:47 2019
Profile=DNS_CUSTOMER-A, Template=DNS_CUSTOMER-A, Domains=18, Report_Only_Domains=0

Here's an example of the syslog format for DNS filtering disallow-list file changes with the MS-MPC services card:

Jan 23 13:34:34 cliq (FPC Slot 1, PIC Slot 1) 2018-01-23 21:34:33: {s1}[jservices-urlf]:
JSERVICES_URLF_FILE_UPDATE_NOTICE: DNS_SR_FILE_UPDATE_NOTICE, File_Name=dnsf1_hashed.txt,
File_Version=20170314_01, Domain_Filter_File_Updated=Tue Jan 23 13:34:34 2018 Profile=webf-prof-1,
Template=dnsf-temp-1, Domains=4, Report_Only_Domains=1

DNS Filtering Summary Report Statistics Syslog Format

Summary report statistics syslog format Stats will be reported in syslog with the following format:

Here's an example summary report syslog message for MX-SPC3 Next Gen Services DNS filtering:

Feb 25 11:50:39 curve junos-url-filter: JSERVICES_URLF_REPORT_STATS: DNS_SR_REPORT_STATS,
Tag=, svc-set-name=s1, TCP_DNS_Packets=0, TCP_DNS_Non_Segmented=0, TCP_DNS_Segmented=0,
Count=1

Here's an example summary report syslog message for MS-MPC services card DNS filtering:

Mar 8 12:20:41 iphone3gs (FPC Slot 5, PIC Slot 1) 2019-03-08 20:20:40: {ATT-Zone1}[jservices-urlf]:
JSERVICES_URLF_REPORT_STATS: DNS_SR_REPORT_STATS, TCP_DNS_Packets=0,
TCP_DNS_Non_Segmented=0, TCP_DNS_Segmented=0, Count=169

DNS Filtering Per-Client-IP Statistics Syslog Format

[Table 46 on page 328](#) describes the syslog fields for MX-SPC3 DNS filtering per-client-IP statistics that is reported per-PIC, per-profile for all known client IP addresses known to the system.

Table 46: Per-Client-IP Statistics Syslog Format

Field Name	Description	Example
Time Stamp	Time when log entry was generated	Oct 27 10:04:17
Router Name	Host name of the router generating the record	Jnpr-router-01
Log Handle	Log handle to identify the log category	junos-url-filter
Match	Log for per-Client IP stats	JSERVICES_URLF_CLIENT_IP_STATS: DNS_SR_CLIENT_IP_STATS

Table 46: Per-Client-IP Statistics Syslog Format (*continued*)

Field Name	Description	Example
Tag	Log-prefix configured	Tag=<value>
svc-set-name	Service-set name	svc-set-name=<value>
Client-IP	IP address of the client	Client-IP=1.1.1.1
Profile	Profile Name [The Web filter profile name as configured]	Profile=profile_01
Template	Template Name [The DNS filter template name as configured]	Template=template_01
Term	Term Name [The DNS filter term name as configured]	Term=term_01
A_Req	DNS A-Record Requests Processed	A_Req=10
AAAA_Req	DNS AAAA-Record Requests Processed	AAAA_Req=10
MX_Req	DNS MX-Record Requests Processed	MX_Req=4
CNAME_Req	DNS CNAME-Record Requests Processed	CNAME_Req=4
SRV_Req	DNS SRV-Record Requests Processed	SRV_Req=4
TXT_Req	DNS TXT-Record Requests Processed	TXT_Req=4
ANY_Req	DNS ANY-Record Requests Processed	ANY_Req=4
A_Req_SH	DNS A-Record Requests sink-holed	A_Req_SH =5
AAAA_Req_SH	DNS AAAA-Record Requests sink-holed	AAAA_Req_SH=5

Table 46: Per-Client-IP Statistics Syslog Format (*continued*)

Field Name	Description	Example
MX_Req_SH	DNS MX-Record Requests Sink-holed	MX_Req_SH=4
CNAME_Req_SH	DNS CNAME-Record Requests Sink-holed	CNAME_Req_SH=4
SRV_Req_SH	DNS SRV-Record Requests Sink-holed	SRV_Req_SH=4
TXT_Req_SH	DNS TXT-Record Requests Sink-holed	TXT_Req_SH=4
ANY_Req_SH	DNS ANY-Record Requests Sink-holed	ANY_Req_SH=4
Req_Rep	DNS Requests reported	Req_Rep=5

Here's an example per-client-IP-statitics for MX-SPC3 DNS filtering:

Feb 25 11:50:39 curve junos-url-filter: JSERVICES_URLF_CLIENT_IP_STATS: DNS_SR_CLIENT_IP_STATS, Tag=tag, svc-set-name=s1, Client-IP=2.2.2.3, Profile=webf-prof-1, Template=dnsf-temp-1, Term=dnsf-term-1, A_Req=0, AAAA_Req=0, MX_Req=0, CNAME_Req=0, SRV_Req=0, TXT_Req=0, ANY_Req=2, A_Req_SH=0, AAAA_Req_SH=0, MX_Req_SH=0, CNAME_Req_SH=0, SRV_Req_SH=0, TXT_Req_SH=0, ANY_Req_SH=0, Req_Rep=2

Here's an example syslog message for DNS filtering client-IP statistics on MS-MPC services cards:

Mar 7 17:58:54 iphone3gs (FPC Slot 5, PIC Slot 3) 2019-03-08 01:58:54: {dns}[jsservices-urlf]: JSERVICES_URLF_CLIENT_IP_STATS: DNS_SR_CLIENT_IP_STATS, Client-IP=2004:db0:2228:8001::1, Profile=dns-profile1, Template=dns1, Term=3, A_Req=19, AAAA_Req=19, MX_Req=0, CNAME_Req=0, SRV_Req=0, TXT_Req=0, ANY_Req=0, A_Req_SH=19, AAAA_Req_SH=19, MX_Req_SH=0, CNAME_Req_SH=0, SRV_Req_SH=0, TXT_Req_SH=0, ANY_Req_SH=0, Req_Rep=0

RELATED DOCUMENTATION

7

PART

URL Filtering

URL Filtering | 332

URL Filtering

IN THIS CHAPTER

- [URL Filtering Overview | 332](#)
- [Configuring URL Filtering | 335](#)

URL Filtering Overview

IN THIS SECTION

- [URL Filter Database File | 333](#)
- [URL Filter Profile Caveats | 334](#)

You can use URL filtering to determine which Web content is not accessible to users.

Components of this feature include the following:

- URL filter database file
- Configuration of one or more templates (up to eight per profile)
- URL Filter Plug-in (jservices-urlf)
- URL filtering daemon (url-filterd)

The URL filter database file is stored on the Routing Engine and contains all the disallowed URLs. Configured *templates* define which traffic to monitor, what criteria to match, and which actions to take. You configure the templates and the location of the URL filter database file in a *profile*.

Starting in Junos OS Release 17.2R2 and 17.4R1, for Adaptive Services, you can disable the filtering of HTTP traffic that contains an embedded IP address (for example, `http://10.1.1.1`) belonging to a disallowed domain name in the URL filter database. Starting in Junos OS Release 19.3R2, this same functionality is supported for Next Gen Services on MX240, MX480, and MX960.

To enable the URL filtering feature, you must configure **jservices-urlf** as the **package-name** at the **[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]** hierarchy level. Once enabled, **jservices-urlf** maintains the URL filtering profile and receives all traffic to be filtered, the filtering criteria, and the action to be taken on the filtered traffic.

The URL filtering daemon (**url-filterd**), which also resides on the Routing Engine, resolves the domain name of each URL in the URL filter database to a list of IPv4 and IPv6 addresses. It then downloads the list of IP addresses to the service PIC, which runs **jservices-urlf**. Then **url-filterd** interacts with the Dynamic Firewall process (**dfwd**) to install filters on the Packet Forwarding Engine to punt the selected traffic from the Packet Forwarding Engine to the service PIC.

As new HTTP and HTTPS traffic reaches the router, a decision is made based on the information in the URL filter database file. The filtering rules are checked and either the router accepts the traffic and passes it on or blocks the traffic. If the traffic is blocked, one of the following configured actions is taken:

- An HTTP redirect is sent to the user.
- A custom page is sent to the user.
- An HTTP status code is sent to the user.
- A TCP reset is sent.

Accept is also an option. In this case, the traffic is not blocked.

For more details on the URL filtering feature, see the following sections:

URL Filter Database File

The URL filter database file contains entries of URLs and IP addresses. Create the URL filter database file in the format indicated in [Table 47 on page 333](#) and locate it on the Routing Engine in the **/var/db/url-filterd** directory.

Table 47: URL Filter Database File Format

Entry	Description	Example
FQDN	Fully qualified domain name.	www.badword.com/jjj/bad.jpg
URL	Full string URL without the Layer 7 protocol.	www.yahoo.com/*badword*/
IPv4 address	HTTP request on a specific IPv4 address.	10.1.1.199
IPv6 address	HTTP request on a specific IPv6 address.	1::1

You must specify a custom URL filter database in the profile. If needed, you can also assign a custom URL filter database file with any template, and that database takes precedence over the database configured at the profile level.

If you change the contents of the URL filter database file, use the **request services (url-filter | web-filter) update** command. Other commands to help maintain the URL filter database file include the following:

- **request services (url-filter | web-filter) delete**
- **request services (url-filter | web-filter) force**
- **request services (url-filter | web-filter) validate**

URL Filter Profile Caveats

The URL filter profile consists of from one to eight templates. Each template consists of a set of configured logical interfaces where traffic is monitored for URL filtering and one or more terms.

A *term* is a set of match criteria with actions to be taken if the match criteria is met. You must configure at least one term to configure URL filtering. Each term consists of a **from** statement and a **then** statement, where the **from** statement defines the source IP prefixes and destination ports that are monitored. The **then** statement specifies the action to be taken. If you omit the **from** statement, any source IP prefix and any destination port are considered to match. But you can omit only one **from** statement per template or per profile.

Example configuration of multiple terms without from statements

```
template1 {
  client-interfaces [ xe-4/0/3.35 xe-4/0/3.36 ];
  server-interfaces xe-4/0/0.31;
  dns-source-interface xe-4/0/0.1;
  dns-routing-instance data_vr;
  routing-instance data_vr2;
  dns-server 50.0.0.3;
  dns-retries 3;
  url-filter-database url_database.txt;
  term term1 {
    then {
      tcp-reset;
    }
  }
  term term2 {
    then {
```



```
        redirect-url www.google.com;
    }
}
}
```

If you omit more than one **from** statement per template, you will get the following error message on commit:

```
URLFD_CONFIG_FAILURE: Configuration not valid:
Cannot have two wild card terms in template template1
error: configuration check-out failed
```

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, this same functionality is supported for Next Gen Services on MX240, MX480, and MX960.
17.2R2	Starting in Junos OS Release 17.2R2 and 17.4R1, for Adaptive Services, you can disable the filtering of HTTP traffic that contains an embedded IP address (for example, http://10.1.1.1) belonging to a disallowed domain name in the URL filter database.

RELATED DOCUMENTATION

[Configuring URL Filtering](#) | 335

request services url-filter update url-filter-database file

request services url-filter force dns-resolution

request services url-filter delete gencfg-data

request services url-filter validate

Configuring URL Filtering

URL filtering is configured on a service PIC. The interfaces you are dealing with are services interfaces (which use the **ms** prefix) or aggregated multiservices (AMS) interfaces (which use the **ams** prefix). For

more information on AMS interfaces, see the *Adaptive Services Interfaces User Guide for Routing Devices* starting with *Understanding Aggregated Multiservices Interfaces*.

To configure the URL filtering feature, you must first configure **jservices-urlf** as the **package-name** at the **[edit chassis fpc slot-number pic pic-number adaptive-services service-package extension-provider]** hierarchy level. For more information on configuring the **extension-provider package package-name** configuration statement, see the *package (Loading on PIC)* statement.

A URL filtering *profile* is a collection of templates. Each template consists of a set of criteria that defines which URLs are disallowed and how the recipient is notified.

To configure the URL profile:

1. Assign a name to the URL profile.

```
[edit]
user@host# edit services (web-filter | url-filter) profile profile-name
```

Starting in Junos OS Release 18.3R1, for Adaptive Services, configure the profile at the **[edit services web-filter]** hierarchy level. Before Junos OS Release 18.3R1, configure the profile at the **[edit services url-filter]** hierarchy level. Starting in Junos OS Release 19.3R2, this same functionality is available for Next Gen Series on MX240, MX480, and MX960.

2. Specify the name of the URL filter database to use.

```
[edit services (web-filter | url-filter) profile profile-name]
user@host# set url-filter-database filename
```

3. Configure one or more templates for the profile.

To configure each template:

- a. Name the template.

```
[edit services (web-filter | url-filter) profile profile-name]
user@host# set (url-filter-template template-name | template template-name)
```

NOTE: Starting in Junos OS Release 18.3R1, configure the template with the **url-filter-template** statement. Before Junos OS Release 18.3R1, configure the template with the **template** statement.

- b. Go to that new template hierarchy level.

```
[edit services (web-filter | url-filter) profile profile-name]
user@host# edit (url-filter-template template-name | template template-name)
```

- c. Specify the name of the URL filter database to use.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set url-filter-database filename
```

- d. Specify the loopback interface for which the source IP address is picked for sending DNS queries.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set dns-source-interface loopback-interface-name
```

- e. Disable the filtering of HTTP traffic that contains an embedded IP address (for example, http://10.1.1.1) belonging to a disallowed domain name in the URL filter database.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set disable-url-filtering
```

- f. Configure the DNS resolution time interval in minutes.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set dns-resolution-interval minutes
```

- g. Configure the number of retries for a DNS query in case the query fails or times out.

```
[edit services (web-filter | url-filter) profile profile-name]
user@host# set dns-retries number
```

- h. Specify the IP addresses (IPv4 or IPv6) of DNS servers to which the DNS queries are sent.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
```

```
user@host# set dns-server [ip-address]
```

- i. Specify the client-facing logical interfaces on which the URL filtering is configured.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set client-interfaces [ client-interface-name ]
```

- j. Specify the server-facing logical interfaces on which the URL filtering is configured.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set server-interfaces [ server-interface-name ]
```

- k. Specify the routing instance on which the URL filtering is configured.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set routing-instance routing-instance-name
```

- l. Specify the routing instance on which the DNS server is reachable.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# dns-routing-instance dns-routing-instance-name
```

4. Configure the term information.

Terms are used in filters to segment the policy or filter into small match and action pairs.

- a. Name the term.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# set term term-name
```

- b. Go to the new term hierarchy level.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name)]
user@host# edit term term-name
```

- c. Specify the source IP address prefixes for traffic you want to filter.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name) term term-name]
user@host# set from src-ip-prefix [prefix]
```

- d. Specify the destination ports for traffic you want to filter.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name) term term-name]
user@host# set from dest-port [port]
```

- e. Configure an action to take.

```
[edit services (web-filter | url-filter) profile profile-name (url-filter-template template-name | template
template-name) term term-name]
user@host# set then action
```

The action can be one of the following:

custom-page *custom-page*—Send a custom page string to the user.

http-status-code *http-status-code*—Send an HTTP status code to the user.

redirect-url *redirect-url*—Send an HTTP redirect to the user.

tcp-reset—Send a TCP reset to the user.

5. Associate the URL profile with a next-hop service set.

NOTE: For URL filtering, you must configure the service set as a next-hop service set.

```
[edit]
user@host# set services service-set service-set-name (web-filter-profile profile-name | url-filter-profile
profile-name)
```

```

user@host# set services service-set service-set-name next-hop-service inside-service-interface
interface-name.unit-number
user@host# set services service-set service-set-name next-hop-service outside-service-interface
interface-name.unit-number

```

NOTE: The service interface can also be of the **ams** prefix. If you are using **ams** interfaces at the **[edit services service-set *service-set-name*]** hierarchy level for the URL filter, you must also configure the **load-balancing-options hash-keys** statement at the **[edit interfaces *ams-interface-name* unit *number*]** hierarchy level. .

NOTE: Starting in Junos OS Release 18.3R1, configure the service set with the **web-filter-profile** statement. Before Junos OS Release 18.3R1, configure the service set with the **url-filter-profile** statement.

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, this same functionality is available for Next Gen Serices on MX240, MX480, and MX960.
18.3R1	Starting in Junos OS Release 18.3R1, for Adaptive Services. configure the profile at the [edit services web-filter] hierarchy level. Before Junos OS Release 18.3R1, configure the profile at the [edit services url-filter] hierarchy level.

RELATED DOCUMENTATION

[URL Filtering Overview](#) | [332](#)

Configuring Service Sets to be Applied to Services Interfaces

8

PART

Integration of Juniper Sky ATP and Web filtering on MX Routers

Integration of Juniper Sky ATP and Web filtering on MX Routers | 342

Integration of Juniper Sky ATP and Web filtering on MX Routers

IN THIS CHAPTER

- [Integration of Juniper ATP Cloud and Web filtering on MX Routers | 342](#)

Integration of Juniper ATP Cloud and Web filtering on MX Routers

IN THIS SECTION

- [Overview | 342](#)
- [Configuring the Web Filter Profile for Sampling | 347](#)

Overview

IN THIS SECTION

- [Benefits | 343](#)
- [Understanding Policy Enforcer and Juniper ATP Cloud | 343](#)
- [Security Intelligence \(SecIntel\) - Overview | 344](#)
- [Web Filtering \(URL-Filterd\) - Overview | 345](#)

Juniper Sky™ Advanced Threat Prevention (Juniper ATP Cloud) is integrated with MX series routers to protect all hosts in your network against evolving security threats by employing cloud-based threat detection software with a next-generation firewall system.

This topic provides an overview of Juniper ATP Cloud, Policy Enforcer, Security Intelligence, Web filtering, and their benefits when integrated on MX Series routers (MX240, MX480 and MX960).

Benefits

- Simplifies deployment and enhances the anti-threat capabilities when integrated with the MX routers.
- Delivers protection against “zero-day” threats using a combination of tools to provide robust coverage against sophisticated, evasive threats.
- Checks inbound and outbound traffic with policy enhancements that allow users to stop malware, quarantine infected systems, prevent data exfiltration, and disrupt lateral movement.
- Supports High Availability to provide uninterrupted service.
- Provides scalability to handle increasing loads that require more computing resources, increased network bandwidth to receive more customer submissions, and a large storage for malware.
- Provides deep inspection, actionable reporting, and inline malware blocking.

Understanding Policy Enforcer and Juniper ATP Cloud

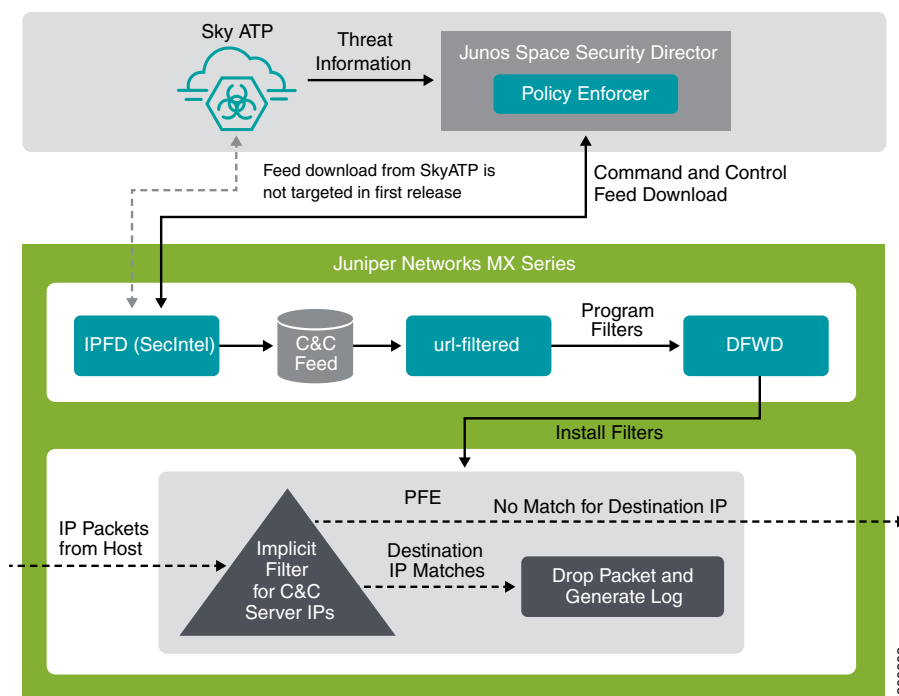
Juniper Networks Security Director comprises a feature called the Policy Enforcer (PE) that enables it to learn from threat conditions, automate the policy creation, and to dynamically deploy enforcement to Juniper devices in the network.

[Figure 10 on page 344](#) illustrates the traffic flow between the PE, the Juniper ATP Cloud, and the MX router which functions as a firewall.

- Policy Enforcer (PE) learns from threat conditions, automates the policy creation, and deploys enforcement to Juniper devices in the network.
- Juniper Sky™ Advanced Threat Prevention (Juniper ATP Cloud) protects all hosts in your network by employing cloud-based threat detection software with a next-generation firewall system.
- MX router fetches the threat intelligence feeds from Policy Enforcer (PE) and implements those policies to quarantine compromised hosts. It comprises of the following important components:
 - Security Intelligence process
 - Web Filtering process
 - Firewall process

Figure 10: System Architecture

ERROR: Unresolved graphic fileref="" not found in
 "///cmsxml/default/main/supplemental/STAGING/images/".



To understand the functionality of the system architecture consider the following example—if a user downloads a file from the Internet and that file passes through an MX firewall, the file can be sent to the Juniper ATP Cloud cloud for malware inspection (depending on your configuration settings.) If the file is determined to be malware, PE identifies the IP address and MAC address of the host that downloaded the file. Based on a user-defined policy, that host can be put into a quarantine VLAN or blocked from accessing the Internet.

MX Series routers (MX240, MX480, and MX960) can be integrated with the Juniper ATP Cloud to prevent compromised hosts (botnets) from communicating with command and control servers:

- Starting in Junos OS Release 18.4R1 with the Adaptive Services as an Inline security capability
- Starting in Junos OS Release 19.3R2 with the Next Gen Services as an Inline security capability

Security Intelligence (SecIntel) - Overview

The Security Intelligence process (IPFD), is responsible for downloading the security intelligence feeds and parsing from the feed connector or ATP Cloud cloud feed server. The IPFD process on the MX platforms fetches the command and control IPv4/IPv6 feeds from Policy Enforcer. C&C feeds are essentially a list of servers that are known command and control servers for botnets. The list also includes servers that are

known sources for malware downloads. The information thus fetched is saved in a file (`urlf_si_cc_db.txt`) created under the `/var/db/url-filterd` directory.

The file format of the disallowed IPs sent by IPFD to the web filtering process is as follows:

IPv4 address | IPv6 address, threat-level.

The ***threat-level*** is an integer ranging from 1 to 10 to indicate the threat level of files scanned for malware and for infected hosts. Here, 1 represents the lowest threat level and 10 represents the highest threat level.

For example: 178.10.19.20, 4

Here, 178.10.19.20 indicates the disallowed IP and 4 indicates the ***threat-level***.

The C&C feed database is synced onto the backup Routing Engine. IPFD then shares the information to the web filtering process (`url-filterd`). The web filtering process reads the file contents and configures the filters accordingly.

Configuring Security Intelligence to Download the CC Feed from Policy Enforcer

To download the command and control IPv4/IPv6 feeds from Juniper ATP Cloud/Policy Enforcer, include the **security-intelligence** statement at the **[edit services]** hierarchy as shown in the following example:

```
security-intelligence {
  authentication {
    auth-token 7QGSBL5ZRKR5UHUZ2X2R6QLHB656D5EN;
  }
  url https://10.92.83.245:443/api/v1/manifest.xml;
  traceoptions {
    file security-intelligence.log size 1g;
    level all;
    flag all;
  }
}
```

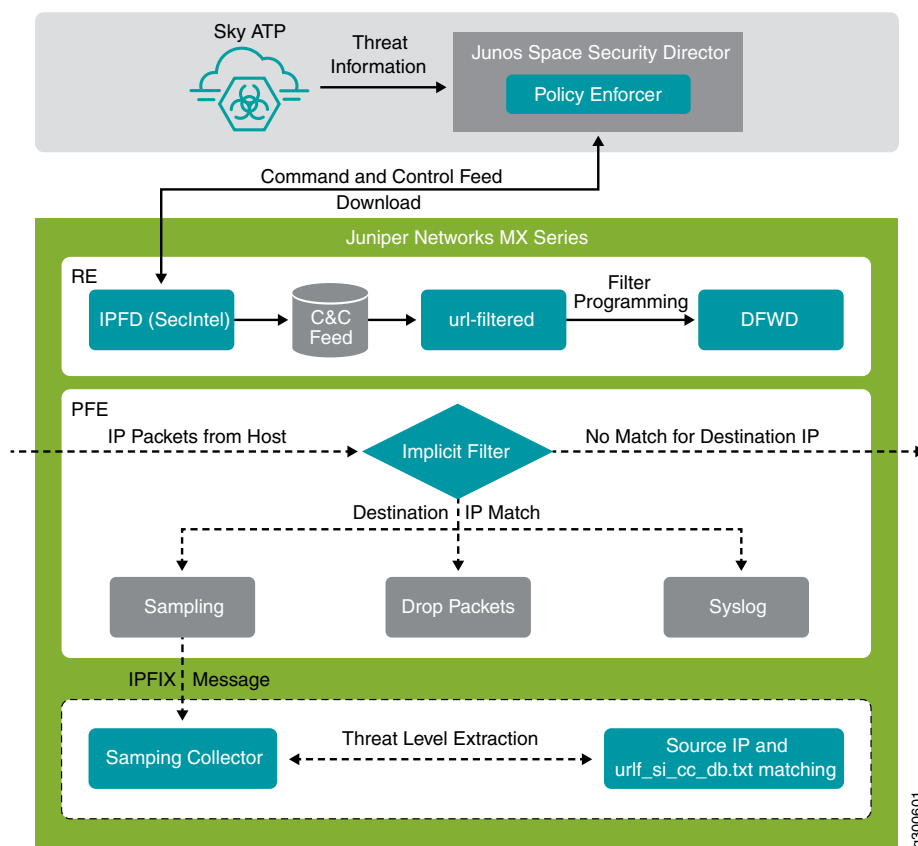
Web Filtering (URL-Filterd) - Overview

The web filtering process reads the file contents fetched from the IPFD and configures the filters on the Packet Forwarding Engine accordingly. The web filtering process enforces the command and control feeds by programming the filters in the Packet Forwarding Engine to block the packets destined to the blocked IP addresses and to generate logs for reporting the incident.

[Figure 11 on page 346](#) illustrates the way C&C feed is fetched by the IPFD and then processed by the web filtering process.

Figure 11: Web Filtering

ERROR: Unresolved graphic fileref="" not found in
 "///cmsxml/default/main/supplemental/STAGING/images/".



The web filter profile can have more than one templates. Each template consists of a set of configured logical interfaces for Web filtering and one or more terms. A term is a set of match criteria with actions to be taken if the match criteria is met. To configure the web filter profile to use dynamically fetched C&C feed, you can configure the **security-intelligence-policy** command under the **[edit services web-filter profile profile-name]** hierarchy level. You need not configure a term for a **security-intelligence-policy** based web filter profiles.

You can configure the following threat level actions for the web filter profile at the **edit web-filter profile profile-name security-intelligence-policy threat-level threat-level threat-action** hierarchy level:

- drop
- drop-and-log
- log

You can configure only one **threat-action** for each **threat level**. If the **threat-action** is not configured for a particular **threat level**, the default **threat-action** is **accept**.

SEE ALSO

[security-intelligence-policy | 667](#)

[security-intelligence | 665](#)

Configuring the Web Filter Profile for Sampling

Starting in Junos OS Release 19.3R1, web filtering process (url-filterd) supports inline sampling of packets as a threat level action. The packets are dropped, logged, and sampled based on the threat-action you configure. For scaled scenarios, sampling of packets is preferred over the logging option. Along with the existing threat level actions, you can configure the following threat level actions on the web filter profile at the **edit web-filter profile *profile-name* security-intelligence-policy threat-level *threat-level* threat-action** hierarchy level:

- **drop-and-sample**
- **drop-log-and-sample**
- **log-and-sample**
- **sample**

The inline flow monitoring samples the packets and sends the flow records in IPFIX format to a flow collector. You can derive the threat level for the sampled packets received at the external collector by matching the received IP from the sampled packets with the corresponding IP entry in **/var/db/url-filterd/urllf_si_cc_db.txt**. You can configure sampling using any of the following methods:

- Associate a sampling instance with the FPC on which the media interface is present at the **[edit chassis]** hierarchy level. If you are configuring sampling of IPv4 flows, IPv6 flows, or VPLS flows, you can configure the flow hash table size for each family.
- Configure the template properties for inline flow monitoring at the **[edit services flow-monitoring]** hierarchy level.
- Configure a sampling instance and associate the flow-server IP address, port number, flow export rate, and specify the collectors at the **[edit forwarding-options]** hierarchy level.

Associate a Sampling Instance with the FPC

To associate the defined instance with a particular FPC, MPC, or DPC, you include the **sampling-instance** statement at the **[edit chassis fpc number]** hierarchy level, as shown in the following example:

```

chassis {
  redundancy {
    graceful-switchover;
  }
  fpc 0 {
    pic0 {
      inline-services {
        bandwidth 10g;
      }
    }
  }
  pic 2 {
    inline-services {
      bandwidth 10g;
    }
  }
  pic 3 {
    inline-services {
      bandwidth 10g;
    }
  }
  sampling-instance 1to1;
  inline-services {
    flow-table-size {
      ipv4-flow-table-size 5;
      ipv6-flow-table-size 5;
    }
  }
}

```

Configure a Sampling Instance and Associate the Template With the Sampling Instance.

To configure the template properties for inline flow monitoring, include the following statements at the **edit services flow-monitoring** hierarchy level as shown in the following example:

```

services {
  flow-monitoring {
    version-ipfix {
      template ipv4 {
        flow-active-timeout 60;
      }
    }
    flow-inactive-timeout 60;
    template-refresh-rate {
      packets 48000;
      seconds 60;
    }
  }
}

```

```

    }
    option-refresh-rate {
        packets 48000;
        seconds 60;
    }
    ipv4-template;
    template ipv6 {
        flow-active-timeout 60;
        flow-inactive-timeout 60;
        template-refresh-rate {
            packets 48000;
            seconds 60;
        }
        ipv6-template;
    }
}

```

Configure the sample instance and associate the flow-server IP address and other parameters.

To configure a sampling instance and associate the flow-server IP address and other parameters. include the following statements at the **[edit forwarding-options]** hierarchy, as shown in the following example:

```

forwarding-options {
    sampling {
        traceoptions {
            file ipfix.log size 10k;
        }
    }
    instance {
        1to1 {
            input {
                rate 1;
            }
        }
    }
    family inet {
        output {
            flow-server 192.168.9.194;
            port 2055;;
            autonomous-system-type origin;
            version-ipfix {
                template {
                    ipv4;
                }
            }
        }
    }
}

```

```

    inline-jflow {
      source-address 192.168.9.195;
    }
  }
}
family inet6 {
  output {
    flow-server 192.168.9.194;
    port 2000;
    autonomous-system-type origin;
    version-ipfix {
      template {
        ipv6;
      }
    }
  }
  inline-jflow {
    source-address 192.168.9.195;
  }
}
}
}
}

```

Example: Configuring Web-filter Profile to Define Different Threat-Levels

```

web-filter {
  profile Profile1 ;
  security-intelligence-policy{
    file-type txt;
    threat-level 7 {
      threat-action {
        log-and-sample;
      }
    }
    threat-level 8 {
      threat-action {
        drop-log-and-sample;
      }
    }
    threat-level 10 {
      threat-action {
        drop-log-and-sample;
      }
    }
  }
}

```



```

}
threat-level 5{
  threat-action {
    drop-log-and-sample;
  }
}
threat-level 6 {
  threat-action {
    drop-log-and-sample;
  }
}
threat-level 9{
  threat-action {
    drop-log-and-sample;
  }
}
}
url-filter-template template1 {
  client-interfaces ge-0/0/4.0;
  client-routing-instance inet.0;
}
}
traceoptions {
  file webfilter_log size 1g;
  level all;
  flag all;
}
}
}

```

SEE ALSO

[security-intelligence-policy](#) | **667**

Configuring Traffic Sampling on MX, M and T Series Routers

9

PART

Aggregated Multiservices Interfaces

Enabling Load Balancing and High Availability Using Multiservices Interfaces | **353**

Enabling Load Balancing and High Availability Using Multiservices Interfaces

IN THIS CHAPTER

- [Understanding Aggregated Multiservices Interfaces for Next Gen Services | 353](#)
- [Configuring Aggregated Multiservices Interfaces | 359](#)
- [Configuring Load Balancing on AMS Infrastructure | 361](#)
- [Configuring Warm Standby for Services Interfaces | 365](#)

Understanding Aggregated Multiservices Interfaces for Next Gen Services

IN THIS SECTION

- [Aggregated Multiservices Interface | 353](#)
- [IPv6 Traffic on AMS Interfaces Overview | 356](#)
- [Member Failure Options and High Availability Settings | 357](#)
- [Warm Standby Redundancy | 358](#)

This topic provides an overview of using the Aggregated Multiservices Interfaces feature with the MX-SPC3 services card for Next Gen Services. It contains the following sections:

Aggregated Multiservices Interface

In Junos OS, you can combine multiple services interfaces to create a bundle of services interfaces that can function as a single interface. Such a bundle of interfaces is known as an *aggregated multiservices interface* (AMS), and is denoted as *amsN* in the configuration, where *N* is a unique number that identifies an AMS interface (for example, *ams0*). Starting in Junos OS Release 19.3R2, AMS interfaces are supported on the Next Gen Services MX-SPC3 services card.

AMS configuration provides higher scalability, improved performance, and better failover and load-balancing options.

An AMS configuration enables service sets to support multiple services PICs by associating an AMS bundle with a service set. For Next Gen Services, the MX-SPC3 services card supports up to two PICs and you can have a maximum of eight MX-SPC3 services cards in your chassis. This enables a Next Gen Services AMS bundle to have up to 16 services PICs as member interfaces and you can distribute services among the member interfaces.

Member interfaces are identified as mams in the configuration. The chassisd process in routers that support AMS configuration creates a mams entry for every multiservices interface on the router.

When you configure services options at the ams interface level, the options apply to all member interfaces (mams) for the ams interface.

The options also apply to service sets configured on services interfaces corresponding to the ams interface's member interfaces. All settings are per PIC. For example, session-limit applies per member and not at an aggregate level.

NOTE: You cannot configure services options at both the ams (aggregate) and member-interface level. If services options are configured on **vms-x/y/z**, they also apply to service sets on **mams-x/y/z**.

When you want services options settings to apply uniformly to all members, configure services options at the ams interface level. If you need different settings for individual members, configure services options at the member interface level.

NOTE: Per-member drop of traffic and per-member next-hop configuration is required for NAT64. For NAPT44, this per-member specification allows arbitrary hash keys, providing better load-balancing options to allow dynamic NAT operations to be performed. For NAT64, NAPT44, and dynamic NAT44, it is not possible to determine which member allocates the dynamic NAT address. To ensure that reverse flow packets arrive at the same member as the forward flow packets, pool-address-based routes are used to steer reverse flow packets.

NOTE: If you modify a NAT pool that is being used by a service set assigned to an AMS interface, you must deactivate and activate the service set before the NAT pool changes take effect.

Traffic distribution over the member interfaces of an AMS interface can occur in either a round-robin fashion or hash-based. You can configure the following hash key values to regulate the traffic distribution:

source-ip, **destination-ip** , and **protocol**. For services that require traffic symmetry, you must configure symmetrical hashing. Symmetrical hashing configuration ensures that both forward and reverse traffic is routed through the same member interface.

If the service set is applied on the Gigabit Ethernet or 10-Gigabit Ethernet interface (interface-style service set) that functions as the NAT inside interface, then the hash keys used for load balancing might be configured in such a way that the ingress key is set as destination IP address and the egress key is set as source IP address. Because the source IP address undergoes NAT processing, it is not available for hashing the traffic in the reverse direction. Therefore, load balancing does not happen on the same IP address and forward and reverse traffic does not map to the same PIC. With the hash keys reversed, load balancing occurs correctly.

With next-hop services, for forward traffic, the ingress key on the inside interface load -balances traffic, and for reverse traffic, the ingress key on the outside interface load -balances traffic or per-member next hops steer reverse traffic. With interface-style services, the ingress key load-balances forward traffic and the egress key load-balances forward traffic or per-member next hops steer reverse traffic. Forward traffic is traffic entering from the inner side of a service set and reverse traffic is traffic entering from the outer side of a service set. The forward key is the hash key used for the forward direction of traffic and the reverse key is the hash key used for the reverse direction of traffic (depends on whether it relates to interface services or next-hop services style.)

With stateful firewalls, you can configure the following combinations of forward and reverse keys for load balancing. In the following combinations presented for hash keys, FOR-KEY refers to the forward key, REV-KEY denotes the reverse key, SIP signifies source IP address, DIP signifies destination IP address, and PROTO refers to protocol such as IP.

- FOR-KEY: SIP, REV-KEY: DIP
- FOR-KEY: SIP,PROTO REV-KEY: DIP, PROTO
- FOR-KEY: DIP, REV-KEY: SIP
- FOR-KEY: DIP,PROTO REV-KEY: SIP, PROTO
- FOR-KEY: SIP,DIP REV-KEY: SIP, DIP
- FOR-KEY: SIP,DIP,PROTO REV-KEY: SIP, DIP,PROTO

With static NAT configured as basic NAT44 or destination NAT44, and with stateful firewall configured or not, if the forward direction of traffic must undergo NAT processing, configure the hash keys as follows:

- FOR-KEY: DIP, REV-KEY: SIP
- FOR-KEY: DIP,PROTO REV-KEY: SIP, PROTO

If the reverse direction of traffic must undergo NAT processing, configure the hash keys as follows:

- FOR-KEY: SIP, REV-KEY: DIP
- FOR-KEY: SIP,PROTO REV-KEY: DIP, PROTO

With dynamic NAT configured, and with stateful firewall configured or not, only the forward direction traffic can undergo NAT. The forward hash key can be any combination of SIP, DIP, and protocol, and the reverse hash key is ignored.

NOTE: The Junos OS AMS configuration supports IPv4 and IPv6 traffic.

IPv6 Traffic on AMS Interfaces Overview

You can use AMS interfaces for IPv6 traffic. To configure IPv6 support for an AMS interface, include the **family inet6** statement at the **[edit interfaces *ams-interface-name* unit 1]** hierarchy level. When **family inet** and **family inet6** are set for an AMS interface subunit, the **hash-keys** is configured at service-set level for interface style and at IFL level for next-hop style.

When a member interface of an AMS bundle fails, traffic destined to the failed member is redistributed among the remaining active members. The traffic (flows or sessions) traversing through the existing active members is unaffected. If M members are currently active, the expected result is that only about $1/M$ fraction of the traffic (flows/sessions) is impacted because that amount of traffic is shifted from the failed member to remain active members. When the failed member interface comes back online, only a fraction of the traffic is redistributed to the new member. If N members are currently active, the expected result is that only about $1/(N+1)$ fraction of the traffic (flows/sessions) is impacted because that amount of traffic moves to the new restored member. The $1/M$ and $1/(N+1)$ values assume that the flows are uniformly distributed among members, because a packet-hash is used to load-balance and because traffic usually contains a typical random combination of IP addresses (or any other fields that are used as load-balancing keys).

Similar to IPv4 traffic, for IPv6 packets, an AMS bundle must contain members of only one services PIC type.

The number of flows distributed, in an ideal environment, can be $1/N$ in a best-case scenario when the N th member goes up or down. However, this assumption considers that the hash keys load-balance the real or dynamic traffic. For example, consider a real-world deployment where member A is serving only one flow, whereas member B is serving 10 flows. If member B goes down, then the number of flows disrupted is 10/11. The NAT pool-split behavior is designed to utilize the benefits of the rehash-minimization feature. The splitting of a NAT pool is performed for dynamic NAT scenarios (dynamic NAT, NAT64, and NAPT44).

If the original and redistributed flows are defined as follows:

- Member-original-flows—The traffic mapped to a member when all members are up.
- Member-redistributed-flows—The additional traffic mapped to a member when some other member fails. These traffic flows might need to be rebalanced when member interfaces come up and go down.

With the preceding definitions of the original and redistributed flows for member interfaces, the following observations apply:

- The member-original-flows of a member stay intact as long as that member is up. Such flows are not impacted when other members move between the up and down states.
- The member-redistributed-flows of a member can change when other members go up or down. This change of flows occurs because these additional flows need to be rebalanced among all active members. Therefore, the member-redistributed-flow can vary a lot based on other members going down or up. Although it might seem that when a member goes down, the flows on active-members are preserved, and that when a member goes up, flows on active-members are not preserved in an effective way, this behavior is only because of static or hash-based rebalancing of traffic among active members.

The rehash-minimization feature handles the operational changes in a member interface status only (such as member offline or member Junos OS reset). It does not handle changes in configuration. For example, addition or deletion, or activation and deactivation, of member interfaces at the **[edit interfaces amsN load-balancing-options member-interface mams-a/b/0]** hierarchy level requires the member PICs to be bounced. Twice NAT or hairpinning is not supported, similar to IPv4 support for AMS interfaces.

Member Failure Options and High Availability Settings

Because multiple service interfaces are configured as part of an AMS bundle, AMS configuration also provides for failover and high availability support. You can either configure one of the member interfaces as a backup interface that becomes active when any one of the other member interfaces goes down, or configure the AMS in such a way that when one of the member interfaces goes down, the traffic assigned to that interface is shared across the active interfaces.

The **member-failure-options** configuration statement enables you to configure how to handle traffic when a member interface fails. One option is to redistribute the traffic immediately among the other member interfaces. However, redistribution of traffic involves recalculating the hash tags, and might cause some disruption in traffic on all the member interfaces.

The other option is to configure the AMS to drop all traffic that is assigned to the failed member interface. With this you can optionally configure an interval, **rejoin-timeout**, for the AMS to wait for the failed interface to come back online after which the AMS can redistribute the traffic among other member interfaces. If the failed member interface comes back online before the configured wait time, traffic continues unaffected on all member interfaces, including the interface that has come back online and resumed the operations.

You can also control the rejoining of the failed interface when it comes back online. If you do not include the **enable-rejoin** statement in the **member-failure-options** configuration, the failed interface cannot rejoin the AMS when it comes back online. In such cases, you can manually rejoin that to the AMS by executing the **request interfaces revert interface-name** operational mode command.

The **rejoin-timeout** and **enable-rejoin** statements enable you to minimize traffic disruptions when member interfaces flap.

NOTE: When **member-failure-options** are not configured, the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

The **high-availability-options** configuration enables you to designate one of the member interfaces as a backup interface. The backup interface does not participate in routing operations as long as it remains a backup interface. When a member interface fails, the backup interface handles the traffic assigned to the failed interface. When the failed interface comes back online, it becomes the new backup interface.

In a many-to-one configuration (N:1), a single backup interface supports all other member interfaces in the group. If any of the member interfaces fails, the backup interface takes over. In this stateless configuration, data is not synchronized between the backup interface and the other member interfaces.

When both **member-failure-options** and **high-availability-options** are configured for an AMS, the **high-availability-options** configuration takes precedence over the **member-failure-options** configuration. If a second failure occurs before the failed interface comes back online to be the new backup, the **member-failure-options** configuration takes effect.

Warm Standby Redundancy

Starting in Junos OS Release 19.3R2, the N:1 warm standby option is supported on the MX-SPC3 if you are running Next Gen Services. Each warm standby AMS interface contains two members; one member is the service interface you want to protect, called the primary interface, and one member is the secondary (backup) interface. The primary interface is the active interface and the backup interface does not handle any traffic unless the primary interface fails.

To configure warm standby on an AMS interface, you use the **redundancy-options** statement. You cannot use the **load-balancing-options** statement in a warm standby AMS interface.

To switch from the primary interface to the secondary interface, issue the **request interface switchover amsN** command.

To revert to the primary interface from the secondary interface, issue the **request interface revert amsN** command.

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, AMS interfaces are supported on the Next Gen Services MX-SPC3 services card.
19.3R2	Starting in Junos OS Release 19.3R2, the N:1 warm standby option is supported on the MX-SPC3 if you are running Next Gen Services.

Configuring Aggregated Multiservices Interfaces

The aggregated multiservices (AMS) interface configuration in Junos OS enables you to combine services interfaces from multiple PICs to create a bundle of interfaces that can function as a single interface. You identify the PIC that you want to act as the backup.

1. Create an aggregated multiservices interface and add member interfaces. Starting in Junos OS Release 19.3R2, an MX-SPC3 Next Gen Services AMS interface can have up to 16 member interfaces with a maximum of 8 MX-SPC3 services cards with up to 2 PICs on each card. Starting with Junos OS Release 16.2, an MS-MPC AMS interface can have up to 36 member interfaces. In Junos OS Release 16.1 and earlier, an AMS interface can have a maximum of 24 member interfaces.

NOTE: The member interface format is **mams-*a*/*b*/0**, where *a* is the Flexible PIC Concentrator (FPC) slot number and *b* is the PIC slot number.

```
[edit interfaces]
user@host# set interface-name load-balancing-options member-interface mams-a/b/0
user@host# set interface-name load-balancing-options member-interface mams-a/b/0
```

For example on an MS-MPC, which can have up to four PICs:

```
[edit interfaces]
user@host# set ams1 load-balancing-options member-interface mams-1/1/0
user@host# set ams1 load-balancing-options member-interface mams-1/2/0
```

For example on an MX-SPC3, which can have up to two PICs:

```
[edit interfaces]
user@host# set ams1 load-balancing-options member-interface mams-1/0/0
user@host# set ams1 load-balancing-options member-interface mams-1/1/0
```

2. Configure logical units for the AMS interface.

```
[edit interfaces]
user@host# set interface-name unit logical-unit-number family family
user@host# set interface-name unit logical-unit-number family family
```

For example:

```
[edit interfaces]
user@host# set ams1 unit 1 family inet
user@host# set ams1 unit 2 family inet6
```

3. Configure member failure options.

```
[edit interfaces interface-name]
user@host# set load-balancing-options member-failure-options drop-member-traffic rejoin-timeout seconds
user@host# set load-balancing-options member-failure-options drop-member-traffic enable-rejoin
```

For example:

```
[edit interfaces ams1]
user@host# set load-balancing-options member-failure-options drop-member-traffic rejoin-timeout 1000
user@host# set load-balancing-options member-failure-options drop-member-traffic enable-rejoin
```

4. Configure the preferred backup.

```
[edit interfaces interface-name]
user@host# set load-balancing-options high-availability-options many-to-one preferred-backup
preferred-backup
```

For example:

```
[edit interfaces ams1]
user@host# set load-balancing-options high-availability-options many-to-one preferred-backup mams-1/2/0
```

5.

NOTE: This step is not applicable to the Next Gen Services MX-SPC3 services card in the MX240, MX480 or MX960 chassis.

If the AMS interface has more than 24 member interfaces, set the service PIC boot timeout value to 240 or 300 seconds for every services PIC on the MX Series router. We recommend that you use a value of 240.

NOTE: Starting with Junos OS Release 16.2, an AMS interface can have up to 36 member interfaces. In Junos OS Release 16.1 and earlier, an AMS interface could have a maximum of 24 member interfaces.

```
[edit interfaces interface-name multiservice-options]
user@host# set pic-boot-timeout (240 | 300);
```

For example:

```
[edit interfaces sp-1/1/0 multiservice-options]
user@host# set pic-boot-timeout 240
```

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, an MX-SPC3 Next Gen Services AMS interface can have up to 16 member interfaces with a maximum of 8 MX-SPC3 services cards with up to 2 PICs on each card.
16.2	Starting with Junos OS Release 16.2, an MS-MPC AMS interface can have up to 36 member interfaces.

RELATED DOCUMENTATION

| [Understanding Aggregated Multiservices Interfaces for Next Gen Services](#) | 353

Configuring Load Balancing on AMS Infrastructure

Configuring load balancing requires an aggregated multiservices (AMS) system. AMS involves grouping several services PICs together. An AMS configuration eliminates the need for separate routers within a system. The primary benefit of having an AMS configuration is the ability to support load balancing of traffic across multiple services PICs.

AMS is supported on the MS-MPC and MS-MIC. Starting in Junos OS Release 19.3R2, AMS interfaces are also supported on the MX-SPC3 if you are running Next Gen Services.

High availability (HA) is supported on AMS infrastructure on all MX Series 5G Universal Routing Platforms. AMS has several benefits:

- Support for configuring behavior if a services PIC that is part of the AMS configuration fails
- Support for specifying hash keys for each service set in either direction
- Support for adding routes to individual PICs within the AMS system

Configuring AMS Infrastructure

AMS supports load balancing across multiple service sets. All ingress or egress traffic for a service set can be load balanced across different services PICs. To enable load balancing, you have to configure an aggregate interface with existing services interfaces.

To configure failure behavior in AMS, include the **member-failure-options** statement:

```
[edit interfaces ams1]
load-balancing-options {
  member-failure-options {
    drop-member-traffic {
      rejoin-timeout rejoin-timeout;
    }
    redistribute-all-traffic {
      enable-rejoin;
    }
  }
}
```

If a PIC fails, you can configure the traffic to the failed PIC to be redistributed by using the **redistribute-all-traffic** statement at the **[edit interfaces *interface-name* load-balancing-options member-failure-options]** hierarchy level. If the **drop-member-traffic** statement is used, all traffic to the failed PIC is dropped. Both options are mutually exclusive.

NOTE: If **member-failure-options** is not explicitly configured, the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Only **mams-** interfaces (services interfaces that are part of AMS) can be aggregated. After an AMS interface has been configured, you cannot configure the individual constituent **mams-** interfaces. A **mams-** interface cannot be used as an **ams** interface (this is not applicable to Next Gen Services MX-SPC3). AMS supports IPv4 (**family inet**) and IPv6 (**family inet6**). You cannot configure addresses on an AMS interface. Network Address Translation (NAT) is the only application that runs on AMS infrastructure at this time.

NOTE: You cannot configure unit 0 on an AMS interface.

To support multiple applications and different types of translation, AMS infrastructure supports configuring hashing for each service set. You can configure the hash keys separately for ingress and egress. The default configuration uses source IP, destination IP, and the protocol for hashing; incoming-interface for ingress and outgoing-interface for egress are also available.

NOTE: When using AMS in a load-balanced setup for the NAT solution, the number of NAT IP addresses must be greater than or equal to the number of active mams-interfaces you have added to the AMS bundle.

Configuring High Availability

In an AMS system configured with high availability, a designated services PIC acts as a backup for other active PICs that are part of the AMS system in a many-to-one (N:1) backup configuration. In a N:1 backup configuration, one PIC is available as backup for all other active PICs. If any of the active PICs fail, the backup PIC takes over for the failed PIC. In an N:1 (stateless) backup configuration, traffic states and data structures are not synchronized between the active PICs and the backup PIC.

An AMS system also supports a one-to-one (1:1) configuration. In the case of 1:1 backup, a backup interface is paired with a single active interface. If the active interface fails, the backup interface takes over. In a 1:1 (stateful) configuration, traffic states and data structures are synchronized between the active PICs and the backup PIC. Stateful synchronization is required for high availability of IPsec connections. For IPsec connections, AMS supports 1:1 configuration only.

NOTE: IPsec connections are not supported on the MX-SPC3 in this release.

High availability for load balancing is configured by adding the **high-availability-options** statement at the **[edit interfaces *interface-name* load-balancing-options]** hierarchy level.

To configure N:1 high availability, include the **high-availability-options** statement with the **many-to-one** option:

```
[edit interfaces ams1]
load-balancing-options {
  high-availability-options {
```

```

many-to-one {
    preferred-backup preferred-backup;
}
}

```

Starting in Junos OS Release 16.1, you can configure stateful 1:1 high availability on an MS-MPC. To configure stateful 1:1 high availability, at the **[edit interfaces *interface-name* load-balancing-options]** hierarchy level, include the **high-availability-options** statement with the **one-to-one** option:

NOTE: The Next Gen Services MX-SPC3 services card does not support AMS 1:1 high availability.

```

[edit interfaces ams1]
load-balancing-options {
    high-availability-options {
        one-to-one {
            preferred-backup preferred-backup;
        }
    }
}

```

Load Balancing Network Address Translation Flows

Network Address Translation (NAT) has been programmed as a plug-in and is a function of load balancing and high availability. The plug-in runs on AMS infrastructure. All flows for translation are automatically distributed to different services PICs that are part of the AMS infrastructure. In case of failure of an active services PIC, the configured backup PIC takes over the NAT pool resources of the failed PIC. The hashing method selected depends on the type of NAT. Using NAT on AMS infrastructure has a few limitations:

- NAT flows to failed PICs cannot be restored.
- There is no support for IPv6 flows.

IPv6 address pools are not supported with AMS, however NAT64 is supported with AMS, so that IPv6 flows enter AMS.

NAT64 is supported for Next Gen Services on the MX-SPC3 services card, there is no support of NAT66. IPv6 flows for different NAT services are supported except where the translation is required to be IPv6 to IPv6 or IPv4 to IPv6.

- Twice NAT is not supported for load balancing on MS-MPC cards.

Twice NAT is supported for load balancing on the Next Gen Services MX-SPC3 services card.

- Deterministic NAT uses warm-standby AMS configuration and can distribute the load using multiple AMS bundles in warm-standby mode.

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, AMS interfaces are also supported on the MX-SPC3 if you are running Next Gen Services.
16.1	Starting in Junos OS Release 16.1, you can configure stateful 1:1 high availability on an MS-MPC.

Configuring Warm Standby for Services Interfaces

You can configure an N:1 warm standby option for MS-MPCs, MS-MICs, and MX-SPC3s by creating multiple aggregated multiservices (AMS) interfaces, each of which contains the service interface you want to backup and the service interface that acts as the backup. The same backup service interface can be used in all these AMS interfaces. Starting in Junos OS Release 19.3R2, the N:1 warm standby option is also supported on the MX-SPC3 if you are running Next Gen Services.

To configure warm standby for services interfaces:

1. Create an AMS interface.

```
[edit interfaces]
user@host# set amsN
```

The variable *N* is a unique number, such as 0 or 1.

2. Specify the primary service interface that you want to backup.

```
[edit interfaces amsN]
user@host# set redundancy-options primary mams-a/b/0
```

The variable *a* is the FPC slot number and *b* is the PIC slot number for the primary service interface.

3. Specify the secondary service interface, which backs up the primary interface.

```
[edit interfaces amsN]
user@host# set redundancy-options secondary mams-a/b/0
```

The variable *a* is the FPC slot number and *b* is the PIC slot number for the secondary service interface.

- 4. Repeat Steps 1 through 3 to create an AMS interface for each service interface that you want to backup. You can use the same secondary service interface in each AMS interface.

Release History Table

Release	Description
19.3R2	Starting in Junos OS Release 19.3R2, the N:1 warm standby option is also supported on the MX-SPC3 if you are running Next Gen Services.

RELATED DOCUMENTATION

| *Understanding Aggregated Multiservices Interfaces*

10

PART

Inter-Chassis Services PIC High Availability

[Inter-Chassis Services PIC High Availability Overview and Configuration](#) | 368

Inter-Chassis Services PIC High Availability Overview and Configuration

IN THIS CHAPTER

- [Next Gen Services Inter-chassis High Availability Overview for NAT, Stateful Firewall, and IDS Flows | 368](#)
- [Inter-Chassis Stateful Synchronization for Long Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 381](#)
- [Inter-Chassis Services Redundancy Overview for Next Gen Services | 388](#)
- [Configuring Inter-Chassis Services Redundancy for Next Gen Services | 391](#)

Next Gen Services Inter-chassis High Availability Overview for NAT, Stateful Firewall, and IDS Flows

IN THIS SECTION

- [Inter-chassis High Availability Overview for NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 368](#)
- [Example: Next Gen Services Inter-Chassis Stateful High Availability for NAT and Stateful Firewall \(MX-SPC3\) | 369](#)

Inter-chassis High Availability Overview for NAT, Stateful Firewall, and IDS Flows for Next Gen Services

Carrier-grade NAT, stateful firewall, and IDS flows can be configured with a dual-chassis, redundant data path. Although intra-chassis high availability can be used in an MX Series device by employing the AMS interfaces, this method only deals locally with services PIC failures. If for any reason traffic is switched to a backup router due to some other failure in the router, the session state from the services PIC is lost

unless you configure synchronization of the services session states with a services PIC on the backup router.

Inter-chassis high availability provides this synchronization, and controls switchovers between the services PICs in the redundancy pair. Inter-chassis high availability is a primary-secondary model, not an active-active cluster. Only one services PIC in a redundancy pair, the current master, receives traffic to be serviced.

To configure interchassis high availability for NAT, stateful firewall, and IDS, you configure:

1. Stateful synchronization, which replicates the session state from the master services PICs on the master to the backup services PIC on the other chassis.
2. Inter-chassis services redundancy, which controls mastership switchovers in the services PIC redundancy pair, based on monitored events. Most operators would not want to employ stateful synchronization without also implementing services redundancy.

Benefits

Interchassis high availability provides automatic switchovers from a services PIC on one chassis to a services PIC on another chassis, while providing uninterrupted services for customer traffic.

Example: Next Gen Services Inter-Chassis Stateful High Availability for NAT and Stateful Firewall (MX-SPC3)

IN THIS SECTION

- [Requirements | 369](#)
- [Overview | 369](#)
- [Configuration | 370](#)

This example shows how to configure Next Gen Services inter-chassis high availability for stateful firewall and NAT services.

Requirements

This example uses the following hardware and software components:

- Two MX480 routers with MX-SPC3 services cards
- Junos OS Release 19.3R2, 19.4R1 or later

Overview

Two MX 3D routers are identically configured to facilitate stateful failover for firewall and NAT services in case of a chassis failure.

Configuration

IN THIS SECTION

- [Configuring Interfaces for Chassis 1. | 372](#)
- [Configure Routing Information for Chassis 1 | 374](#)
- [Configuring NAT and Stateful Firewall for Chassis 1 | 375](#)
- [Configuring the Service Set | 377](#)
- [Configuring Interfaces for Chassis 2 | 378](#)
- [Configure Routing Information for Chassis 2 | 380](#)

To configure inter-chassis high availability for this example, perform these tasks:

CLI Quick Configuration

To quickly configure this example on the routers, copy the following commands and paste them into the router terminal window after removing line breaks and substituting interface information specific to your site.

NOTE: The following configuration is for chassis 1.

```
[edit]
set interfaces vms-4/0/0 redundancy-options redundancy-peer ipaddress 5.5.5.2
set interfaces vms-4/0/0 redundancy-options routing-instance HA
set interfaces vms-4/0/0 unit 10 ip-address-owner service-plane
set interfaces vms-4/0/0 unit 10 family inet address 5.5.5.1/32
set interfaces vms-4/0/0 unit 20 family inet
set interfaces vms-4/0/0 unit 20 service-domain inside
set interfaces vms-4/0/0 unit 30 family inet
set interfaces vms-4/0/0 unit 30 service-domain outside
set interfaces ge-2/0/0 vlan-tagging
set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.1/24
set routing-instances HA instance-type vrf
set routing-instances HA interface ge-2/0/0.0
set routing-instances HA interface vms-4/0/0.10
set routing-instances HA route-distinguisher 1:1
set policy-options policy-statement dummy term 1 then reject
set routing-instances HA vrf-import dummy
```

```

set routing-instances HA vrf-export dummy
set routing-instances HA routing-options static route route 5.5.5.1/32 next-hop vms-4/0/0.10
set routing-instances HA routing-options static route route 5.5.5.2/32 next-hop 20.1.1.2
set services nat pool p2 address 32.0.0.0/24
set services nat pool p2 port automatic random-allocation
set services nat pool p2 address-allocation round-robin
set services nat rule r2 match-direction input
set services nat rule r2 term t1 from source-address 129.0.0.0/8
set services nat rule r2 term t1 from source-address 128.0.0.0/8
set services nat rule r2 term t1 then translated source-pool p2
set services nat rule r2 term t1 then translated translation-type napt-44
set services nat rule r2 term t1 then translated address-pooling paired
set services nat rule r2 term t1 then syslog
set services stateful-firewall rule r2 match-direction input
set services stateful-firewall rule r2 term t1 from source-address any-unicast
set services stateful-firewall rule r2 term t1 then accept
set services stateful-firewall rule r2 term t1 then syslog
set services service-set ss2 replicate-services replication-threshold 180
set services service-set ss2 replicate-services stateful-firewall
set services service-set ss2 replicate-services nat
set services service-set ss2 stateful-firewall-rules r2
set services service-set ss2 nat-rules r2
set services service-set ss2 next-hop-service inside-service-interface vms-4/0/0.20
set services service-set ss2 next-hop-service outside-service-interface vms-4/0/0.30
set services service-set ss2 syslog host local class session-logs
set services service-set ss2 syslog host local class stateful-firewall-logs
set services service-set ss2 syslog host local class nat-logs

```

NOTE: The following configuration is for chassis 2. The NAT, stateful firewall, and service-set information must be identical for chassis 1 and 2.

```

set interfaces vms-4/0/0 redundancy-options routing-instance HA
set interfaces vms-4/0/0 unit 10 ip-address-owner service-plane
set interfaces vms-4/0/0 unit 10 family inet address 5.5.5.2/32
set interfaces vms-4/0/0 unit 20 family inet
set interfaces vms-4/0/0 unit 20 service-domain inside
set interfaces vms-4/0/0 unit 30 family inet
set interfaces vms-4/0/0 unit 30 service-domain outside
set interfaces ge-2/0/0 vlan-tagging
set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.2/24
set routing-instances HA instance-type vrf

```

```

set routing-instances HA interface ge-2/0/0.0
set routing-instances HA interface vms-4/0/0.10
set routing-instances HA route-distinguisher 1:1
set routing-instances HA vrf-import dummy
set routing-instances HA vrf-export dummy
set routing-instances HA routing-options static route 5.5.5.2/32 next-hop vms-4/0/0.10
set routing-instances HA routing-options static route 5.5.5.1/32 next-hop 20.1.1.1
set services nat pool p2 address 32.0.0.0/24
set services nat pool p2 port automatic random-allocation
set services nat pool p2 address-allocation round-robin
set services nat rule r2 match-direction input
set services nat rule r2 term t1 from source-address 129.0.0.0/8
set services nat rule r2 term t1 from source-address 128.0.0.0/8
set services nat rule r2 term t1 then translated source-pool p2
set services nat rule r2 term t1 then translated translation-type napt-44
set services nat rule r2 term t1 then translated address-pooling paired
set services nat rule r2 term t1 then syslog
set services stateful-firewall rule r2 match-direction input
set services stateful-firewall rule r2 term t1 from source-address any-unicast
set services stateful-firewall rule r2 term t1 then accept
set services stateful-firewall rule r2 term t1 then syslog
set services service-set ss2 replicate-services replication-threshold 180
set services service-set ss2 replicate-services stateful-firewall
set services service-set ss2 replicate-services nat
set services service-set ss2 stateful-firewall-rules r2
set services service-set ss2 nat-rules r2
set services service-set ss2 next-hop-service inside-service-interface vms-4/0/0.20
set services service-set ss2 next-hop-service outside-service-interface vms-4/0/0.30
set services service-set ss2 syslog host local class session-logs
set services service-set ss2 syslog host local class stateful-firewall-logs
set services service-set ss2 syslog host local class nat-logs

```

Configuring Interfaces for Chassis 1.

Step-by-Step Procedure

The interfaces for each of the HA pair of routers are configured identically with the exception of the following service PIC options:

- **redundancy-options redundancy-peer ipaddress *address***
- **unit *unit-number* family inet address *address*** of a unit, other than 0, that contains the **ip-address-owner service-plane** option

To configure interfaces:

1. Configure the redundant service PIC on chassis 1.

```
[edit interfaces]
user@host# set interfaces vms-4/0/0 redundancy-options redundancy-peer ipaddress 5.5.5.2
user@host# set interfaces vms-4/0/0 redundancy-options routing-instance HA
user@host# set interfaces vms-4/0/0 unit 10 ip-address-owner service-plane
user@host# set interfaces vms-4/0/0 unit 10 family inet address 5.5.5.1/32
user@host# set interfaces vms-4/0/0 unit 20 family inet
user@host# set interfaces vms-4/0/0 unit 20 service-domain inside
user@host# set interfaces vms-4/0/0 unit 30 family inet
user@host# set interfaces vms-4/0/0 unit 30 service-domain outside
```

2. Configure the interfaces for chassis 1 that are used as interchassis links for synchronization traffic.

```
user@host# set interfaces ge-2/0/0 vlan-tagging
user@host# set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.1/24
```

3. Configure remaining interfaces as needed.

Results

```
user@host# show interfaces
```

```
ge-2/0/0 {
    vlan-tagging;
    unit 0 {
        vlan-id 100;
        family inet {
            address 20.1.1.1/24;
        }
    }
}

vms-4/0/0 {
    redundancy-options {
        redundancy-peer {
            ipaddress 5.5.5.2;
        }
        routing-instance HA;
    }
    unit 10 {
        ip-address-owner service-plane;
        family inet {
            address 5.5.5.1/32;
        }
    }
}
```

```

    }
    unit 20 {
        family inet;
        family inet6;
        service-domain inside;
    }
    unit 30 {
        family inet;
        family inet6;
        service-domain outside;
    }
}

```

Configure Routing Information for Chassis 1

Step-by-Step Procedure

Detailed routing configuration is not included for this example. A routing instance is required for the HA synchronization traffic between the chassis as follows:

- Configure routing instances for Chassis 1.

```

user@host# set routing-instances HA instance-type vrf
user@host# set routing-instances HA interface ge-2/0/0.0
user@host# set routing-instances HA interface vms-4/0/0.10
user@host# set routing-instances HA route-distinguisher 1:1
user@host# set policy-options policy-statement dummy term 1 then reject
user@host# set routing-instances HA vrf-import dummy
user@host# set routing-instances HA vrf-export dummy
user@host# set routing-instances HA routing-options static route route 5.5.5.1/32 next-hop vms-4/0/0.10
user@host# set routing-instances HA routing-options static route route 5.5.5.2/32 next-hop 20.1.1.2

```

Results

```
user@host# show routing-instances
```

```

HA {
    instance-type vrf;
    interface ge-2/0/0.0;
    interface vms-4/0/0.10;
    route-distinguisher 1:1;
    vrf-import dummy;
}

```



```

vrf-export dummy;
routing-options {
    static {
        route 5.5.5.1/32 next-hop vms-4/0/0.10;
        route 5.5.5.2/32 next-hop 20.1.1.2;
    }
}

```

Configuring NAT and Stateful Firewall for Chassis 1

Step-by-Step Procedure

Configure NAT and stateful firewall identically on both routers. To configure NAT and stateful firewall:

1. Configure NAT as needed.

```

user@host# set services nat pool p2 address 32.0.0.0/24
user@host# set services nat pool p2 port automatic random-allocation
user@host# set services nat pool p2 address-allocation round-robin
user@host# set services nat rule r2 match-direction input
user@host# set services nat rule r2 term t1 from source-address 129.0.0.0/8
user@host# set services nat rule r2 term t1 from source-address 128.0.0.0/8
user@host# set services nat rule r2 term t1 then translated source-pool p2
user@host# set services nat rule r2 term t1 then translated translation-type napt-44
user@host# set services nat rule r2 term t1 then translated address-pooling paired
user@host# set services nat rule r2 term t1 then syslog

```

2. Configure stateful firewall as needed.

```

user@host# set services stateful-firewall rule r2 match-direction input
user@host# set services stateful-firewall rule r2 term t1 from source-address any-unicast
user@host# set services stateful-firewall rule r2 term t1 then accept
user@host# set services stateful-firewall rule r2 term t1 then syslog

```

Results

user@host# **show services nat**

```

nat {
    pool p2 {
        address 32.0.0.0/24;
    }
}

```

```

        port {
            automatic {
                random-allocation;
            }
        }
        address-allocation round-robin;
    }
    rule r2 {
        match-direction input;
        term t1 {
            from {
                source-address {
                    129.0.0.0/8;
                    128.0.0.0/8;
                }
            }
            then {
                translated {
                    source-pool p2;
                    translation-type {
                        napt-44;
                    }
                    address-pooling paired;
                }
                syslog;
            }
        }
    }
}

```

user@host **show services stateful-firewall**

```

rule r2 {
    match-direction input;
    term t1 {
        from {
            source-address {
                any-unicast;
            }
        }
        then {
            accept;
        }
    }
}

```

```

        syslog;
    }
}

```

Configuring the Service Set

Step-by-Step Procedure

Configure the the service set identically on both routers. To configure the service set:

1. Configure the service set replication options.

```

user@host# set services service-set ss2 replicate-services replication-threshold 180
user@host# set services service-set ss2 replicate-services stateful-firewall
user@host# set services service-set ss2 replicate-services nat

```

2. Configure references to NAT and stateful firewall rules for the service set.

```

user@host# set services service-set ss2 stateful-firewall-rules r2
user@host# set services service-set ss2 nat-rules r2

```

3. Configure next-hop service interface on the vms-PIC.

```

user@host# set services service-set ss2 next-hop-service inside-service-interface vms-4/0/0.20
user@host# set services service-set ss2 next-hop-service outside-service-interface vms-4/0/0.30

```

4. Configure desired logging options.

```

user@host# set services service-set ss2 syslog host local class session-logs
user@host# set services service-set ss2 syslog host local class stateful-firewall-logs
user@host# set services service-set ss2 syslog host local class nat-logs

```

Results

```

user@host# show services service-set ss2

```

```

syslog {
    host local {

```

```

        class {
            session-logs;
            inactive: stateful-firewall-logs;
            nat-logs;
        }
    }
}
replicate-services {
    replication-threshold 180;
    stateful-firewall;
    nat;
}
stateful-firewall-rules r2;
inactive: nat-rules r2;
next-hop-service {
    inside-service-interface vms-3/0/0.20;
    outside-service-interface vms-3/0/0.30;
}
}

```

Configuring Interfaces for Chassis 2

Step-by-Step Procedure

The interfaces for each of the HA pair of routers are configured identically with the exception of the following service PIC options:

- **redundancy-options redundancy-peer ipaddress address**
- **unit unit-number family inet address address** of a unit, other than 0, that contains the **ip-address-owner service-plane** option

1. Configure the redundant service PIC on chassis 2.

The **redundancy-peer ipaddress** points to the address of the unit (unit 10) on vms-4/0/0 on chassis on chassis 1 that contains the **ip-address-owner service-plane** statement.

```

[edit interfaces]
set interfaces vms-4/0/0 redundancy-options redundancy-peer ipaddress 5.5.5.1
user@host# set interfaces vms-4/0/0 redundancy-options routing-instance HA
user@host# set interfaces vms-4/0/0 unit 10 ip-address-owner service-plane
user@host# set interfaces vms-4/0/0 unit 10 family inet address 5.5.5.2/32
user@host# set interfaces vms-4/0/0 unit 20 family inet
user@host# set interfaces vms-4/0/0 unit 20 service-domain inside
user@host# set interfaces vms-4/0/0 unit 30 family inet

```

```
user@host# set interfaces vms-4/0/0 unit 30 service-domain outside
```

2. Configure the interfaces for chassis 2 that are used as interchassis links for synchronization traffic

```
user@host# set interfaces ge-2/0/0 vlan-tagging
user@host# set interfaces ge-2/0/0 unit 0 vlan-id 100 family inet address 20.1.1.2/24
```

3. Configure remaining interfaces for chassis 2 as needed.

Results

```
user@host# show interfaces
```

```
vms-4/0/0 {
  redundancy-options {
    redundancy-peer {
      ipaddress 5.5.5.1;
    }
    routing-instance HA;
  }
  unit 0 {
    family inet;
  }
  unit 10 {
    ip-address-owner service-plane;
    family inet {
      address 5.5.5.2/32;
    }
  }
}
ge-2/0/0 {
  vlan-tagging;
  unit 0 {
    vlan-id 100;
    family inet {
      address 20.1.1.2/24;
    }
  }
  unit 10 {
    vlan-id 10;
    family inet {
      address 2.10.1.2/24;
    }
  }
}
```

Configure Routing Information for Chassis 2

Step-by-Step Procedure

Detailed routing configuration is not included for this example. A routing instance is required for the HA synchronization traffic between the two chassis and is included here.

- Configure routing instances for chassis 2.

```
user@host# set routing-instances HA instance-type vrf
user@host# set routing-instances HA interface ge-2/0/0.0
user@host# set routing-instances HA interface vms-4/0/0.10
user@host# set routing-instances HA route-distinguisher 1:1
user@host# set policy-options policy-statement dummy term 1 then reject
user@host# set routing-instances HA vrf-import dummy
user@host# set routing-instances HA vrf-export dummy
user@host# set routing-instances HA routing-options static route 5.5.5.2/32 next-hop vms-4/0/0.10
user@host# set routing-instances HA routing-options static route 5.5.5.1/32 next-hop 20.1.1.1
```

NOTE: The following configuration steps are *identical* to the steps shown for chassis 1.

- Configuring NAT and Stateful Firewall
- Configuring the Service Set

Results

user@host# **show services routing-instances**

```
HA {
    instance-type vrf;
    interface xe-2/2/0.0;
    interface vms-4/0/0.10;
    route-distinguisher 1:1;
    vrf-import dummy;
    vrf-export dummy;
    routing-options {
        static {
            route 5.5.5.2/32 next-hop vms-4/0/0.10;
            route 5.5.5.1/32 next-hop 20.1.1.1;
        }
    }
}
```

RELATED DOCUMENTATION

[Inter-Chassis Stateful Synchronization for Long Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 381](#)

[Inter-Chassis Services Redundancy Overview for Next Gen Services | 388](#)

Inter-Chassis Stateful Synchronization for Long Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services

IN THIS SECTION

- [Inter-Chassis Stateful Synchronization Overview | 381](#)
- [Configuring Inter-Chassis Stateful Synchronization for Long- Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 382](#)

Inter-Chassis Stateful Synchronization Overview

Stateful synchronization replicates the state of long-lived NAT, stateful firewall, and IDS sessions on the master services PIC and sends it to the backup services PIC, which is on a different MX Series chassis. By default, long lived sessions are defined as having been active on the services PIC for at least 180 seconds, though you can configure this to a higher value.

The following restrictions apply:

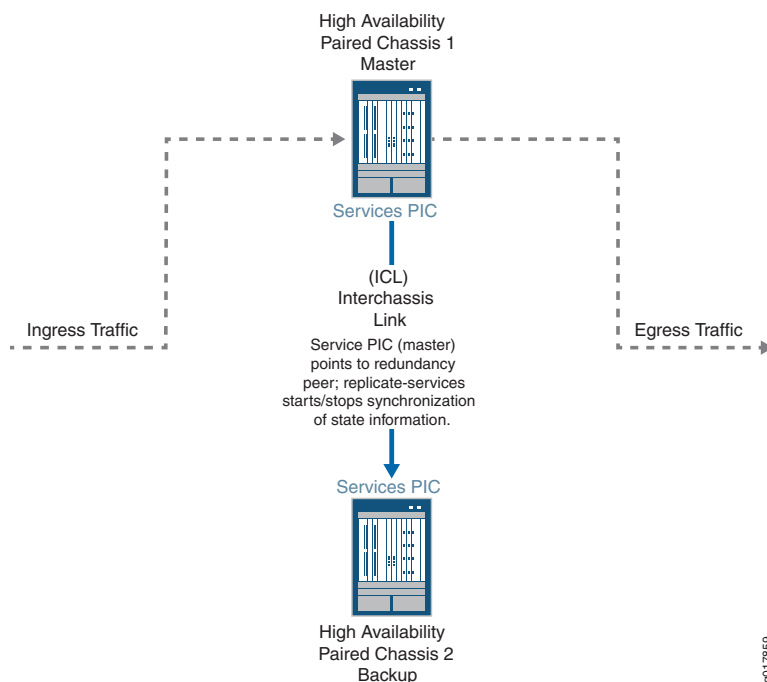
- NAPT44 is the only translation type supported.

Replicating state information for the port block allocation (PBA), endpoint-independent mapping (EIM), or endpoint-independent filters (EIF) features are supported for Next Gen Services.

When configuring a service set for NAT, stateful firewall, or IDS that belongs to a stateful synchronization setup, you must use a next-hop service set, and the NAT, stateful firewall, and IDS configurations for the service set must be identical on both MX Series chassis.

[Figure 12 on page 382](#) shows the stateful synchronization topology.

Figure 12: Stateful Sync Topology



Benefits

Interchassis stateful synchronization of the services session state allows uninterrupted services when a switchover occurs from a services PIC on one chassis to a services PIC on another chassis.

SEE ALSO

[Configuring Inter-Chassis Stateful Synchronization for Long- Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services | 382](#)

Configuring Inter-Chassis Stateful Synchronization for Long- Lived NAT, Stateful Firewall, and IDS Flows for Next Gen Services

IN THIS SECTION

- [Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with non-AMS Interface | 383](#)
- [Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with AMS Interface | 385](#)

Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with non-AMS Interface

To configure stateful synchronization inter-chassis high availability for NAT, stateful firewall, and IDS flows for Next Gen Services when the services interfaces are not AMS, perform the following configuration steps on each chassis of the high availability pair.

1. Specify the IP address of the vms- interface. This address is used by the TCP channel between the HA pairs.

```
[edit interfaces interface-name redundancy-options]
user@host# set redundancy-local data-address address
```

For example:

```
[edit interfaces vms-1/0/0 redundancy-options]
user@host# set redundancy-local data-address 192.0.2.2
```

When you configure the other chassis, this is the address you use for the **redundancy-peer ipaddress**.

2. Specify the IP address of the remote services interface. This address is used by the TCP channel between the HA pairs.

```
[edit interfaces interface-name redundancy-options]
user@host# set redundancy-peer ipaddress address
```

For example:

```
[edit interfaces vms-1/0/0 redundancy-options]
user@host# set redundancy-peer ipaddress 192.0.2.1
```

When you configure the other chassis, this is the address you use for the **redundancy-local data-address**.

3. Configure the length of time that the flow remains active for replication, in seconds.

```
[edit interfaces interface-name redundancy-options]
user@host# set replication-threshold seconds
```

For example:

```
[edit interfaces vms-1/0/0 redundancy-options]
user@host# set replication-threshold 60
```

4. Configure a unit other than 0, and assign it the IP address of the local services interface that you configured with the **redundancy-local data-address** option.

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

For example:

```
[edit interfaces vms-1/0/0]
user@host# set unit 10 family inet address 192.0.2.2/32
```

5. For ease of management, we recommend you create a special routing instance with **instance-type vrf** to host the HA synchronization traffic between the MX Series high availability pair. Then specify the name of the special routing instance to apply to the HA synchronization traffic between the high availability pair.

```
[edit interfaces interface-name redundancy-options]
user@host# set routing-instance instance-name
```

6. Configure the inside and outside interface units, which are used by the next-hop service set. Use different unit numbers for the inside and outside units, and do not use 0 or the unit number used in Step 4.

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number family (inet | inet6)
user@host# set interfaces interface-name unit logical-unit-number service-domain inside
user@host# set interfaces interface-name unit logical-unit-number family (inet | inet6)
user@host# set interfaces interface-name unit logical-unit-number service-domain outside
```

For example:

```
[edit]
user@host# set interfaces vms-1/0/0 unit 100 family inet
user@host# set interfaces vms-1/0/0 unit 100 family inet6
user@host# set interfaces vms-1/0/0 unit 100 service-domain inside
user@host# set interfaces vms-1/0/0 unit 1000 family inet
user@host# set interfaces vms-1/0/0 unit 1000 family inet6
user@host# set interfaces vms-1/0/0 unit 1000 service-domain outside
```

7. Configure the next-hop service set that contains the NAT rules, stateful firewall rules, or IDS screens. The service set must be configured identically on each chassis of the high availability pair. The NAT rules, stateful firewall rules, and IDS screens must also be configured identically on each chassis.

For example:

```
user@host#set service-set internal-nat next-hop-service inside-service-interface vms-1/0/0.100
user@host#set service-set internal-nat next-hop-service outside-service-interface vms-1/0/0.1000
user@host#set service-set internal-nat next-hop-service nat-rules internal-nat1
```

8. Repeat these steps for the other chassis of the high availability pair.

SEE ALSO

[Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with AMS Interface](#) | 385

Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with AMS Interface

To configure stateful synchronization inter-chassis high availability for NAT, stateful firewall, and IDS flows for Next Gen Services for an AMS services interface, perform the following configuration steps on each chassis of the high availability pair.

1. Configure a services vms- interface for every member of the AMS interface:
 - a. Specify the IP address of the vms- interface. This address is used by the TCP channel between the HA pairs.

```
[edit interfaces interface-name redundancy-options]
user@host# set redundancy-local data-address address
```

For example:

```
[edit interfaces vms-1/0/0 redundancy-options]
user@host# set redundancy-local data-address 192.0.2.2
```

When you configure the other chassis, this is the address you use for the **redundancy-peer ipaddress**.

- b. Specify the IP address of the remote services interface. This address is used by the TCP channel between the HA pairs.

```
[edit interfaces interface-name redundancy-options]
user@host# set redundancy-peer ipaddress address
```

For example:

```
[edit interfaces vms-1/0/0 redundancy-options]
user@host# set redundancy-peer ipaddress 192.0.2.1
```

When you configure the other chassis, this is the address you use for the **redundancy-local data-address**.

- c. Configure the length of time that the flow remains active for replication, in seconds.

```
[edit interfaces interface-name redundancy-options]
user@host# set replication-threshold seconds
```

For example:

```
[edit interfaces vms-1/0/0 redundancy-options]
user@host# set replication-threshold 60
```

- d. Configure a unit other than 0, and assign it the IP address of the local services interface that you configured with the **redundancy-local data-address** option.

```
[edit interfaces interface-name]
user@host# set unit logical-unit-number family inet address address
```

For example:

```
[edit interfaces vms-1/0/0]
user@host# set unit 10 family inet address 192.0.2.2/32
```

- e. For ease of management, we recommend you create a special routing instance with **instance-type vrf** to host the HA synchronization traffic between the MX Series high availability pair. Then specify the name of the special routing instance to apply to the HA synchronization traffic between the high availability pair.

```
[edit interfaces interface-name redundancy-options]
user@host# set routing-instance instance-name
```

2. Create the AMS interface and add the member interfaces you configured in Step [1](#).

```
[edit interfaces]
```

```
user@host# set interface-name load-balancing-options [member-interface mams-a/b/0]
```

where the *interface-name* is *amsN*, and *a* is the FPC slot number and *b* is the PIC slot number for each member interface.

For example:

```
[edit interfaces]
user@host# set ams0 load-balancing-options member-interface mams-1/0/0
user@host# set ams0 load-balancing-options member-interface mams-1/1/0
```

3. Configure the inside interface for the AMS interface, which is used by the next-hop service set:
 - a. Configure the family for the inside interface. Do not use 0 for the unit number.

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number service-domain inside
user@host# set interfaces interface-name unit logical-unit-number family (inet | inet6)
```

For example:

```
[edit]
user@host# set interfaces ams0 unit 100 service-domain inside
user@host# set interfaces ams0 unit 100 family inet
user@host# set interfaces ams0 unit 100 family inet6
```

- b. Configure the hash key to regulate distribution for the inside interface.

```
[edit set interfaces interface-name unit logical-unit-number]
user@host# load-balancing-options hash-keys ingress-key [source-ip destination-ip]
```

4. Configure the outside interface for the AMS interface, which is used by the next-hop service set. Do not use 0 or the same unit number that you used for the inside interface.
 - a. Configure the family for the outside interface.

```
[edit]
user@host# set interfaces interface-name unit logical-unit-number service-domain outside
user@host# set interfaces interface-name unit logical-unit-number family (inet | inet6)
```

For example:

```
[edit]
user@host# set interfaces ams0 unit 1000 service-domain outside
user@host# set interfaces ams0 unit 1000 family inet
user@host# set interfaces ams0 unit 1000 family inet6
```

- b. Configure the hash key to regulate distribution for the outside interface.

```
[edit set interfaces interface-name unit logical-unit-number]
user@host# load-balancing-options hash-keys ingress-key [source-ip destination-ip]
```

5. Configure the next-hop service set that contains the NAT rules, stateful firewall rules, or IDS screens. The service set must be configured identically on each chassis of the high availability pair. The NAT rules, stateful firewall rule, and IDS screens must also be configured identically on each chassis.

For example:

```
user@host#set service-set internal-nat next-hop-service inside-service-interface ams0.100
user@host#set service-set internal-nat next-hop-service outside-service-interface ams0.1000
user@host#set service-set internal-nat next-hop-service nat-rules internal-nat1
```

6. Repeat these steps for the other chassis of the high availability pair.

SEE ALSO

[Configuring Inter-Chassis Stateful Synchronization for Next Gen Services with non-AMS Interface | 383](#)

Inter-Chassis Services Redundancy Overview for Next Gen Services

IN THIS SECTION

- [Introduction to Inter-Chassis Services Redundancy | 389](#)
- [Benefits | 389](#)
- [Services Redundancy Components | 389](#)
- [Services Redundancy Operation | 390](#)

Introduction to Inter-Chassis Services Redundancy

Interchassis redundancy for services is controlled by the services redundancy daemon (SRD). The SRD lets you specify events that trigger a switchover between the master and standby services PICs, which are on two different MX Series chassis. The SRD monitors conditions, and performs a switchover when an event occurs. Inter-chassis services redundancy is a primary-secondary model, not an active-active cluster. Only one services PIC in a redundancy pair, the current master, receives traffic to be serviced.

You can configure redundancy based on the following monitored events:

- Link down events.
- FPC and PIC reboots.
- Routing protocol daemon (rpd) aborts and restarts.
- Peer gateway events, including requests to acquire or release mastership, or to broadcast warnings.

Benefits

Inter-chassis services redundancy provides automatic switchovers from a services PIC on one chassis to a services PIC on another chassis when a monitored event occurs.

Services Redundancy Components

The following configurable components control services redundancy processing:

- **Redundancy Event**—A monitored critical event that triggers the redundancy peers to acquire or release mastership or to create a warning, and to add or delete signal routes.

One monitored interface can be part of only one redundancy event, but one redundancy event can have multiple monitored interfaces.

- **Redundancy Policy**—A policy that defines the set of actions taken when a redundancy event occurs. Available actions include acquisition or release of mastership, creation of a warning, and addition or deletion of signal routes. You can configure a maximum of 256 redundancy policies. A redundancy policy can have a maximum of 256 interface-down events.

One redundancy event can be part of only one redundancy policy, but one redundancy policy can have multiple redundancy events. For example, redundancy policy RP1 can include redundancy events RE1 and RE2. Redundancy events RE1 and RE2 cannot be included in redundancy policies other than RP1.

- **Redundancy Set**—A collection of one or more redundancy policies that is assigned to one or more service sets on each MX Series chassis of the redundant pair, and the redundancy group that is associated with the redundancy set. At a given time, a particular redundancy set can be active on only one gateway, but not all redundancy sets have to be active on the same gateway. For example, redundancy set A can be

active on gateway 1 while redundancy set B is active on gateway 2. You can configure a maximum of 128 redundancy sets.

One service set can be assigned only one redundancy set, but multiple service sets can be assigned the same redundancy set.

One redundancy policy can be part of only one redundancy set, but one redundancy set can have multiple redundancy policies. For example, redundancy set RS1 can include redundancy policies RP1 and RP2. Redundancy policies RP1 and RP2 cannot be included in redundancy sets other than RS1. A redundancy set can have a maximum of 16 redundancy policies.

- **Redundancy Group**—The redundancy group identifies the associated ICCP redundancy group. A one-to-one relationship exists between a redundancy set and a redundancy group. One redundancy set can be part of only one redundancy group. You can configure a maximum of 16 redundancy groups. A maximum of 16 redundancy sets can be associated with the same redundancy group.
- **Signal routes**—Static routes that are added or deleted by services redundancy processing, based on mastership state changes.
- **Routing Policies**—Policies that advertise routes based on the existence or non-existence of signal routes.
- **VRRP (Virtual Router Redundancy Protocol) route tracking**—Tracks whether a reachable signal route exists in the routing table of the routing instance in the configuration. Based on the reachability of the tracked route, VRRP route tracking dynamically changes the priority of the VRRP group.

Services Redundancy Operation

Services redundancy operates as follows:

1. The services redundancy daemon runs on the Routing Engine. It continuously monitors configured redundancy events.
2. When a redundancy event is detected, the services redundancy daemon:
 - a. Adds or removes signal routes specified in the redundancy policy.
 - b. Switches services to the standby.
 - c. Updates stateful synchronization roles as needed.
3. Resulting route changes cause:
 - a. The routing policy connected to this route to advertise routes differently.
 - b. VRRP to change advertised priorities.

To summarize the switchover process:

1. A critical event occurs.
2. The services redundancy daemon adds or removes a signal route.

3. A routing policy advertises routes differently. VRRP changes advertised priorities.
4. Services switch over to the standby.
5. Stateful synchronization is updated accordingly.

NOTE: The order of routing priorities must match the order of services mastership.

If a redundancy policy action is release-mastership and the redundancy peer's state is wait, the mastership-release fails. If a redundancy policy action is release-mastership-force, the mastership release succeeds even if the redundancy peer's state is warned.

Similarly, if a redundancy policy action on the standby is acquire-mastership and the local state is wait, the mastership-release fails. If a redundancy policy action is acquire-mastership-force, the mastership release succeeds even if the standby state is wait.

You can also use a manual command to trigger a redundancy policy that releases or acquires mastership.

If gateway 1, the chassis that is configured with the lower IP address, is the master chassis and you deactivate the services redundancy daemon on it, a switchover to gateway 2 occurs. If gateway 2, the chassis that is configured with the higher IP address, is the master chassis and you deactivate the services redundancy daemon on it, a switchover does not occur.

RELATED DOCUMENTATION

[Configuring Inter-Chassis Services Redundancy for Next Gen Services](#) | 391

Configuring Inter-Chassis Services Redundancy for Next Gen Services

IN THIS SECTION

- [Configuring Non-Stop Services Redundancy for Next Gen Services Service Set](#) | 392
- [Configuring One-Way Services Redundancy for Next Gen Services Service Set](#) | 398

This topic describes how to configure interchassis-services redundancy for Next Gen Services. This topic contains a procedure for configuring non-stop services redundancy (automatic switchovers in both

directions) and a procedure for one-way redundancy (automatic switchovers only from the original master to the original standby).

You can also use a manual request command to release or acquire mastership:

```
request services redundancy-set redundancy-set trigger redundancy-event event-name <force>
```

The command automatically triggers the specified redundancy event. You must create a configuration that assigns the redundancy event to a redundancy policy that either releases or acquires mastership. You must also assign the redundancy policy to the redundancy set used in the command.

Configuring Non-Stop Services Redundancy for Next Gen Services Service Set

Non-stop services redundancy gives you automatic services switchovers between the MX Series routers when a critical event occurs. Automatic switchovers from gateway1 to gateway2 and from gateway2 to gateway1 take place without manual intervention.

To configure non-stop services redundancy for a service set, perform the following steps on both gateway1 and gateway2:

1. Configure one or more redundancy events to monitor the conditions that trigger a services switchover to the peer gateway.
 - a. Configure a name for the redundancy event.

```
[edit services]
user@host# set event-options redundancy-event event-name
```

For example:

```
[edit services]
user@host# set event-options redundancy-event RELS_MSHIP_CRIT_EV
```

- b. Specify any interfaces that trigger a services switchover when the interface goes down.

```
[edit services event-options redundancy-event event-name]
user@host# set monitor link-down [interface-name]
```

- c. Specify that a process routing daemon restart request triggers a services switchover.

```
[edit services event-options redundancy-event event-name]
user@host# set monitor process routing restart
```

- d. Specify that a process routing daemon abort request triggers a services switchover.

```
[edit services event-options redundancy-event event-name]
user@host# set monitor process routing abort
```

- e. Specify that a request from the peer to acquire ownership triggers a services switchover.

```
[edit services event-options redundancy-event event-name]
user@host# set monitor peer mastership-acquire
```

2. Configure a redundancy policy that releases mastership and deletes a static route when the redundancy event conditions are met.

- a. Configure a name for the policy.

```
user@host# edit policy-options redundancy-policy policy-name
```

For example:

```
user@host# edit policy-options redundancy-policy RLS_MSHIP_POL
```

- b. Specify the redundancy events that release mastership.

```
[edit policy-options redundancy-policy policy-name]
user@host# set redundancy-events [event-list]
```

For example:

```
[edit policy-options redundancy-policy RLS_MSHIP_POL]
user@host# set redundancy-events RELS_MSHIP_CRIT_EV
```

If you want to be able to run the **request services redundancy-set *redundancy-set* trigger redundancy-event *event-name* <force>** to manually release mastership, include that *event-name* in the redundancy policy. The redundancy event itself does not need to be configured, because it is triggered by the request command.

For example:

```
[edit policy-options redundancy-policy RLS_MSHIP_POL]
user@host# set redundancy-events [RELS_MSHIP_CRIT_EV RELS_MSHIP_MANUAL_EV]
```

- c. Release mastership.

```
[edit policy-options redundancy-policy policy-name]
user@host# set then release-mastership
```

- d. Delete the static route.

```
[edit policy-options redundancy-policy policy-name]
user@host# set then delete-static-route destination (receive | next-hop next-hop) routing-instance
routing-instance
```

3. Configure a redundancy event to identify when the peer gateway releases mastership.

```
[edit services]
user@host# set event-options redundancy-event event-name monitor peer release-mastership
```

For example:

```
[edit services]
user@host# set event-options redundancy-event PEER_RELS_MSHIP_EV monitor peer release-mastership
```

4. Configure a redundancy policy that acquires mastership from the peer gateway and adds a static route.

- a. Configure a name for the policy.

```
user@host# edit policy-options redundancy-policy policy-name
```

For example:

```
user@host# edit policy-options redundancy-policy ACQU_MSHIP_POL
```

- b. Specify the redundancy events that acquire mastership.

```
[edit policy-options redundancy-policy policy-name]
user@host# set redundancy-events [event-list]
```

For example:

```
[edit policy-options redundancy-policy ACQU_MSHIP_POL]
user@host# set redundancy-events PEER_RELS_MSHIP_EV
```

If you want to be able to run the **request services redundancy-set *redundancy-set* trigger redundancy-event *event-name* <force>** to manually acquire mastership, include that *event-name* in the redundancy policy. The redundancy event itself does not need to be configured, because it is triggered by the request command.

For example:

```
[edit policy-options redundancy-policy ACQU_MSHIP_POL]
user@host# set redundancy-events [PEER_RELS_MSHIP_EV ACQU_MSHIP_MANUAL_EV]
```

c. Acquire mastership.

```
[edit policy-options redundancy-policy policy-name]
user@host# set then acquire-mastership
```

d. Add a static route.

```
[edit policy-options redundancy-policy policy-name]
user@host# set then add-static-route destination (receive | next-hop next-hop) routing-instance
routing-instance
```

5. Configure the redundancy set.

a. Configure a name for the redundancy set.

```
[edit services]
user@host# set redundancy-set redundancy-set
```

For example:

```
[edit services]
user@host# set redundancy-set 1
```

b. Specify the redundancy group ID for the redundancy set.

```
[edit services redundancy-set redundancy-set]
user@host# set redundancy-group redundancy-group
```

For example:

```
[edit services redundancy-set 1]
user@host# set redundancy-group 1
```

The redundancy group ID is the same redundancy group ID configured for the ICCP daemon (iccpd) through the existing ICCP configuration hierarchy. For example,

```
iccp {
  local-ip-addr 1.1.1.1;
  peer 2.2.2.2 {
    redundancy-group-id-list 1;
    liveness-detection {
      minimum-interval 1000;
    }
  }
}
```

- c. Specify the redundancy policy that releases mastership and the redundancy policy that acquires mastership.

```
[edit services redundancy-set redundancy-set]
user@host# set redundancy-policy [redundancy-policy-list]
```

For example:

```
[edit services redundancy-set 1]
user@host# set redundancy-policy [ACQU_MSHIP_POL RLS_MSHIP_POL]
```

- d. Configure the frequency of health check probes of the redundancy set, in seconds.

```
[edit services redundancy-set redundancy-set]
user@host# set healthcheck-timer-interval healthcheck-timer-interval
```

The default is 30 seconds.

- e. Configure the maximum wait time for a help check response, in seconds.

```
[edit services redundancy-set redundancy-set]
user@host# set hold-time hold-time
```

The range is 0 through 3600 seconds.

- f. Configure the frequency of srd hello messages, in seconds.

```
[edit services redundancy-set redundancy-set]
user@host# set keepalive keepalive
```

The range is 1 through 60 seconds.

6. Configure routing policies.

- a. Identify signal routes that requires redundancy-related routing changes. Specify the signal route and the routing table that is used.

```
[edit policy-options condition condition-name]
user@host# set if-route-exists signal-route table routing-table
```

For example:

```
[edit policy-options condition switchover-route-exists]
user@host# set if-route-exists 10.45.45.0/24 table bgp1_table
```

- b. To change the local-preference for the signal route, enter it in a policy statement.

```
[edit policy-options policy-statement policy-name]
user@host# set term term from protocol [protocol variables] prefix-list prefix-list condition condition-name
then local-preference preference-value accept
```

- c. To change as-path-prepend values for the signal route, enter them in the policy statement.

```
[edit policy-options policy-statement policy-name]
user@host# set term term from prefix-list prefix-list condition condition-name then as-path-prepend
[as-prepend-values] next-hop self accept
```

7. Configure redundancy for the service set by assigning the redundancy set to the service set.

```
[edit]
user@host# set services service-set service-set-name redundancy-set-id redundancy-set
```

8. Repeat these steps on the peer gateway.

SEE ALSO

[Configuring One-Way Services Redundancy for Next Gen Services Service Set](#) | 398

Configuring One-Way Services Redundancy for Next Gen Services Service Set

One-way services redundancy gives you automatic services switchovers from gateway1, the original master gateway, to gateway2, the original standby gateway. An automatic switchover from gateway 2 to gateway1 does not happen. To switchover from gateway2 to gateway1, you must perform a manual switchover.

1. On gateway1, the initial master, configure one or more redundancy events to monitor the conditions that trigger a services switchover to gateway2, the standby gateway.
 - a. Configure a name for the redundancy event.

```
[edit services]
user@gateway1# set event-options redundancy-event event-name
```

For example:

```
[edit services]
user@gateway1# set event-options redundancy-event RELS_MSHIP_CRIT_EV
```

- b. Specify any interfaces that trigger a services switchover when the interface goes down.

```
[edit services event-options redundancy-event event-name]
user@gateway1# set monitor link-down [interface-name]
```

- c. Specify that a process routing daemon restart request triggers a services switchover.

```
[edit services event-options redundancy-event event-name]
user@gateway1# set monitor process routing restart
```

- d. Specify that a process routing daemon abort request triggers a services switchover.

```
[edit services event-options redundancy-event event-name]
user@gateway1# set monitor process routing abort
```

2. On gateway1, configure a redundancy policy that releases mastership and deletes a static route when the redundancy event conditions are met.
 - a. Configure a name for the policy.


```
user@gateway1# edit policy-options redundancy-policy policy-name
```

For example:

```
user@gateway1# edit policy-options redundancy-policy RLS_MSHIP_POL
```

- b. Specify the redundancy events that release mastership.

```
[edit policy-options redundancy-policy policy-name]
user@gateway1# set redundancy-events [event-list]
```

For example:

```
[edit policy-options redundancy-policy RLS_MSHIP_POL]
user@gateway1# set redundancy-events RELS_MSHIP_CRIT_EV
```

If you want to be able to run the **request services redundancy-set *redundancy-set* trigger redundancy-event *event-name* <force>** to manually release mastership, include that *event-name* in the redundancy policy. The redundancy event itself does not need to be configured, because it is triggered by the request command.

For example:

```
[edit policy-options redundancy-policy RLS_MSHIP_POL]
user@gateway1# set redundancy-events [RELS_MSHIP_CRIT_EV RELS_MSHIP_MANUAL_EV]
```

- c. Release mastership.

```
[edit policy-options redundancy-policy policy-name]
user@gateway1# set then release-mastership force
```

- d. Delete the static route.

```
[edit policy-options redundancy-policy policy-name]
```

```
user@gateway1# set then delete-static-route destination (receive | next-hop next-hop) routing-instance routing-instance
```

3. On gateway1, configure a redundancy policy that acquires mastership from gateway2 when you perform a manual request on gateway1 (**request services redundancy-set *redundancy-set* trigger redundancy-event *event-name* <force>**) .
 - a. Configure a name for the policy.

```
user@gateway1# edit policy-options redundancy-policy policy-name
```

For example:

```
user@gateway1# edit policy-options redundancy-policy ACQU_MSHIP_POL
```

- b. Specify the name of the redundancy event that the manual request uses.

```
[edit policy-options redundancy-policy policy-name]
user@gateway1# set redundancy-events event-name
```

For example:

```
[edit policy-options redundancy-policy ACQU_MSHIP_POL]
user@gateway1# set redundancy-events ACQU_MSHIP_MANUAL_EV
```

The redundancy event itself does not need to be configured, because it is triggered by the request command.

- c. Acquire mastership.

```
[edit policy-options redundancy-policy policy-name]
user@host# set then acquire-mastership
```

4. On gateway1, configure the redundancy set.

- a. Configure a name for the redundancy set.

```
[edit services]
user@gateway1# set redundancy-set redundancy-set
```

For example:

```
[edit services]
user@gateway1# set redundancy-set 1
```

- b. Specify the redundancy group ID for the redundancy set.

```
[edit services redundancy-set redundancy-set]
user@gateway1# set redundancy-group redundancy-group
```

For example:

```
[edit services redundancy-set 1]
user@gateway1# set redundancy-group 1
```

The redundancy group ID is the same redundancy group ID configured for the ICCP daemon (iccpd) through the existing ICCP configuration hierarchy. For example,

```
iccp {
  local-ip-addr 1.1.1.1;
  peer 2.2.2.2 {
    redundancy-group-id-list 1;
    liveness-detection {
      minimum-interval 1000;
    }
  }
}
```

- c. Specify the redundancy policy that releases mastership and the redundancy policy that acquires mastership.

```
[edit services redundancy-set redundancy-set]
user@gateway1# set redundancy-policy [redundancy-policy-list]
```

For example:

```
[edit services redundancy-set 1]
user@gateway1# set redundancy-policy [ ACQU_MSHIP_POL RLS_MSHIP_POL]
```

- d. Configure the frequency of health check probes of the redundancy set, in seconds.

```
[edit services redundancy-set redundancy-set]
```

```
user@gateway1# set healthcheck-timer-interval healthcheck-timer-interval
```

The default is 30 seconds.

- e. Configure the maximum wait time for a help check response, in seconds.

```
[edit services redundancy-set redundancy-set]
user@gateway1# set hold-time hold-time
```

The range is 0 through 3600 seconds.

- f. Configure the frequency of srd hello messages, in seconds.

```
[edit services redundancy-set redundancy-set]
user@gateway1# set keepalive keepalive
```

The range is 1 through 60 seconds.

5. On gateway1, configure routing policies.

- a. Identify signal routes that requires redundancy-related routing changes. Specify the signal route and the routing table that is used.

```
[edit policy-options condition condition-name]
user@gateway1# set if-route-exists signal-route table routing-table
```

For example:

```
[edit policy-options condition switchover-route-exists]
user@gateway1# set if-route-exists 10.45.45.0/24 table bgp1_table
```

- b. To change the local-preference for the signal route, enter it in a policy statement.

```
[edit policy-options policy-statement policy-name]
user@gateway1# set term term from protocol [protocol variables] prefix-list prefix-list condition
condition-name then local-preference preference-value accept
```

- c. To change as-path-prepend values for the signal route, enter them in the policy statement.

```
[edit policy-options policy-statement policy-name]
```

```
user@gateway1# set term term from prefix-list prefix-list condition condition-name then as-path-prepend
[as-prepend-values] next-hop self accept
```

6. On gateway1, configure redundancy for the service set by assigning the redundancy set to the service set.

```
[edit]
user@gateway1# set services service-set service-set-name redundancy-set-id redundancy-set
```

7. On gateway2, the initial standby, configure a redundancy event to identify when the peer gateway releases mastership.

```
[edit services]
user@gateway2# set event-options redundancy-event event-name monitor peer release-mastership
```

For example:

```
[edit services]
user@gateway2# set event-options redundancy-event PEER_RELS_MSHIP_EV monitor peer
release-mastership
```

8. On gateway2, configure a redundancy policy that acquires mastership from the peer gateway and adds a static route.

- a. Configure a name for the policy.

```
user@gateway2# edit policy-options redundancy-policy policy-name
```

For example:

```
user@gateway2# edit policy-options redundancy-policy ACQU_MSHIP_POL
```

- b. Specify the configured redundancy event for the peer gateway mastership release event.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set redundancy-events event-name
```

For example:

```
[edit policy-options redundancy-policy ACQU_MSHIP_POL]
user@gateway2# set redundancy-events PEER_RELS_MSHIP_EV
```

- c. Acquire mastership.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set then acquire-mastership
```

- d. Add a static route.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set then add-static-route destination (receive | next-hop next-hop) routing-instance
routing-instance
```

9. On gateway2, configure a redundancy event to identify when the peer gateway requests mastership.

```
[edit services]
user@gateway2# set event-options redundancy-event event-name monitor peer mastership-acquire
```

For example:

```
[edit services]
user@gateway2# set event-options redundancy-event PEER_MSHIP_ACQU_EV monitor peer
mastership-acquire
```

10. On gateway2, configure a redundancy policy that releases mastership and deletes a static route when gateway1 requests mastership.

- a. Configure a name for the policy.

```
user@gateway2# edit policy-options redundancy-policy policy-name
```

For example:

```
user@gateway2# edit policy-options redundancy-policy RELS-MSHIP_POL
```

- b. Specify the configured redundancy event that identifies when the peer gateway requests mastership.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set redundancy-events event-name
```

For example:

```
[edit policy-options redundancy-policy RELS-MSHIP_POL]
user@gateway2# set redundancy-events PEER_MSHIP_ACQU_EV
```

- c. Release mastership.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set then release-mastership force
```

- d. Delete the static route.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set then delete-static-route destination (receive | next-hop next-hop) routing-instance
routing-instance
```

11. On gateway2, configure one or more redundancy events to monitor the conditions that trigger a warning.

- a. Configure a name for the redundancy event.

```
[edit services]
user@gateway2# set event-options redundancy-event event-name
```

For example:

```
[edit services]
user@gateway2# set event-options redundancy-event WARN_EV
```

- b. Specify any interfaces that trigger a warning when the interface goes down.

```
[edit services event-options redundancy-event event-name]
user@gateway2# set monitor link-down [interface-name]
```

- c. Specify that a process routing daemon restart request triggers a warning.

```
[edit services event-options redundancy-event event-name]
user@gateway2# set monitor process routing restart
```

- d. Specify that a process routing daemon abort request triggers a warning.

```
[edit services event-options redundancy-event event-name]
user@gateway2# set monitor process routing abort
```

12. On gateway2, configure a redundancy policy that broadcasts a warning.

- a. Configure a name for the policy.

```
user@gateway2# edit policy-options redundancy-policy policy-name
```

For example:

```
user@gateway2# edit policy-options redundancy-policy WARN_POL
```

- b. Specify the configured redundancy events that trigger a warning.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set redundancy-events [event-list]
```

For example:

```
[edit policy-options redundancy-policy WARN_POL]
user@gateway2# set redundancy-events WARN_EV
```

- c. Broadcast the warning.

```
[edit policy-options redundancy-policy policy-name]
user@gateway2# set then broadcast-warning
```

13. On gateway2, configure the redundancy set.

- a. Configure a name for the redundancy set.

```
[edit services]
user@gateway2# set redundancy-set redundancy-set
```


For example:

```
[edit services]
user@gateway2# set redundancy-set 1
```

- b. Specify the redundancy group ID for the redundancy set.

```
[edit services redundancy-set redundancy-set]
user@gateway2# set redundancy-group redundancy-group
```

For example:

```
[edit services redundancy-set 1]
user@gateway2# set redundancy-group 1
```

The redundancy group ID is the same redundancy group ID configured for the ICCP daemon (iccpd) through the existing ICCP configuration hierarchy. For example,

```
iccp {
  local-ip-addr 10.1.1.1;
  peer 10.2.2.2 {
    redundancy-group-id-list 1;
    liveness-detection {
      minimum-interval 1000;
    }
  }
}
```

- c. Specify the redundancy policy that releases mastership, the redundancy policy that acquires mastership, and the redundancy policy that triggers a warning.

```
[edit services redundancy-set redundancy-set]
user@gateway2# set redundancy-policy [redundancy-policy-list]
```

For example:

```
[edit services redundancy-set 1]
user@gateway2# set redundancy-policy [ ACQU_MSHIP_POL RLS_MSHIP_POL WARN_POL]
```

- d. Configure the frequency of health check probes of the redundancy set, in seconds.

```
[edit services redundancy-set redundancy-set]
user@gateway2# set healthcheck-timer-interval healthcheck-timer-interval
```

The default is 30 seconds.

- e. Configure the maximum wait time for a health check response, in seconds.

```
[edit services redundancy-set redundancy-set]
user@gateway2# set hold-time hold-time
```

The range is 0 through 3600 seconds.

- f. Configure the frequency of srd hello messages, in seconds.

```
[edit services redundancy-set redundancy-set]
user@gateway2# set keepalive keepalive
```

The range is 1 through 60 seconds.

14. On gateway2, configure routing policies.

- a. Identify signal routes that requires redundancy-related routing changes. Specify the signal route and the routing table that is used.

```
[edit policy-options condition condition-name]
user@gateway2# set if-route-exists signal-route table routing-table
```

For example:

```
[edit policy-options condition switchover-route-exists]
user@gateway2# set if-route-exists 10.45.45.0/24 table bgp1_table
```

- b. To change the local-preference for the signal route, enter it in a policy statement.

```
[edit policy-options policy-statement policy-name]
user@gateway2# set term term from protocol [protocol variables] prefix-list prefix-list condition
condition-name then local-preference preference-value accept
```

- c. To change as-path-prepend values for the signal route, enter them in the policy statement.

```
[edit policy-options policy-statement policy-name]  
user@gateway2# set term term from prefix-list prefix-list condition condition-name then as-path-prepend  
[as-prepend-values] next-hop self accept
```

15. On gateway2, configure redundancy for the service set by assigning the redundancy set to the service set.

```
[edit]  
user@gateway2# set services service-set service-set-name redundancy-set-id redundancy-set
```

SEE ALSO

| [Inter-Chassis Services Redundancy Overview for Next Gen Services](#) | 388

11

PART

Application Layer Gateways

Enabling Traffic to Pass Securely Using Application Layer Gateways | **411**

Enabling Traffic to Pass Securely Using Application Layer Gateways

IN THIS CHAPTER

- [Next Gen Services Application Layer Gateways | 411](#)
- [Configuring Application Sets | 421](#)
- [Configuring Application Properties for Next Gen Services | 421](#)
- [Examples: Configuring Application Protocols | 437](#)
- [Verifying the Output of ALG Sessions | 438](#)

Next Gen Services Application Layer Gateways

IN THIS SECTION

- [RTSP | 411](#)
- [SIP | 412](#)
- [Configuring SIP | 412](#)

This topic describes the Application Layer Gateways (ALGs) supported by Junos OS for Next Gen Services. ALG support includes managing pinholes and parent-child relationships for the supported ALGs.

RTSP

The Real-Time Streaming Protocol (RTSP) controls the delivery of data with real-time properties such as audio and video. The streams controlled by RTSP can use RTP, but it is not required. Media can be transmitted on the same RTSP control stream. This is an HTTP-like text-based protocol, but client and server maintain session information. A session is established using the SETUP message and terminated

using the TEARDOWN message. The transport (the media protocol, address, and port numbers) is negotiated in the setup and the setup-response.

Support for stateful firewall and NAT services requires that you configure the RTSP ALG for TCP port 554.

The ALG monitors the control connection, opens flows dynamically for media (RTP/RTSP) streams, and performs NAT address and port rewrites.

SIP

The Session Initiation Protocol (SIP) is an application layer protocol that can establish, maintain, and terminate media sessions. It is a widely used voice over IP (VoIP) signaling protocol. The SIP ALG monitors SIP traffic and dynamically creates and manages pinholes on the signaling and media paths. The ALG only allows packets with the correct permissions. The SIP ALG also performs the following functions:

- Manages parent-child session relationships.
- Enforces security policies.
- Manages pinholes for VoIP traffic.

The SIP ALG supports the following features:

- Stateful firewall
- Static source NAT
- Dynamic address only source NAT
- Network Address Port Translation (NAPT)

NOTE: SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. SIP sessions on the MS-DPC have no time limit.

Configuring SIP

The Session Initiation Protocol (SIP) is a generalized protocol for communication between endpoints involved in Internet services such as telephony, fax, video conferencing, instant messaging, and file exchange.

The Junos OS provides ALG services in accordance with the standard described in RFC 3261, *SIP: Session Initiation Protocol*. SIP flows under the Junos OS are as described in RFC 3665, *Session Initiation Protocol (SIP) Basic Call Flow Examples*.

NOTE: Before implementing the Junos OS SIP ALG, you should be familiar with certain limitations, discussed in [“Junos OS SIP ALG Limitations” on page 420](#)

The use of NAT in conjunction with the SIP ALG results in changes in SIP header fields due to address translation. For an explanation of these translations, refer to [“SIP ALG Interaction with Network Address Translation” on page 414](#).

To implement SIP on adaptive services interfaces, you configure the **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level with the value **sip**. In addition, there are two other statements you can configure to modify how SIP is implemented:

- You can enable the router to accept any incoming SIP calls for the endpoint devices that are behind the NAT firewall. When a device behind the firewall registers with the proxy that is outside the firewall, the AS or Multiservices PIC maintains the registration state. When the **learn-sip-register** statement is enabled, the router can use this information to accept inbound calls. If this statement is not configured, no inbound calls are accepted; only the devices behind the firewall can call devices outside the firewall.

To configure SIP registration, include the **learn-sip-register** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]  
learn-sip-register;
```

NOTE: The **learn-sip-register** statement is not applicable to the Next Gen Services MX-SPC3.

You can also manually inspect the SIP register by issuing the **show services stateful-firewall sip-register** command; for more information, see the *Junos OS System Basics and Services Command Reference*. The **show services stateful-firewall sip-register** command is not supported for Next Gen Services.

- You can specify a timeout period for the duration of SIP calls that are placed on hold. When a call is put on hold, there is no activity and flows might time out after the configured **inactivity-timeout** period expires, resulting in call state teardown. To avoid this, when a call is put on hold, the flow timer is reset to the **sip-call-hold-timeout** cycle to preserve the call state and flows for longer than the **inactivity-timeout** period.

NOTE: The **sip-call-hold-timeout** statement is not applicable to the Next Gen Services MX-SPC3.

To configure a timeout period, include the **sip-call-hold-timeout** statement at the [edit applications application *application-name*] hierarchy level:

```
[edit applications application application-name]
sip-call-hold-timeout seconds;
```

The default value is 7200 seconds and the range is from 0 through 36,000 seconds (10 hours).

SIP ALG Interaction with Network Address Translation

IN THIS SECTION

- [Outgoing Calls | 415](#)
- [Incoming Calls | 415](#)
- [Forwarded Calls | 416](#)
- [Call Termination | 416](#)
- [Call Re-INVITE Messages | 416](#)
- [Call Session Timers | 416](#)
- [Call Cancellation | 416](#)
- [Forking | 417](#)
- [SIP Messages | 417](#)
- [SIP Headers | 417](#)
- [SIP Body | 419](#)

The Network Address Translation (NAT) protocol enables multiple hosts in a private subnet to share a single public IP address to access the Internet. For outgoing traffic, NAT replaces the private IP address of the host in the private subnet with the public IP address. For incoming traffic, the public IP address is converted back into the private address, and the message is routed to the appropriate host in the private subnet.

Using NAT with the Session Initiation Protocol (SIP) service is more complicated because SIP messages contain IP addresses in the SIP headers as well as in the SIP body. When using NAT with the SIP service, the SIP headers contain information about the caller and the receiver, and the device translates this information to hide it from the outside network. The SIP body contains the Session Description Protocol (SDP) information, which includes IP addresses and port numbers for transmission of the media. The device translates SDP information for allocating resources to send and receive the media.

How IP addresses and port numbers in SIP messages are replaced depends on the direction of the message. For an outgoing message, the private IP address and port number of the client are replaced with the public

IP address and port number of the Juniper Networks firewall. For an incoming message, the public address of the firewall is replaced with the private address of the client.

When an INVITE message is sent out across the firewall, the SIP Application Layer Gateway (ALG) collects information from the message header into a call table, which it uses to forward subsequent messages to the correct endpoint. When a new message arrives, for example an ACK or 200 OK, the ALG compares the “From:, To:, and Call-ID:” fields against the call table to identify the call context of the message. If a new INVITE message arrives that matches the existing call, the ALG processes it as a REINVITE.

When a message containing SDP information arrives, the ALG allocates ports and creates a NAT mapping between them and the ports in the SDP. Because the SDP requires sequential ports for the Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) channels, the ALG provides consecutive even-odd ports. If it is unable to find a pair of ports, it discards the SIP message.

This topic contains the following sections:

Outgoing Calls

When a SIP call is initiated with a SIP request message from the internal to the external network, NAT replaces the IP addresses and port numbers in the SDP and binds the IP addresses and port numbers to the Juniper Networks firewall. Via, Contact, Route, and Record-Route SIP header fields, if present, are also bound to the firewall IP address. The ALG stores these mappings for use in retransmissions and for SIP response messages.

The SIP ALG then opens pinholes in the firewall to allow media through the device on the dynamically assigned ports negotiated based on information in the SDP and the Via, Contact, and Record-Route header fields. The pinholes also allow incoming packets to reach the Contact, Via, and Record-Route IP addresses and ports. When processing return traffic, the ALG inserts the original Contact, Via, Route, and Record-Route SIP fields back into packets.

Incoming Calls

Incoming calls are initiated from the public network to public static NAT addresses or to interface IP addresses on the device. Static NATs are statically configured IP addresses that point to internal hosts; interface IP addresses are dynamically recorded by the ALG as it monitors REGISTER messages sent by internal hosts to the SIP registrar. When the device receives an incoming SIP packet, it sets up a session and forwards the payload of the packet to the SIP ALG.

The ALG examines the SIP request message (initially an INVITE) and, based on information in the SDP, opens gates for outgoing media. When a 200 OK response message arrives, the SIP ALG performs NAT on the IP addresses and ports and opens pinholes in the outbound direction. (The opened gates have a short time-to-live, and they time out if a 200 OK response message is not received quickly.)

When a 200 OK response arrives, the SIP proxy examines the SDP information and reads the IP addresses and port numbers for each media session. The SIP ALG on the device performs NAT on the addresses and port numbers, opens pinholes for outbound traffic, and refreshes the timeout for gates in the inbound direction.

When the ACK arrives for the 200 OK, it also passes through the SIP ALG. If the message contains SDP information, the SIP ALG ensures that the IP addresses and port numbers are not changed from the previous INVITE—if they are, the ALG deletes old pinholes and creates new pinholes to allow media to pass through. The ALG also monitors the Via, Contact, and Record-Route SIP fields and opens new pinholes if it determines that these fields have changed.

Forwarded Calls

A forwarded call is when, for example, user A outside the network calls user B inside the network, and user B forwards the call to user C outside the network. The SIP ALG processes the INVITE from user A as a normal incoming call. But when the ALG examines the forwarded call from B to C outside the network and notices that B and C are reached using the same interface, it does not open pinholes in the firewall, because media will flow directly between user A and user C.

Call Termination

The BYE message terminates a call. When the device receives a BYE message, it translates the header fields just as it does for any other message. But because a BYE message must be acknowledged by the receiver with a 200 OK, the ALG delays call teardown for five seconds to allow time for transmission of the 200 OK.

Call Re-INVITE Messages

Re-INVITE messages add new media sessions to a call and remove existing media sessions. When new media sessions are added to a call, new pinholes are opened in the firewall and new address bindings are created. The process is identical to the original call setup. When one or more media sessions are removed from a call, pinholes are closed and bindings released just as with a BYE message.

Call Session Timers

The SIP ALG uses the Session-Expires value to time out a session if a Re-INVITE or UPDATE message is not received. The ALG gets the Session-Expires value, if present, from the 200 OK response to the INVITE and uses this value for signaling timeout. If the ALG receives another INVITE before the session times out, it resets all timeout values to this new INVITE or to default values, and the process is repeated.

As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the device is protected should one of the following events occur:

- End systems crash during a call and a BYE message is not received.
- Malicious users never send a BYE in an attempt to attack a SIP ALG.
- Poor implementations of SIP proxy fail to process Record-Route and never send a BYE message.
- Network failures prevent a BYE message from being received.

Call Cancellation

Either party can cancel a call by sending a CANCEL message. Upon receiving a CANCEL message, the SIP ALG closes pinholes through the firewall—if any have been opened—and releases address bindings. Before releasing the resources, the ALG delays the control channel age-out for approximately five seconds to

allow time for the final 200 OK to pass through. The call is terminated when the five second timeout expires, regardless of whether a 487 or non-200 response arrives.

Forking

Forking enables a SIP proxy to send a single INVITE message to multiple destinations simultaneously. When the multiple 200 OK response messages arrive for the single call, the SIP ALG parses but updates call information with the first 200 OK messages it receives.

SIP Messages

The SIP message format consists of a SIP header section and the SIP body. In request messages, the first line of the header section is the request line, which includes the method type, request-URI, and protocol version. In response messages, the first line is the status line, which contains a status code. SIP headers contain IP addresses and port numbers used for signaling. The SIP body, separated from the header section by a blank line, is reserved for session description information, which is optional. Junos OS currently supports the SDP only. The SIP body contains IP addresses and port numbers used to transport the media.

SIP Headers

In the following sample SIP request message, NAT replaces the IP addresses in the header fields to hide them from the outside network.

```
INVITE bob@10.150.20.5 SIP/2.0
Via: SIP/2.0/UDP 10.150.20.3:5434
From: alice@10.150.20.3
To: bob@10.150.20.5
Call-ID: a12abcde@10.150.20.3
Contact: alice@10.150.20.3:5434
Route: <sip:netscreen@10.150.20.3:5060>
Record-Route: <sip:netscreen@10.150.20.3:5060>
```

How IP address translation is performed depends on the type and direction of the message. A message can be any of the following:

- Inbound request
- Outbound response
- Outbound request
- Inbound response

[Table 48 on page 418](#) shows how NAT is performed in each of these cases. Note that for several of the header fields the ALG determine more than just whether the messages comes from inside or outside the network. It must also determine what client initiated the call, and whether the message is a request or response.

Table 48: Requesting Messages with NAT Table

Inbound Request (from public to private)	To:	Replace domain with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	Replace ALG address with local address
	Contact:	None
	Record-Route:	None
	Route:	None
Outbound Response (from private to public)	To:	Replace ALG address with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	N/A
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	None

Table 48: Requesting Messages with NAT Table (*continued*)

Outbound Request (from private to public)	To:	None
	From:	Replace local address with ALG address
	Call-ID:	None
	Via:	Replace local address with ALG address
	Request-URI:	None
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	Replace ALG address with local address
Outbound Response (from public to private)	To:	None
	From:	Replace ALG address with local address
	Call-ID:	None
	Via:	Replace ALG address with local address
	Request-URI:	N/A
	Contact:	None
	Record-Route:	Replace ALG address with local address
	Route:	Replace ALG address with local address

SIP Body

The SDP information in the SIP body includes IP addresses the ALG uses to create channels for the media stream. Translation of the SDP section also allocates resources, that is, port numbers to send and receive the media.

The following excerpt from a sample SDP section shows the fields that are translated for resource allocation.

```
o=user 2344234 55234434 IN IP4 10.150.20.3
c=IN IP4 10.150.20.3
m=audio 43249 RTP/AVP 0
```

SIP messages can contain more than one media stream. The concept is similar to attaching multiple files to an e-mail message. For example, an INVITE message sent from a SIP client to a SIP server might have the following fields:

```
c=IN IP4 10.123.33.4
m=audio 33445 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33447 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33449 RTP/AVP 0
```

Junos OS supports up to 6 SDP channels negotiated for each direction, for a total of 12 channels per call.

Junos OS SIP ALG Limitations

The following limitations apply to configuration of the SIP ALG:

- Only the methods described in RFC 3261 are supported.
- Only SIP version 2 is supported.
- TCP is not supported as a transport mechanism for signaling messages for MS-MPCs but is supported for Next Gen Services.
- *Do not configure the SIP ALG when using STUN.* If clients use STUN/TURN to detect the firewall or NAT devices between the caller and responder or proxy, the client attempts to best-guess the NAT device behavior and act accordingly to place the call.
- On MS-MPCs, do not use the endpoint-independent mapping NAT pool option in conjunction with the SIP ALG. Errors will result. This does not apply to Next Gen Services.
- IPv6 signaling data is not supported for MS-MPCs but is supported for Next Gen Services.
- Authentication is not supported.
- Encrypted messages are not supported.
- SIP fragmentation is not supported for MS-MPCs but is supported for Next Gen Services.
- The maximum UDP packet size containing a SIP message is assumed to be 9 KB. SIP messages larger than this are not supported.
- The maximum number of media channels in a SIP message is assumed to be six.
- Fully qualified domain names (FQDNs) are not supported in critical fields.
- QoS is not supported. SIP supports DSCP rewrites.
- High availability is not supported, except for warm standby.
- A timeout setting of never is not supported on SIP or NAT.
- Multicast (forking proxy) is not supported.

RELATED DOCUMENTATION

[ALG Descriptions](#)

[ALGs Available for Junos OS Address Aware NAT](#)

Configuring Application Sets

You can group the applications you have defined into a named object by including the **application-set** statement at the **[edit applications]** hierarchy level with an **application** statement for each application:

```
[edit applications]
  application-set application-set-name {
    application application;
  }
```

For an example of a typical application set, see [“Examples: Configuring Application Protocols” on page 437](#).

RELATED DOCUMENTATION

[ALG Descriptions](#)

[Configuring Application Properties](#)

[Examples: Configuring Application Protocols | 437](#)

[Verifying the Output of ALG Sessions | 438](#)

Configuring Application Properties for Next Gen Services

IN THIS SECTION

- [Configuring an Application Protocol | 422](#)
- [Configuring the Network Protocol | 424](#)
- [Configuring the ICMP Code and Type | 425](#)
- [Configuring Source and Destination Ports | 427](#)
- [Configuring the Inactivity Timeout Period | 428](#)
- [Configuring SIP | 428](#)
- [Configuring an SNMP Command for Packet Matching | 437](#)

To configure application properties, include the **application** statement at the **[edit applications]** hierarchy level:

```
[edit applications]
application application-name {
  application-protocol protocol-name;
  child-inactivity-timeout seconds;
  destination-port port-number;
  gate-timeout seconds;
  icmp-code value;
  icmp-type value;
  inactivity-timeout value;
  protocol type;
  rpc-program-number number;
  snmp-command command;
  source-port port-number;
  ttl-threshold value;
  uuid hex-value;
}
```

You can group application objects by configuring the **application-set** statement; for more information, see [“Configuring Application Sets” on page 421](#).

This section includes the following tasks for configuring applications:

Configuring an Application Protocol

The **application-protocol** statement allows you to specify which of the supported application protocols (ALGs) to configure and include in an application set for service processing. To configure application protocols, include the **application-protocol** statement at the **[edit applications application application-name]** hierarchy level:

```
[edit applications application application-name]
application-protocol protocol-name;
```

[Table 49 on page 422](#) shows the list of supported protocols for Next Gen Services. For more information about specific protocols, see *ALG Descriptions*.

Table 49: Application Protocols Supported by Services Interfaces

Protocol Name	CLI Value	Comments
Bootstrap protocol (BOOTP)	bootp	Supports BOOTP and dynamic host configuration protocol (DHCP).

Table 49: Application Protocols Supported by Services Interfaces (*continued*)

Protocol Name	CLI Value	Comments
Distributed Computing Environment (DCE) remote procedure call (RPC)	dce-rpc	Requires the protocol statement to have the value udp or tcp . Requires a uuid value. You cannot specify destination-port or source-port values.
DCE RPC portmap	dce-rpc-portmap	Requires the protocol statement to have the value udp or tcp . Requires a destination-port value.
Domain Name System (DNS)	dns	Requires the protocol statement to have the value udp . This application protocol closes the DNS flow as soon as the DNS response is received.
Exec	exec	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
FTP	ftp	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
H.323	h323	–
Internet Control Message Protocol (ICMP)	icmp	Requires the protocol statement to have the value icmp or to be unspecified.
IP	ip	–
Login	login	–
NetBIOS	netbios	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.
NetShow	netshow	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
RealAudio	realaudio	–
Real-Time Streaming Protocol (RTSP)	rtsp	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port value.
Session Initiation Protocol	sip	–
SNMP	snmp	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.

Table 49: Application Protocols Supported by Services Interfaces (*continued*)

Protocol Name	CLI Value	Comments
SQLNet	sqlnet	Requires the protocol statement to have the value tcp or to be unspecified. Requires a destination-port or source-port value.
Talk Program	talk	
Trace route	traceroute	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.
Trivial FTP (TFTP)	tftp	Requires the protocol statement to have the value udp or to be unspecified. Requires a destination-port value.
WinFrame	winframe	–

NOTE: You can configure application-level gateways (ALGs) for ICMP and trace route under stateful firewall, NAT, or CoS rules when twice NAT is configured in the same service set. These ALGs cannot be applied to flows created by the Packet Gateway Controller Protocol (PGCP). Twice NAT does not support any other ALGs. NAT applies only the IP address and TCP or UDP headers, but not the payload.

For more information about configuring twice NAT, see *Junos Address Aware Network Addressing Overview*.

RELATED DOCUMENTATION

ALGs Available for Junos OS Address Aware NAT

Configuring the Network Protocol

The **protocol** statement allows you to specify which of the supported network protocols to match in an application definition. To configure network protocols, include the **protocol** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
  protocol type;
```

You specify the protocol type as a numeric value; for the more commonly used protocols, text names are also supported in the command-line interface (CLI). [Table 50 on page 425](#) shows the list of the supported protocols.

Table 50: Network Protocols Supported by Next Gen Services

Network Protocol Type	CLI Value	Comments
External Gateway Protocol (EGP)	egp	–
Generic routing encapsulation (GR)	gre	–
ICMP	icmp	Requires an application-protocol value of icmp .
ICMPv6	icmp6	Requires an application-protocol value of icmp .
Internet Group Management Protocol (IGMP)	igmp	–
TCP	tcp	Requires a destination-port or source-port value unless you specify application-protocol rcp or dce-rcp .
UDP	udp	Requires a destination-port or source-port value unless you specify application-protocol rcp or dce-rcp .

For a complete list of possible numeric values, see RFC 1700, *Assigned Numbers (for the Internet Protocol Suite)*.

NOTE: IP version 6 (IPv6) is not supported as a network protocol in application definitions.

By default, the twice NAT feature can affect IP, TCP, and UDP headers embedded in the payload of ICMP error messages. You can include the **protocol tcp** and **protocol udp** statements with the application statement for twice NAT configurations. For more information about configuring twice NAT, see *Junos Address Aware Network Addressing Overview*.

Configuring the ICMP Code and Type

The ICMP code and type provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. To configure ICMP settings, include the **icmp-code** and **icmp-type** statements at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
icmp-code value;
icmp-type value;
```

You can include only one ICMP code and type value. The **application-protocol** statement must have the value **icmp**. [Table 51 on page 426](#) shows the list of supported ICMP values.

Table 51: ICMP Codes and Types Supported by Services Interfaces

CLI Statement	Description
icmp-code	<p>This value or keyword provides more specific information than icmp-type. Because the value's meaning depends upon the associated icmp-type value, you must specify icmp-type along with icmp-code. For more information, see the <i>Routing Policies, Firewall Filters, and Traffic Policers User Guide</i>.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The keywords are grouped by the ICMP type with which they are associated:</p> <p>parameter-problem: ip-header-bad (0), required-option-missing (1)</p> <p>redirect: redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2)</p> <p>time-exceeded: ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0)</p> <p>unreachable: communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)</p>

Table 51: ICMP Codes and Types Supported by Services Interfaces (*continued*)

CLI Statement	Description
icmp-type	<p>Normally, you specify this match in conjunction with the protocol match statement to determine which protocol is being used on the port. For more information, see the <i>Routing Policies, Firewall Filters, and Traffic Policers User Guide</i>.</p> <p>In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): echo-reply (0), echo-request (8), info-reply (16), info-request (15), mask-request (17), mask-reply (18), parameter-problem (12), redirect (5), router-advertisement (9), router-solicit (10), source-quench (4), time-exceeded (11), timestamp (13), timestamp-reply (14), or unreachable (3).</p>

NOTE: If you configure an interface with an input firewall filter that includes a reject action and with a service set that includes stateful firewall rules, the router executes the input firewall filter before the stateful firewall rules are run on the packet. As a result, when the Packet Forwarding Engine sends an ICMP error message out through the interface, the stateful firewall rules might drop the packet because it was not seen in the input direction.

Possible workarounds are to include a forwarding-table filter to perform the reject action, because this type of filter is executed after the stateful firewall in the input direction, or to include an output service filter to prevent the locally generated ICMP packets from going to the stateful firewall service.

Configuring Source and Destination Ports

The TCP or UDP source and destination port provide additional specification, in conjunction with the network protocol, for packet matching in an application definition. To configure ports, include the **destination-port** and **source-port** statements at the [edit applications application *application-name*] hierarchy level:

```
[edit applications application application-name]
destination-port value;
source-port value;
```

You must define one source or destination port. Normally, you specify this match in conjunction with the **protocol** match statement to determine which protocol is being used on the port.

You can specify either a numeric value or one of the text synonyms listed in [Table 52 on page 428](#).

Table 52: Port Names Supported by Next Gen Services

Port Name	Corresponding Port Number
snmp	161
snmptrap	162

For more information about matching criteria, see the *Routing Policies, Firewall Filters, and Traffic Policers User Guide*.

Configuring the Inactivity Timeout Period

You can specify a timeout period for application inactivity. If the software has not detected any activity during the duration, the flow becomes invalid when the timer expires. To configure a timeout period, include the **inactivity-timeout** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
inactivity-timeout seconds;
```

The default value is 14,400 seconds. The value you configure for an application overrides any global value configured at the **[edit interfaces *interface-name* service-options]** hierarchy level; for more information, see *Configuring Default Timeout Settings for Services Interfaces*.

Configuring SIP

The Session Initiation Protocol (SIP) is a generalized protocol for communication between endpoints involved in Internet services such as telephony, fax, video conferencing, instant messaging, and file exchange.

The Junos OS provides ALG services in accordance with the standard described in RFC 3261, *SIP: Session Initiation Protocol*. SIP flows under the Junos OS are as described in RFC 3665, *Session Initiation Protocol (SIP) Basic Call Flow Examples*.

NOTE: Before implementing the Junos OS SIP ALG, you should be familiar with certain limitations, discussed in [“Junos OS SIP ALG Limitations” on page 420](#)

The use of NAT in conjunction with the SIP ALG results in changes in SIP header fields due to address translation. For an explanation of these translations, refer to [“SIP ALG Interaction with Network Address Translation” on page 414](#).

To implement SIP on adaptive services interfaces, you configure the **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level with the value **sip**. In addition, there are two other statements you can configure to modify how SIP is implemented:

- You can enable the router to accept any incoming SIP calls for the endpoint devices that are behind the NAT firewall. When a device behind the firewall registers with the proxy that is outside the firewall, the AS or Multiservices PIC maintains the registration state. When the **learn-sip-register** statement is enabled, the router can use this information to accept inbound calls. If this statement is not configured, no inbound calls are accepted; only the devices behind the firewall can call devices outside the firewall.

To configure SIP registration, include the **learn-sip-register** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]  
learn-sip-register;
```

NOTE: The **learn-sip-register** statement is not applicable to the Next Gen Services MX-SPC3.

You can also manually inspect the SIP register by issuing the **show services stateful-firewall sip-register** command; for more information, see the *Junos OS System Basics and Services Command Reference*. The **show services stateful-firewall sip-register** command is not supported for Next Gen Services.

- You can specify a timeout period for the duration of SIP calls that are placed on hold. When a call is put on hold, there is no activity and flows might time out after the configured **inactivity-timeout** period expires, resulting in call state teardown. To avoid this, when a call is put on hold, the flow timer is reset to the **sip-call-hold-timeout** cycle to preserve the call state and flows for longer than the **inactivity-timeout** period.

NOTE: The **sip-call-hold-timeout** statement is not applicable to the Next Gen Services MX-SPC3.

To configure a timeout period, include the **sip-call-hold-timeout** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]  
sip-call-hold-timeout seconds;
```

The default value is 7200 seconds and the range is from 0 through 36,000 seconds (10 hours).

SIP ALG Interaction with Network Address Translation

IN THIS SECTION

- [Outgoing Calls | 431](#)
- [Incoming Calls | 431](#)
- [Forwarded Calls | 431](#)
- [Call Termination | 432](#)
- [Call Re-INVITE Messages | 432](#)
- [Call Session Timers | 432](#)
- [Call Cancellation | 432](#)
- [Forking | 432](#)
- [SIP Messages | 433](#)
- [SIP Headers | 433](#)
- [SIP Body | 435](#)

The Network Address Translation (NAT) protocol enables multiple hosts in a private subnet to share a single public IP address to access the Internet. For outgoing traffic, NAT replaces the private IP address of the host in the private subnet with the public IP address. For incoming traffic, the public IP address is converted back into the private address, and the message is routed to the appropriate host in the private subnet.

Using NAT with the Session Initiation Protocol (SIP) service is more complicated because SIP messages contain IP addresses in the SIP headers as well as in the SIP body. When using NAT with the SIP service, the SIP headers contain information about the caller and the receiver, and the device translates this information to hide it from the outside network. The SIP body contains the Session Description Protocol (SDP) information, which includes IP addresses and port numbers for transmission of the media. The device translates SDP information for allocating resources to send and receive the media.

How IP addresses and port numbers in SIP messages are replaced depends on the direction of the message. For an outgoing message, the private IP address and port number of the client are replaced with the public IP address and port number of the Juniper Networks firewall. For an incoming message, the public address of the firewall is replaced with the private address of the client.

When an INVITE message is sent out across the firewall, the SIP Application Layer Gateway (ALG) collects information from the message header into a call table, which it uses to forward subsequent messages to the correct endpoint. When a new message arrives, for example an ACK or 200 OK, the ALG compares the “From:, To:, and Call-ID:” fields against the call table to identify the call context of the message. If a new INVITE message arrives that matches the existing call, the ALG processes it as a REINVITE.

When a message containing SDP information arrives, the ALG allocates ports and creates a NAT mapping between them and the ports in the SDP. Because the SDP requires sequential ports for the Real-Time Transport Protocol (RTP) and Real-Time Control Protocol (RTCP) channels, the ALG provides consecutive even-odd ports. If it is unable to find a pair of ports, it discards the SIP message.

This topic contains the following sections:

Outgoing Calls

When a SIP call is initiated with a SIP request message from the internal to the external network, NAT replaces the IP addresses and port numbers in the SDP and binds the IP addresses and port numbers to the Juniper Networks firewall. Via, Contact, Route, and Record-Route SIP header fields, if present, are also bound to the firewall IP address. The ALG stores these mappings for use in retransmissions and for SIP response messages.

The SIP ALG then opens pinholes in the firewall to allow media through the device on the dynamically assigned ports negotiated based on information in the SDP and the Via, Contact, and Record-Route header fields. The pinholes also allow incoming packets to reach the Contact, Via, and Record-Route IP addresses and ports. When processing return traffic, the ALG inserts the original Contact, Via, Route, and Record-Route SIP fields back into packets.

Incoming Calls

Incoming calls are initiated from the public network to public static NAT addresses or to interface IP addresses on the device. Static NATs are statically configured IP addresses that point to internal hosts; interface IP addresses are dynamically recorded by the ALG as it monitors REGISTER messages sent by internal hosts to the SIP registrar. When the device receives an incoming SIP packet, it sets up a session and forwards the payload of the packet to the SIP ALG.

The ALG examines the SIP request message (initially an INVITE) and, based on information in the SDP, opens gates for outgoing media. When a 200 OK response message arrives, the SIP ALG performs NAT on the IP addresses and ports and opens pinholes in the outbound direction. (The opened gates have a short time-to-live, and they time out if a 200 OK response message is not received quickly.)

When a 200 OK response arrives, the SIP proxy examines the SDP information and reads the IP addresses and port numbers for each media session. The SIP ALG on the device performs NAT on the addresses and port numbers, opens pinholes for outbound traffic, and refreshes the timeout for gates in the inbound direction.

When the ACK arrives for the 200 OK, it also passes through the SIP ALG. If the message contains SDP information, the SIP ALG ensures that the IP addresses and port numbers are not changed from the previous INVITE—if they are, the ALG deletes old pinholes and creates new pinholes to allow media to pass through. The ALG also monitors the Via, Contact, and Record-Route SIP fields and opens new pinholes if it determines that these fields have changed.

Forwarded Calls

A forwarded call is when, for example, user A outside the network calls user B inside the network, and user B forwards the call to user C outside the network. The SIP ALG processes the INVITE from user A as

a normal incoming call. But when the ALG examines the forwarded call from B to C outside the network and notices that B and C are reached using the same interface, it does not open pinholes in the firewall, because media will flow directly between user A and user C.

Call Termination

The BYE message terminates a call. When the device receives a BYE message, it translates the header fields just as it does for any other message. But because a BYE message must be acknowledged by the receiver with a 200 OK, the ALG delays call teardown for five seconds to allow time for transmission of the 200 OK.

Call Re-INVITE Messages

Re-INVITE messages add new media sessions to a call and remove existing media sessions. When new media sessions are added to a call, new pinholes are opened in the firewall and new address bindings are created. The process is identical to the original call setup. When one or more media sessions are removed from a call, pinholes are closed and bindings released just as with a BYE message.

Call Session Timers

The SIP ALG uses the Session-Expires value to time out a session if a Re-INVITE or UPDATE message is not received. The ALG gets the Session-Expires value, if present, from the 200 OK response to the INVITE and uses this value for signaling timeout. If the ALG receives another INVITE before the session times out, it resets all timeout values to this new INVITE or to default values, and the process is repeated.

As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the device is protected should one of the following events occur:

- End systems crash during a call and a BYE message is not received.
- Malicious users never send a BYE in an attempt to attack a SIP ALG.
- Poor implementations of SIP proxy fail to process Record-Route and never send a BYE message.
- Network failures prevent a BYE message from being received.

Call Cancellation

Either party can cancel a call by sending a CANCEL message. Upon receiving a CANCEL message, the SIP ALG closes pinholes through the firewall—if any have been opened—and releases address bindings. Before releasing the resources, the ALG delays the control channel age-out for approximately five seconds to allow time for the final 200 OK to pass through. The call is terminated when the five second timeout expires, regardless of whether a 487 or non-200 response arrives.

Forking

Forking enables a SIP proxy to send a single INVITE message to multiple destinations simultaneously. When the multiple 200 OK response messages arrive for the single call, the SIP ALG parses but updates call information with the first 200 OK messages it receives.

SIP Messages

The SIP message format consists of a SIP header section and the SIP body. In request messages, the first line of the header section is the request line, which includes the method type, request-URI, and protocol version. In response messages, the first line is the status line, which contains a status code. SIP headers contain IP addresses and port numbers used for signaling. The SIP body, separated from the header section by a blank line, is reserved for session description information, which is optional. Junos OS currently supports the SDP only. The SIP body contains IP addresses and port numbers used to transport the media.

SIP Headers

In the following sample SIP request message, NAT replaces the IP addresses in the header fields to hide them from the outside network.

```
INVITE bob@10.150.20.5 SIP/2.0
Via: SIP/2.0/UDP 10.150.20.3:5434
From: alice@10.150.20.3
To: bob@10.150.20.5
Call-ID: a12abcde@10.150.20.3
Contact: alice@10.150.20.3:5434
Route: <sip:netscreen@10.150.20.3:5060>
Record-Route: <sip:netscreen@10.150.20.3:5060>
```

How IP address translation is performed depends on the type and direction of the message. A message can be any of the following:

- Inbound request
- Outbound response
- Outbound request
- Inbound response

[Table 48 on page 418](#) shows how NAT is performed in each of these cases. Note that for several of the header fields the ALG determine more than just whether the messages comes from inside or outside the network. It must also determine what client initiated the call, and whether the message is a request or response.

Table 53: Requesting Messages with NAT Table

Inbound Request (from public to private)	To:	Replace domain with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	Replace ALG address with local address
	Contact:	None
	Record-Route:	None
	Route:	None
Outbound Response (from private to public)	To:	Replace ALG address with local address
	From:	None
	Call-ID:	None
	Via:	None
	Request-URI:	N/A
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	None

Table 53: Requesting Messages with NAT Table (*continued*)

Outbound Request (from private to public)	To:	None
	From:	Replace local address with ALG address
	Call-ID:	None
	Via:	Replace local address with ALG address
	Request-URI:	None
	Contact:	Replace local address with ALG address
	Record-Route:	Replace local address with ALG address
	Route:	Replace ALG address with local address
Outbound Response (from public to private)	To:	None
	From:	Replace ALG address with local address
	Call-ID:	None
	Via:	Replace ALG address with local address
	Request-URI:	N/A
	Contact:	None
	Record-Route:	Replace ALG address with local address
	Route:	Replace ALG address with local address

SIP Body

The SDP information in the SIP body includes IP addresses the ALG uses to create channels for the media stream. Translation of the SDP section also allocates resources, that is, port numbers to send and receive the media.

The following excerpt from a sample SDP section shows the fields that are translated for resource allocation.

```
o=user 2344234 55234434 IN IP4 10.150.20.3
c=IN IP4 10.150.20.3
m=audio 43249 RTP/AVP 0
```

SIP messages can contain more than one media stream. The concept is similar to attaching multiple files to an e-mail message. For example, an INVITE message sent from a SIP client to a SIP server might have the following fields:

```
c=IN IP4 10.123.33.4
m=audio 33445 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33447 RTP/AVP 0
c=IN IP4 10.123.33.4
m=audio 33449 RTP/AVP 0
```

Junos OS supports up to 6 SDP channels negotiated for each direction, for a total of 12 channels per call.

Junos OS SIP ALG Limitations

The following limitations apply to configuration of the SIP ALG:

- Only the methods described in RFC 3261 are supported.
- Only SIP version 2 is supported.
- TCP is not supported as a transport mechanism for signaling messages for MS-MPCs but is supported for Next Gen Services.
- *Do not configure the SIP ALG when using STUN.* If clients use STUN/TURN to detect the firewall or NAT devices between the caller and responder or proxy, the client attempts to best-guess the NAT device behavior and act accordingly to place the call.
- On MS-MPCs, do not use the endpoint-independent mapping NAT pool option in conjunction with the SIP ALG. Errors will result. This does not apply to Next Gen Services.
- IPv6 signaling data is not supported for MS-MPCs but is supported for Next Gen Services.
- Authentication is not supported.
- Encrypted messages are not supported.
- SIP fragmentation is not supported for MS-MPCs but is supported for Next Gen Services.
- The maximum UDP packet size containing a SIP message is assumed to be 9 KB. SIP messages larger than this are not supported.
- The maximum number of media channels in a SIP message is assumed to be six.
- Fully qualified domain names (FQDNs) are not supported in critical fields.
- QoS is not supported. SIP supports DSCP rewrites.
- High availability is not supported, except for warm standby.
- A timeout setting of never is not supported on SIP or NAT.
- Multicast (forking proxy) is not supported.

Configuring an SNMP Command for Packet Matching

You can specify an SNMP command setting for packet matching. To configure SNMP, include the **snmp-command** statement at the **[edit applications application *application-name*]** hierarchy level:

```
[edit applications application application-name]
snmp-command value;
```

The supported values are **get**, **get-next**, **set**, and **trap**. You can configure only one value for matching. The **application-protocol** statement at the **[edit applications application *application-name*]** hierarchy level must have the value **snmp**.

Examples: Configuring Application Protocols

The following example shows an application protocol definition describing a special FTP application running on port 78:

```
[edit applications]
application my-ftp-app {
  application-protocol ftp;
  protocol tcp;
  destination-port 78;
  timeout 100; # inactivity timeout for FTP service
}
```

The following example shows a special ICMP protocol (**application-protocol icmp**) of type 8 (ICMP echo):

```
[edit applications]
application icmp-app {
  application-protocol icmp;
  protocol icmp;
  icmp-type icmp-echo;
}
```

The following example shows a possible application set:

```
[edit applications]
application-set basic {
  http;
  ftp;
```

```
telnet;  
nfs;  
icmp;  
}
```

The software includes a predefined set of well-known application protocols. The set includes applications for which the TCP and UDP destination ports are already recognized by stateless firewall filters.

RELATED DOCUMENTATION

ALG Descriptions

[Configuring Application Sets | 421](#)

Configuring Application Properties

[Verifying the Output of ALG Sessions | 438](#)

Verifying the Output of ALG Sessions

IN THIS SECTION

- [FTP Example | 438](#)
- [RTSP ALG Example | 444](#)
- [System Log Messages | 448](#)

This section contains examples of successful output from ALG sessions and information on system log configuration. You can compare the results of your sessions to check whether the configurations are functioning correctly.

FTP Example

IN THIS SECTION

- [Sample Output | 439](#)
- [FTP System Log Messages | 441](#)

- Analysis | 442
- Troubleshooting Questions | 444

This example analyzes the output during an active FTP session. It consists of four different flows; two are control flows and two are data flows. The example consists of the following parts:

Sample Output

MS-MPC Card

For MS-MPCs, the following is a complete sample output from the **show services stateful-firewall conversations application-protocol ftp** operational mode command:

```
user@host>show services stateful-firewall conversations application-protocol ftp
```

```
Interface: ms-1/3/0, Service set: CLBJI1-AAF001
Conversation: ALG protocol: ftp
  Number of initiators: 2, Number of responders: 2
```

Flow	State	Dir	Frm count
TCP	1.1.79.2:14083 ->	2.2.2.2:21	Watch I 13
NAT source	1.1.79.2:14083	-> 194.250.1.237:50118	
TCP	1.1.79.2:14104 ->	2.2.2.2:20	Forward I 3
NAT source	1.1.79.2:14104	-> 194.250.1.237:50119	
TCP	2.2.2.2:21 ->	194.250.1.237:50118	Watch O 12
NAT dest	194.250.1.237:50118	-> 1.1.79.2:14083	
TCP	2.2.2.2:20 ->	194.250.1.237:50119	Forward O 5
NAT dest	194.250.1.237:50119	-> 1.1.79.2:14104	

For each flow, the first line shows flow information, including protocol (TCP), source address, source port, destination address, destination port, flow state, direction, and frame count.

- The state of a flow can be **Watch**, **Forward**, or **Drop**:
 - A **Watch** flow state indicates that the control flow is monitored by the ALG for information in the payload. NAT processing is performed on the header and payload as needed.
 - A **Forward** flow forwards the packets without monitoring the payload. NAT is performed on the header as needed.
 - A **Drop** flow drops any packet that matches the 5 tuple.
- The frame count (**Frm count**) shows the number of packets that were processed on that flow.

The second line shows the NAT information.

- **source** indicates source NAT.
- **dest** indicates destination NAT.
- The first address and port in the NAT line are the original address and port being translated for that flow.
- The second address and port in the NAT line are the translated address and port for that flow.

MX-SPC3 Card

On the MX-SPC3 services card, the following is a complete sample output from the **show services sessions application-protocol ftp** operational mode command:

```
user@host>show services sessions application-protocol ftp
```

```
Session ID: 536870917, Service-set: ssl, Policy name: pl/131085, Timeout: 1, Valid
Logical system: root-logical-system
Resource information : FTP ALG, 1, 1
  In: 12.10.10.10/35281 --> 22.20.20.3/8204;tcp, Conn Tag: 0x0, If: vms-2/0/0.100,
Pkts: 6, Bytes: 320,
  Out: 22.20.20.3/8204 --> 60.1.1.2/48747;tcp, Conn Tag: 0x0, If: vms-2/0/0.200,
Pkts: 9, Bytes: 8239,

Session ID: 536870919, Service-set: ssl, Policy name: pl/131085, Timeout: 29, Valid
Logical system: root-logical-system
Resource information : FTP ALG, 1, 0
  In: 12.10.10.10/44194 --> 22.20.20.3/21;tcp, Conn Tag: 0x0, If: vms-2/0/0.100,
Pkts: 13, Bytes: 585,
  Out: 22.20.20.3/21 --> 60.1.1.2/48660;tcp, Conn Tag: 0x0, If: vms-2/0/0.200,
Pkts: 11, Bytes: 650,
Total sessions: 2
```

For each session:

- The first line shows flow information, including session ID, service-set name, policy name, session timeout, logical system name, and its state.
- The second line, **Resource information**, indicates the session is created by ALG, including the ALG name (FTP ALG) and ASL group id, which is 1 and the ASL resource id, which is 0 for control session and 1 for data session.
- The third line **In** is forward flow and the fourth line **Out** is reverse flow, including the source address, source port, destination address, destination port, protocol (TCP), session conn-tag, incoming for **In** and outgoing for **Out** interface, received frame count and bytes. NAT is performed on the header as needed.

FTP System Log Messages

System log messages are generated during an FTP session. For more information about system logs, see [“System Log Messages” on page 448](#).

MS-MPC Card

The following system log messages are generated during creation of the FTP control flow:

- Rule Accept system log:

```
Oct 27 11:42:54 (FPC Slot 1, PIC Slot 1) {ss_ftp}[FWNAT]: ASP_SFW_RULE_ACCEPT:
proto 6 (TCP) application: ftp, fe-3/3/3.0:1.1.1.2:4450 -> 2.2.2.2:21, Match SFW
accept rule-set:, rule: ftp, term: 1
```

- Create Accept Flow system log:

```
Oct 27 11:42:54 (FPC Slot 1, PIC Slot 1) {ss_ftp}[FWNAT]: ASP_SFW_CREATE_ACCEPT_FLOW: proto 6 (TCP)
application: ftp, fe-3/3/3.0:1.1.1.2:4450 -> 2.2.2.2:21, creating forward or watch flow
```

- System log for data flow creation:

```
Oct 27 11:43:30 (FPC Slot 1, PIC Slot 1) {ss_ftp}[FWNAT]: ASP_SFW_FTP_ACTIVE_ACCEPT: proto 6 (TCP)
application: ftp, so-2/1/2.0:2.2.2.2:20 -> 1.1.1.2:50726, Creating FTP active mode forward flow
```

MX-SPC3 Card

The following system log messages are generated during creation of the FTP control flow:

- System log for FTP control session creation:

```
Mar 23 23:58:54 esst480r RT_FLOW: RT_FLOW_SESSION_CREATE_USF: Tag svc-set-name
ssl: session created 20.1.1.2/52877->30.1.1.2/21 0x0 junos-ftp
20.1.1.2/52877->30.1.1.2/21 0x0 N/A N/A N/A N/A 6 p1 ssl-ZoneIn ssl-ZoneOut
818413576 N/A(N/A) ge-1/0/2.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1 N/A
```

```
Mar 23 23:59:00 esst480r junos-alg: RT_ALG_FTP_ACTIVE_ACCEPT: application:ftp
data, vms-3/0/0.0 30.1.1.2:20 -> 20.1.1.2:33947 (TCP)
```

- System log for FTP data session creation:

```
Mar 23 23:59:00 esst480r RT_FLOW: RT_FLOW_SESSION_CREATE_USF: Tag svc-set-name
ssl: session created 30.1.1.2/20->20.1.1.2/33947 0x0 junos-ftp-data
```

```
30.1.1.2/20->20.1.1.2/33947 0x0 N/A N/A N/A N/A 6 pl ssl-ZoneOut ssl-ZoneIn
818413577 N/A(N/A) ge-1/1/6.0 FTP-DATA UNKNOWN UNKNOWN Infrastructure File-Servers
2 N/A
```

- System log for FTP data session destroy:

```
Mar 23 23:59:02 esst480r RT_FLOW: RT_FLOW_SESSION_CLOSE_USF: Tag svc-set-name
ssl: session closed TCP FIN: 30.1.1.2/20->20.1.1.2/33947 0x0 junos-ftp-data
30.1.1.2/20->20.1.1.2/33947 0x0 N/A N/A N/A N/A 6 pl ssl-ZoneOut ssl-ZoneIn
818413577 2954(4423509) 281(14620) 2 FTP-DATA UNKNOWN N/A(N/A) ge-1/1/6.0 No
Infrastructure File-Servers 2 N/A
```

- System log for FTP control session destroy:

```
Mar 23 23:59:39 esst480r RT_FLOW: RT_FLOW_SESSION_CLOSE_USF: Tag svc-set-name
ssl: session closed Closed by junos-tcp-clt-emul: 20.1.1.2/52877->30.1.1.2/21
0x0 junos-ftp 20.1.1.2/52877->30.1.1.2/21 0x0 N/A N/A N/A N/A 6 pl ssl-ZoneIn
ssl-ZoneOut 818413576 23(1082) 18(1176) 45 UNKNOWN UNKNOWN N/A(N/A) ge-1/0/2.0
No N/A N/A -1 N/A
```

Analysis

Control Flows

MS-MPC Card

The control flows are established after the three-way handshake is complete.

- Control flow from FTP client to FTP server. TCP destination port is 21.

```
TCP          1.1.79.2:14083 ->          2.2.2.2:21    Watch    I          13
NAT source   1.1.79.2:14083  ->    194.250.1.237:50118
```

- Control flow from FTP server to FTP client. TCP source port is 21.

```
TCP          2.2.2.2:21    -> 194.250.1.237:50118 Watch    O          12
NAT dest     194.250.1.237:50118 ->          1.1.79.2:14083
```

MX-SPC3 Card

The control flows are established after the three-way handshake is complete.

- Control session from FTP client to FTP server, TCP destination port is 21.

```

Session ID: 536870919, Service-set: ssl, Policy name: pl/131085, Timeout: 29,
Valid
Logical system: root-logical-system
Resource information : FTP ALG, 1, 0
  In: 12.10.10.10/44194 --> 22.20.20.3/21;tcp, Conn Tag: 0x0, If: vms-2/0/0.100,
  Pkts: 13, Bytes: 585,
  Out: 22.20.20.3/21 --> 60.1.1.2/48660;tcp, Conn Tag: 0x0, If: vms-2/0/0.200,
  Pkts: 11, Bytes: 650,

```

- Data session from FTP client to FTP server, it's for FTP passive mode.

```

Session ID: 536870917, Service-set: ssl, Policy name: pl/131085, Timeout: 1,
Valid
Logical system: root-logical-system
Resource information : FTP ALG, 1, 1
  In: 12.10.10.10/35281 --> 22.20.20.3/8204;tcp, Conn Tag: 0x0, If: vms-2/0/0.100,
  Pkts: 6, Bytes: 320,
  Out: 22.20.20.3/8204 --> 60.1.1.2/48747;tcp, Conn Tag: 0x0, If: vms-2/0/0.200,
  Pkts: 9, Bytes: 8239,

```

- Data session from FTP server to FTP client, it's for FTP active mode:

```

Session ID: 549978117, Service-set: ssl, Policy name: pl/131085, Timeout: 1,
Valid
Logical system: root-logical-system
Resource information : FTP ALG, 1, 1
  In: 22.20.20.3/20 --> 60.1.1.3/6049;tcp, Conn Tag: 0x0, If: vms-2/0/0.200,
  Pkts: 10, Bytes: 8291,
  Out: 12.10.10.10/33203 --> 22.20.20.3/20;tcp, Conn Tag: 0x0, If: vms-2/0/0.100,
  Pkts: 5, Bytes: 268,

```

Data Flows

A data port of 20 is negotiated for data transfer during the course of the FTP control protocol. These two flows are data flows between the FTP client and the FTP server:

TCP	1.1.79.2:14104 ->	2.2.2.2:20	Forward	I	3
NAT source	1.1.79.2:14104	->	194.250.1.237:50119		
TCP	2.2.2.2:20 ->	194.250.1.237:50119	Forward	O	5
NAT dest	194.250.1.237:50119	->	1.1.79.2:14104		

Troubleshooting Questions

1. How do I know if the FTP ALG is active?
 - The ALG protocol field in the conversation should display **ftp**.
 - There should be a valid frame count (**Frm count**) in the control flows.
 - A valid frame count in the data flows indicates that data transfer has taken place.
2. What do I need to check if the FTP connection is established but data transfer does not take place?
 - Most probably, the control connection is up, but the data connection is down.
 - Check the conversations output to determine whether both the control and data flows are present.
3. How do I interpret each flow? What does each flow mean?
 - FTP control flow initiator flow—Flow with destination port 21
 - FTP control flow responder flow—Flow with source port ;21
 - FTP data flow initiator flow—Flow with destination port 20
 - FTP data flow responder flow—Flow with source port 20

RTSP ALG Example

IN THIS SECTION

- [Sample Output for MS-MPCs | 444](#)
- [Sample Output for MX-SPC3 Services Card | 445](#)
- [Analysis | 446](#)
- [Troubleshooting Questions | 446](#)

The following is an example of an RTSP conversation. The application uses the RTSP protocol for control connection. Once the connection is set up, the media is sent using UDP protocol (RTP).

This example consists of the following:

Sample Output for MS-MPCs

Here is the output from the **show services stateful-firewall conversations** operational mode command:

```
user@host# show services stateful-firewall conversations
```

```
Interface: ms-3/2/0, Service set: svc_set
Conversation: ALG protocol: rtsp
  Number of initiators: 5, Number of responders: 5
```

Flow	State	Dir	Frm	count			
TCP	1.1.1.3:58795	->	2.2.2.2:554	Watch	I	7	
UDP	1.1.1.3:1028	->	2.2.2.2:1028	Forward	I	0	
UDP	1.1.1.3:1029	->	2.2.2.2:1029	Forward	I	0	
UDP	1.1.1.3:1030	->	2.2.2.2:1030	Forward	I	0	
UDP	1.1.1.3:1031	->	2.2.2.2:1031	Forward	I	0	
TCP	2.2.2.2:554	->	1.1.1.3:58795	Watch	O	5	
UDP	2.2.2.2:1028	->	1.1.1.3:1028	Forward	O	6	
UDP	2.2.2.2:1029	->	1.1.1.3:1029	Forward	O	0	
UDP	2.2.2.2:1030	->	1.1.1.3:1030	Forward	O	3	
UDP	2.2.2.2:1031	->	1.1.1.3:1031	Forward	O	0	

Sample Output for MX-SPC3 Services Card

Here is the output from the **show services sessions application-protocol rtsp** operational mode command:

```
user@host# run show services sessions application-protocol rtsp
```

```
Session ID: 1073741828, Service-set: sset1, Policy name: pl/131081, Timeout: 116,
Valid
```

```
Logical system: root-logical-system
```

```
Resource information : RTSP ALG, 1, 0
```

```
In: 31.0.0.2/33575 --> 41.0.0.2/554;tcp, Conn Tag: 0x0, If: vms-4/0/0.1, Pkts:
8, Bytes: 948,
```

```
Out: 41.0.0.2/554 --> 131.10.0.1/7777;tcp, Conn Tag: 0x0, If: vms-4/0/0.2, Pkts:
6, Bytes: 1117,
```

```
Session ID: 1073741829, Service-set: sset1, Policy name: pl/131081, Timeout: 120,
Valid
```

```
Logical system: root-logical-system
```

```
Resource information : RTSP ALG, 1, 1
```

```
In: 41.0.0.2/35004 --> 131.10.0.1/7780;udp, Conn Tag: 0x0, If: vms-4/0/0.2, Pkts:
220, Bytes: 79200,
```

```
Out: 31.0.0.2/30004 --> 41.0.0.2/35004;udp, Conn Tag: 0x0, If: vms-4/0/0.1, Pkts:
0, Bytes: 0,
```

```
Session ID: 1073741830, Service-set: sset1, Policy name: pl/131081, Timeout: 120,
Valid
```

```
Logical system: root-logical-system
```

```
Resource information : RTSP ALG, 1, 4
```

```
In: 41.0.0.2/35006 --> 131.10.0.1/7781;udp, Conn Tag: 0x0, If: vms-4/0/0.2, Pkts:
```

```

220, Bytes: 174240,
  Out: 31.0.0.2/30006 --> 41.0.0.2/35006;udp, Conn Tag: 0x0, If: vms-4/0/0.1, Pkts:
0, Bytes: 0,
Total sessions: 3

```

Analysis

An RTSP conversation should consist of TCP flows corresponding to the RTSP control connection. There should be two flows, one in each direction, from client to server and from server to client:

TCP	1.1.1.3:58795 ->	2.2.2.2:554	Watch	I	7
TCP	2.2.2.2:554 ->	1.1.1.3:58795	Watch	O	5

- The RTSP control connection for the initiator flow is sent from destination port 554.
- The RTSP control connection for the responder flow is sent from source port 554.

The UDP flows correspond to RTP media sent over the RTSP connection.

Troubleshooting Questions

1. Media does not work when the RTSP ALG is configured. What do I do?

- Check RTSP conversations to see whether both TCP and UDP flows exist.
- The ALG protocol should be displayed as **rtsp**.

NOTE: The state of the flow is displayed as **Watch**, because the ALG processing is taking place and the client is essentially “watching” or processing payload corresponding to the application. For FTP and RTSP ALG flows, the control connections are always **Watch** flows.

2. How do I check for ALG errors?

- You can check for errors by issuing the following command. Each ALG has a separate field for ALG packet errors.

```
user@host# show services stateful-firewall statistics extensive
```

```

Interface: ms-3/2/0
Service set: svc_set
New flows:
  Accepts: 1347, Discards: 0, Rejects: 0
Existing flows:

```



```

    Accepts: 144187, Discards: 0, Rejects: 0
Drops:
    IP option: 0, TCP SYN defense: 0
    NAT ports exhausted: 0
Errors:
    IP: 0, TCP: 276
    UDP: 0, ICMP: 0
    Non-IP packets: 0, ALG: 0
IP errors:
    IP packet length inconsistencies: 0
    Minimum IP header length check failures: 0
    Reassembled packet exceeds maximum IP length: 0
    Illegal source address: 0
    Illegal destination address: 0
    TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
    Land attack: 0
    Non-IPv4 packets: 0, Bad checksum: 0
    Illegal IP fragment length: 0
    IP fragment overlap: 0
    IP fragment reassembly timeout: 0
    Unknown: 0
TCP errors:
    TCP header length inconsistencies: 0
    Source or destination port number is zero: 0
    Illegal sequence number and flags combinations: 0
    SYN attack (multiple SYN messages seen for the same flow): 276
    First packet not a SYN message: 0
    TCP port scan (TCP handshake, RST seen from server for SYN): 0
    Bad SYN cookie response: 0
UDP errors:
    IP data length less than minimum UDP header length (8 bytes): 0
    Source or destination port number is zero: 0
    UDP port scan (ICMP error seen for UDP flow): 0
ICMP errors:
    IP data length less than minimum ICMP header length (8 bytes): 0
    ICMP error length inconsistencies: 0
    Duplicate ping sequence number: 0
    Mismatched ping sequence number: 0
ALG errors:
    BOOTP: 0, DCE-RPC: 0, DCE-RPC portmap: 0
    DNS: 0, Exec: 0, FTP: 0
    ICMP: 0
    Login: 0, NetBIOS: 0, NetShow: 0
    RPC: 0, RPC portmap: 0

```

```
RTSP: 0, Shell: 0
SNMP: 0, SQLNet: 0, TFTP: 0
Traceroute: 0
```

System Log Messages

IN THIS SECTION

- [System Log Configuration | 448](#)
- [System Log Output | 449](#)

Enabling system log generation and checking the system log are also helpful for ALG flow analysis. This section contains the following:

System Log Configuration

You can configure the enabling of system log messages at a number of different levels in the Junos OS CLI. As shown in the following sample configurations, the choice of level depends on how specific you want the event logging to be and what options you want to include. For details on the configuration options, see the *Junos OS Administration Library* (system level) or the *Junos OS Services Interfaces Library for Routing Devices* (all other levels).

1. At the topmost global level:

```
user@host# show system syslog
file messages {
  any any;
}
```

2. At the service set level:

```
user@host# show services service-set svc_set
syslog {
  host local {
    services any;
  }
}
stateful-firewall-rules allow_rtsp;
interface-service {
```

```
service-interface ms-3/2/0;
}
```

3. At the service rule level:

```
user@host# show services stateful-firewall rule allow_rtsp
match-direction input-output;
term 0 {
  from {
    applications junos-rtsp;
  }
  then {
    accept;
    syslog;
  }
}
```

System Log Output

System log messages are generated during flow creation, as shown in the following examples:

The following system log message indicates that the ASP matched an accept rule:

```
Oct 25 16:11:37 (FPC Slot 3, PIC Slot 2) {svc_set}[FWNAT]: ASP_SFW_RULE_ACCEPT: proto 6 (TCP) application:
rtsp, ge-2/0/1.0:1.1.1.2:35595 -> 2.2.2.2:554, Match SFW accept rule-set: , rule: allow_rtsp, term: 0
```

For a complete listing of system log messages, see the [System Log Explorer](#).

RELATED DOCUMENTATION

ALG Descriptions

[Configuring Application Sets | 421](#)

Configuring Application Properties

[Examples: Configuring Application Protocols | 437](#)

12

PART

NAT, Stateful Firewall, and IDS Flows

[Inline NAT Services Overview and Configuration](#) | **451**

Inline NAT Services Overview and Configuration

IN THIS CHAPTER

- [Inline Static Source NAT Overview | 451](#)
- [Configuring Inline Static Source NAT44 for Next Gen Services | 452](#)
- [Inline Static Destination NAT Overview | 455](#)
- [Configuring Inline Static Destination NAT for Next Gen Services | 455](#)
- [Inline Twice Static NAT Overview | 459](#)
- [Configuring Inline Twice Static NAT44 for Next Gen Services | 460](#)

Inline Static Source NAT Overview

Inline static source NAT uses the capabilities of the MPC line card to perform address translation, eliminating the need for a services card.

Static source NAT performs a one-to-one static mapping of the original private domain host source address to a public source address. A block of external addresses is set aside for this mapping, and source addresses are translated as hosts in a private domain originate sessions to the external domain. Static source NAT does not perform port mapping. For packets outbound from the private network, static source NAT translates source IP addresses and related fields such as IP, TCP, UDP, and ICMP header checksums. For inbound packets, static source NAT translates the destination IP address and the checksums.

Benefits

- Allows hosts in the private network to connect with the external domain, while hiding the private network.
- Eliminates the need for a services card
- Supports more NAT flows than a services card

Configuring Inline Static Source NAT44 for Next Gen Services

1. [Configuring the Source Pool for Inline Static Source NAT44 | 452](#)
2. [Configuring the NAT Rule for Inline Static Source NAT44 | 453](#)
3. [Configuring the Service Set for Inline Static Source NAT44 | 454](#)
4. [Configuring Inline Services and an Inline Services Interface | 454](#)

Configuring the Source Pool for Inline Static Source NAT44

To configure the source pool for inline static source NAT44:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. Configure a one-to-one static mapping of the original source addresses to the addresses in the source pool by specifying the first address from the matching source-address prefix that is in the source NAT rule.

```
[edit services nat source pool nat-pool-name]  
user@host# set host-address-base ip-address
```

4. To allow the IP addresses of a NAT source pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]  
user@host# set allow-overlapping-pools
```

Configuring the NAT Rule for Inline Static Source NAT44

To configure the NAT source rule for inline static source NAT44:

1. Configure the NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

4. Specify the NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

5. Configure the generation of a syslog when traffic matches the NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Inline Static Source NAT44

To configure the service set for inline static source NAT44:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service set, which requires a single service interface, or a next-hop service set, which requires an inside and outside service interface.

- To configure an interface service set:

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface si-slot-number/pic-number/O.logical-unit-number
```

- To configure a next-hop service set:

```
[edit services service-set service-set-name]
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface vms-slot-number/pic-number/O.logical-unit-number
outside-service-interface si-slot-number/pic-number/O.logical-unit-number
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

Configuring Inline Services and an Inline Services Interface

To enable inline services and an inline services interface:

1. Enable inline services for the FPC and PIC slot, and define the amount of bandwidth to dedicate to inline services.

```
[edit chassis si-fpc slot-number pic number]
user@host# set inline-services bandwidth (1g | 10g | 20g | 30g | 40g | 100g)
```

2. Configure the inline services logical interface or interfaces.

- If you are using an interface service set, configure one logical unit:


```
[edit interfaces si-slot-number/pic-number/0
user@host# set unit logical-unit-number family family
```

- If you are using a next-hop service set, configure two logical units and define the inside and outside interfaces:

```
[edit interfaces si-slot-number/pic-number/0
user@host# set unit logical-unit-number family family
user@host# set unit logical-unit-number service-domain inside
user@host# set unit logical-unit-number family family
user@host# set unit logical-unit-number service-domain outside
```

Inline Static Destination NAT Overview

Inline static destination NAT uses the capabilities of the MPC line card to perform address translation, eliminating the need for a services card.

Static destination NAT translates the IPv4 destination address of an incoming packet to the IPv4 address of a private server. This redirects traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

Static destination NAT uses a one-to-one mapping between the original address and the translated address; the mapping is configured statically.

Benefits

- Allows external traffic to communicate with a private host without revealing the host's private IP address
- Does not require port mapping
- Eliminates the need for a services card
- Supports more NAT flows than a services card

Configuring Inline Static Destination NAT for Next Gen Services

1. [Configuring the Destination Pool for Inline Static Destination NAT | 456](#)
2. [Configuring the NAT Rule for Inline Static Destination NAT | 456](#)

3. [Configuring the Service Set for Inline Static Destination NAT | 458](#)
4. [Configuring Inline Services and an Inline Services Interface | 458](#)

Configuring the Destination Pool for Inline Static Destination NAT

To configure the destination pool for inline static destination NAT:

1. Create a destination pool.

```
user@host# edit services nat destination pool nat-pool-name
```

2. Define the addresses or subnets to which destination addresses are translated.

```
[edit services nat destination pool nat-pool-name]  
user@host# set address address-prefix
```

3. To allow the IP addresses of a NAT destination pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]  
user@host# set allow-overlapping-pools
```

Configuring the NAT Rule for Inline Static Destination NAT

To configure the NAT destination for static destination NAT:

1. Configure the NAT rule name.

```
[edit services destination source]  
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]  
user@host# set match-direction (in | out)
```

3. Specify the source addresses of traffic that the NAT rule applies to.

To specify one address or prefix value:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

4. Specify the destination addresses that the NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

5. Specify the NAT pool that contains the destination addresses for translated traffic.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat pool nat-pool-name
```

6. Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Inline Static Destination NAT

To configure the service set for inline static destination NAT:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service set, which requires a single service interface, or a next-hop service set, which requires an inside and outside service interface.

- To configure an interface service set:

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface si-slot-number/pic-number/O.logical-unit-number
```

- To configure a next-hop service set:

```
[edit services service-set service-set-name]
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface si-slot-number/pic-number/O.logical-unit-number
outside-service-interface si-slot-number/pic-number/O.logical-unit-number
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

Configuring Inline Services and an Inline Services Interface

To enable inline services and an inline services interface:

1. Enable inline services for the FPC and PIC slot, and define the amount of bandwidth to dedicate to inline services.

```
[edit chassis si-fpc slot-number pic number port number]
user@host# set inline-services bandwidth (1g | 10g | 20g | 30g | 40g | 100g)
```

2. Configure the inline services logical interface or interfaces.

- If you are using an interface service set, configure one logical unit:

```
[edit interfaces si-slot-number/pic-number/0
user@host# set unit logical-unit-number family family
```

- If you are using a next-hop service set, configure two logical units and define the inside and outside interfaces:

```
[edit interfaces si-slot-number/pic-number/0
user@host# set unit logical-unit-number family family
user@host# set unit logical-unit-number service-domain inside
user@host# set unit logical-unit-number family family
user@host# set unit logical-unit-number service-domain outside
```

Inline Twice Static NAT Overview

Inline twice static NAT uses the capabilities of the MPC line card to perform address translation, eliminating the need for a services card.

Twice static NAT translates both the source and destination IP addresses. An addresses is translated with a one-to-one static mapping to an address in a pool. Port mapping is not performed.

The original private domain host source address is translated to a public source address.

The destination address is translated to the IPv4 address of a private server. This redirects traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

Benefits

- Allows hosts in the private network to connect with the external domain, while hiding the private network.
- Hides a private network
- Allows external traffic to communicate with a private host without revealing the host's private IP address
- Does not require port mapping
- Eliminates the need for a services card
- Supports more NAT flows than a services card

Configuring Inline Twice Static NAT44 for Next Gen Services

1. [Configuring the Source and Destination Pools for Inline Twice Static NAT44 | 460](#)
2. [Configuring the NAT Rules for Inline Twice Static NAT44 | 461](#)
3. [Configuring the Service Set for Inline Twice Static NAT44 | 463](#)
4. [Configuring Inline Services and an Inline Services Interface | 464](#)

Configuring the Source and Destination Pools for Inline Twice Static NAT44

To configure the source and destination pools for inline twice static NAT44:

1. Create a source pool.

```
user@host# edit services nat source pool nat-pool-name
```

2. Define the addresses or subnets to which source addresses are translated.

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix
```

or

```
[edit services nat source pool nat-pool-name]  
user@host# set address address-prefix to address address-prefix
```

3. Configure a one-to-one static mapping of the original source addresses to the addresses in the source pool by specifying the first address from the matching source-address prefix that is in the source NAT rule.

```
[edit services nat source pool nat-pool-name]  
user@host# set host-address-base ip-address
```

4. Create a destination pool. Do not use the same name that you used for the source pool.

```
user@host# edit services nat destination pool nat-pool-name
```

5. Define the addresses or subnets to which destination addresses are translated.

```
[edit services nat destination pool nat-pool-name]
user@host# set address address-prefix
```

6. To allow the IP addresses of a NAT pool to overlap with IP addresses in pools used in other service sets, configure **allow-overlapping-pools**.

```
[edit services nat]
user@host# set allow-overlapping-pools
```

Configuring the NAT Rules for Inline Twice Static NAT44

To configure the source and destination NAT rules for twice static NAT44:

1. Configure the source NAT rule name.

```
[edit services nat source]
user@host# set rule-set rule-set-name rule rule-name
```

2. Specify the traffic direction to which the source NAT rule set applies.

```
[edit services nat source rule-set rule-set-name]
user@host# set match-direction (in | out)
```

3. Specify the addresses that are translated by the source NAT rule.

To specify one address or prefix value:

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set match source-address-name address-name
```

4. Specify the source NAT pool that contains the addresses for translated traffic.

```
[edit services nat source rule-set rule-set-name rule rule-name]
user@host# set then source-nat pool nat-pool-name
```

5. Configure the generation of a syslog when traffic matches the source NAT rule conditions.

```
[edit services nat source rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

6. Configure the destination NAT rule name.

```
[edit services nat destination]
user@host# set rule-set rule-set-name rule rule-name
```

7. Specify the traffic direction to which the destination NAT rule set applies.

```
[edit services nat destination rule-set rule-set-name]
user@host# set match-direction (in | out | in-out)
```

8. Specify the destination addresses of traffic that the destination NAT rule applies to.

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address address
```

To specify a range of addresses, configure an address book global address with the desired address range, and assign the global address to the NAT rule:

```
[edit services address-book global]
user@host# set address address-name range-address lower-limit to upper-limit
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address-name address-name
```

To specify any unicast address:

```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set match destination-address any-unicast
```

9. Specify the destination NAT pool that contains the destination addresses for translated traffic.


```
[edit services nat destination rule-set rule-set-name rule rule-name]
user@host# set then destination-nat pool nat-pool-name
```

10. Configure the generation of a syslog when traffic matches the destination NAT rule match conditions.

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
user@host# set syslog
```

Configuring the Service Set for Inline Twice Static NAT44

To configure the service set for inline static NAT44:

1. Define the service set.

```
[edit services]
user@host# edit service-set service-set-name
```

2. Configure either an interface service set, which requires a single service interface, or a next-hop service set, which requires an inside and outside service interface.

- To configure an interface service set:

```
[edit services service-set service-set-name]
user@host# set interface-service service-interface si-slot-number/pic-number/0.logical-unit-number
```

- To configure a next-hop service set:

```
[edit services service-set service-set-name]
[edit services service-set service-set-name]
user@host# set next-hop-service inside-service-interface si-slot-number/pic-number/0.logical-unit-number
outside-service-interface vms-slot-number/pic-number/0.logical-unit-number
```

3. Specify the NAT rule sets to be used with the service set.

```
[edit services service-set service-set-name]
user@host# set nat-rule-sets rule-set-name
```

Configuring Inline Services and an Inline Services Interface

To enable inline services and an inline services interface:

1. Enable inline services for the FPC and PIC slot, and define the amount of bandwidth to dedicate to inline services.

```
[edit chassis fpc slot-number pic number]
user@host# set inline-services bandwidth (1g | 10g | 20g | 30g | 40g | 100g)
```

2. Configure the inline services logical interface or interfaces.

- If you are using an interface service set, configure one logical unit:

```
[edit interfaces si-slot-number/pic-number/0]
user@host# set unit logical-unit-number family family
```

- If you are using a next-hop service set, configure two logical units and define the inside and outside interfaces:

```
[edit interfaces si-slot-number/pic-number/0]
user@host# set unit logical-unit-number family family
user@host# set unit logical-unit-number service-domain inside
user@host# set unit logical-unit-number family family
user@host# set unit logical-unit-number service-domain outside
```

13

PART

Configuration Statements

Configuration Statements | **466**

Configuration Statements

IN THIS CHAPTER

- address (Address Book Next Gen Services) | 473
- address (NAT Pool Next Gen Services) | 474
- address-pooling (Source NAT Next Gen Services) | 475
- aggregations (IDS Screen Next Gen Services) | 476
- alarm-without-drop (IDS Screen Next Gen Services) | 477
- white-list | 478
- allow-overlapping-pools (NAT Next Gen Services) | 479
- application (NAT Next Gen Services) | 480
- application-profile (Services CoS Next Gen Services) | 481
- application-protocol | 483
- application-set | 485
- applications (Services ALGs) | 486
- automatic (Source NAT Next Gen Services) | 487
- bad-option (IDS Screen Next Gen Services) | 488
- block-allocation (Source NAT Next Gen Services) | 489
- block-frag (IDS Screen Next Gen Services) | 490
- by-destination (IDS Screen Next Gen Services) | 491
- bypass-traffic-on-exceeding-flow-limits | 493
- by-protocol (IDS Screen Next Gen Services) | 494
- by-source (IDS Screen Next Gen Services) | 496
- category (System Logging) | 498
- child-inactivity-timeout | 500
- clat-prefix (Source NAT Next Gen Services) | 501
- clear-dont-fragment-bit (NAT Next Gen Services) | 502
- close-timeout | 503
- cos-rule-sets (Service Set Next Gen Services) | 504
- cos-rules (Service Set Next Gen Services) | 505
- cpu-load-threshold | 506

- [cpu-throttle \(Next Gen Services\) | 507](#)
- [data \(FTP\) | 508](#)
- [description \(Security Policies Next Gen Services\) | 509](#)
- [destination-address \(NAT Next Gen Services\) | 510](#)
- [destination-address-name \(NAT Next Gen Services\) | 511](#)
- [destination-prefix \(Destination NAT Next Gen Services\) | 512](#)
- [deterministic \(Source NAT Next Gen Services\) | 513](#)
- [deterministic-nat-configuration-log-interval \(Source NAT Next Gen Services\) | 514](#)
- [disable-global-timeout-override | 515](#)
- [dns-filter | 516](#)
- [dns-filter-template | 518](#)
- [drop-member-traffic \(Aggregated Multiservices\) | 521](#)
- [dscp \(Services CoS\) | 522](#)
- [ds-lite | 523](#)
- [ei-mapping-timeout \(Source NAT Next Gen Services\) | 525](#)
- [enable-asymmetric-traffic-processing \(Service Set Next Gen Services\) | 526](#)
- [enable-rejoin \(Aggregated Multiservices\) | 527](#)
- [enable-subscriber-analysis \(Services Options VMS Interfaces\) | 528](#)
- [event-rate \(Next Gen Services Service-Set Local System Logging\) | 529](#)
- [file \(Next Gen Services Global System Logging\) | 530](#)
- [files \(Next Gen Services Global System Logging\) | 531](#)
- [filename \(Next Gen Services Global System Logging\) | 532](#)
- [filtering-type \(Source NAT Next Gen Services\) | 533](#)
- [fin-no-ack \(IDS Screen Next Gen Services\) | 534](#)
- [flag \(Next Gen Services Global System Logging\) | 535](#)
- [format \(Next Gen Services Service-Set Remote System Logging\) | 536](#)
- [forwarding-class \(Services PIC Classifiers\) | 537](#)
- [forwarding-class \(Services PIC Classifiers\) | 538](#)
- [forwarding-class \(Services PIC Classifiers\) | 539](#)
- [fragment \(IDS Screen Next Gen Services\) | 540](#)
- [fragment-limit | 541](#)
- [ftp \(Services CoS Next Gen Services\) | 542](#)
- [gate-timeout | 543](#)
- [global-dns-stats-log-timer | 544](#)

- group (Traffic Load Balancer) | 545
- hash-keys (Interfaces) | 547
- header-integrity-check (Next Gen Services) | 549
- high-availability-options (Aggregated Multiservices) | 551
- host (Next Gen Services Service-Set Remote System Logging) | 552
- host-address-base (Source NAT Next Gen Services) | 553
- inactivity-timeout | 554
- inactivity-asymm-tcp-timeout (Service Set Next Gen Services) | 555
- icmp (IDS Screen Next Gen Services) | 556
- icmp-type | 557
- icmpv6-malformed (IDS Screen Next Gen Services) | 558
- ip (IDS Screen Next Gen Services) | 559
- ipv6-extension-header (IDS Screen Next Gen Services) | 561
- limit-session (IDS Screen Next Gen Services) | 563
- inline-services (PIC level) | 565
- ipv6-extension-header (IDS Screen Next Gen Services) | 566
- instance (Traffic Load Balancer) | 568
- interface-service (Services Interfaces) | 570
- land (IDS Screen Next Gen Services) | 571
- large (IDS Screen Next Gen Services) | 572
- limit-session (IDS Screen Next Gen Services) | 573
- load-balancing-options (Aggregated Multiservices) | 575
- local-category (Next Gen Services Service-Set Local System Logging) | 577
- local-log-tag (Next Gen Services Service-Set System Logging) | 579
- loose-source-route-option (IDS Screen Next Gen Services) | 580
- many-to-one (Aggregated Multiservices) | 581
- map-e | 583
- mapping-timeout (Source NAT Next Gen Services) | 585
- mapping-type (Source NAT Next Gen Services) | 586
- match (Next Gen Services Global System Logging) | 587
- match (Services CoS Next Gen Services) | 588
- match (Stateful Firewall Rule Next Gen Services) | 590
- match-direction (NAT Next Gen Services) | 591
- match-rules-on-reverse-flow (Next Gen Services) | 592

- max-session-setup-rate (Service Set) | 593
- max-sessions-per-subscriber (Service Set Next Gen Services) | 594
- maximum | 595
- member-failure-options (Aggregated Multiservices) | 596
- member-interface (Aggregated Multiservices) | 599
- mode (Next Gen Services Service-Set System Logging) | 601
- name (Next Gen Services Global System Logging) | 602
- nat-options (Next Gen Services) | 603
- nat-rule-sets (Service Set Next Gen Services) | 604
- next-hop-service | 605
- no-remote-trace (Next Gen Services Global System Logging) | 606
- no-translation (Source NAT Next Gen Services) | 607
- no-world-readable (Next Gen Services Global System Logging) | 608
- off (Destination NAT Next Gen Services) | 609
- open-timeout | 610
- pcp-rules | 611
- ping-death (IDS Screen Next Gen Services) | 612
- policy (Services CoS Next Gen Services) | 613
- policy (Stateful Firewall Rules Next Gen Services) | 615
- pool (Destination NAT Next Gen Services) | 616
- pool (Source NAT Next Gen Services) | 617
- pool (NAT Rule Next Gen Services) | 618
- pool-default-port-range (Source NAT Next Gen Services) | 619
- pool-utilization-alarm (Source NAT Next Gen Services) | 620
- port (Source NAT Next Gen Services) | 621
- port-forwarding (Destination NAT Next Gen Services) | 622
- port-forwarding-mappings (Destination NAT Rule Next Gen Services) | 623
- port-round-robin (Source NAT Next Gen Services) | 624
- ports-per-session | 625
- preserve-parity (Source NAT Next Gen Services) | 626
- preserve-range (Source NAT Next Gen Services) | 627
- profile (Traffic Load Balancer) | 628
- profile (Web Filter) | 632
- protocol (Applications) | 635

- range (Source NAT Next Gen Services) | 637
- rate | 638
- real-service (Traffic Load Balancer) | 639
- reassembly-timeout | 640
- record-route-option (IDS Screen Next Gen Services) | 641
- redistribute-all-traffic (Aggregated Multiservices) | 642
- redundancy-event (Services Redundancy Daemon) | 643
- redundancy-options (Aggregated Multiservices) | 645
- redundancy-options (Stateful Synchronization) | 646
- redundancy-policy (Interchassis Services Redundancy) | 648
- redundancy-set | 650
- redundancy-set-id (Service Set) | 652
- rejoin-timeout (Aggregated Multiservices) | 653
- rpc-program-number | 654
- rtlog (Next Gen Services Global System Logging) | 655
- rule (Destination NAT Next Gen Services) | 656
- rule (Services CoS Next Gen Services) | 657
- rule (PCP) | 659
- rule (Source NAT Next Gen Services) | 661
- rule-set (Services CoS Next Gen Services) | 662
- rule-set (Softwires Next Gen Services) | 663
- secure-nat-mapping (Source NAT Next Gen Services) | 664
- security-intelligence | 665
- security-intelligence-policy | 667
- security-option (IDS Screen Next Gen Services) | 668
- server (pcp) | 669
- service-domain | 671
- service-interface (Services Interfaces) | 672
- services-options (Next Gen Services Interfaces) | 673
- service-set (Interfaces) | 677
- service-set (Services) | 678
- service-set-options (Next Gen Services Services) | 682
- session-limit | 683
- session-limit (Service Set Next Gen Services) | 684

- session-timeout (Service Set Next Gen Services) | 685
- severity (Next Gen Services Service-Set Remote System Logging) | 686
- sip (Services CoS Next Gen Services) | 688
- size (Next Gen Services Global System Logging) | 689
- snmp-command | 690
- snmp-trap-thresholds (Next Gen Services) | 691
- softwire-name (Next Gen Services) | 692
- softwires (Next Gen Services) | 694
- softwire-name (Next Gen Services) | 695
- softwire-options | 697
- softwire-types (Next Gen Services) | 698
- softwires-rule-set (Service Set Next Gen Services) | 702
- source-address (Next Gen Services Service-Set Remote System Logging) | 703
- source-address (NAT Next Gen Services) | 704
- source-address-name (NAT Next Gen Services) | 705
- source-port | 706
- source-route-option (IDS Screen Next Gen Services) | 707
- stateful-firewall-rules (Service Set Next Gen Services) | 708
- stateful-firewall-rule-set (Next Gen Services) | 709
- stateful-firewall-rule-sets (Service Set Next Gen Services) | 710
- stream (Next Gen Services Service-Set Remote System Logging) | 711
- stream-option (IDS Screen Next Gen Services) | 712
- strict-source-route-option (IDS Screen Next Gen Services) | 713
- syn-ack-ack-proxy (IDS Screen Next Gen Services) | 714
- syn-fin (IDS Screen Next Gen Services) | 715
- syn-frag (IDS Screen Next Gen Services) | 716
- syslog (Services CoS) | 717
- syslog (Next Gen Services Service-Set System Logging) | 718
- tcp-no-flag (IDS Screen Next Gen Services) | 719
- tcp-session (Service Set Next Gen Services) | 720
- tcp-tickles (Service Set Next Gen Services) | 721
- tear-drop (IDS Screen Next Gen Services) | 722
- then (Services CoS Next Gen Services) | 723
- then (Stateful Firewall Rule Next Gen Services) | 725

- timestamp-option (IDS Screen Next Gen Services) | 726
- traceoptions (Next Gen Services Service-Set Flow) | 727
- traceoptions (Traffic Load Balancer) | 731
- traceoptions (Next Gen Services Global System Logging) | 734
- traceoptions (Next Gen Services Softwires) | 735
- traffic-load-balance (Traffic Load Balancer) | 737
- transport (Next Gen Services Syslog Message Security) | 739
- ttl-threshold | 740
- unknown-protocol (IDS Screen Next Gen Services) | 741
- url-filter | 742
- url-filter-profile | 744
- url-filter-template | 745
- uuid | 747
- v6rd | 748
- video (Application Profile) | 749
- video (Application Profile) | 750
- virtual-service (Traffic Load Balancer) | 751
- voice | 753
- voice (Application Profile) | 754
- web-filter | 755
- web-filter-profile | 757
- winnuke (IDS Screen Next Gen Services) | 758
- world-readable (Next Gen Services Global System Logging) | 759

address (Address Book Next Gen Services)

Syntax

```
address address-name range-address lower-limit to upper-limit
```

Hierarchy Level

```
[edit services address-book global]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a range of addresses that can be referenced in the **match** stanza of a NAT rule.

Options

lower-limit—The lower end of the address range.

upper-limit—The upper end of the address range.

Required Privilege Level

interface—To view this statement in the configuration.

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

address (NAT Pool Next Gen Services)

Syntax

```
address address-prefix | address address-prefix to address address-prefix;
```

Hierarchy Level

```
[edit services nat destination pool nat-pool-name],  
[edit services nat source pool nat-pool-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Define the addresses or subnets to which source addresses or destination addresses are translated. You can configure a single address, an address range, a single subnet, or a subnet range.

Options

address *address-prefix*—A single address or subnet.

address *address-prefix* to address *address-prefix*—An address range or a subnet range.

Required Privilege Level

interface—To view this statement in the configuration.

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

address-pooling (Source NAT Next Gen Services)

Syntax

```
address-pooling {  
    no-paired;  
}
```

Hierarchy Level

```
[edit services nat source pool pool-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Allow address-pooling no-paired for a source pool without port translation

Options

no-paired— Allow address-pooling no-paired for a source pool without port translation.

Required Privilege Level

interface—To view this statement in the configuration.

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

aggregations (IDS Screen Next Gen Services)

Syntax

```
aggregations {
  destination-prefix-ipv6-mask prefix-length;
  destination-prefix-mask prefix-length;
  source-prefix-ipv6-mask prefix-length;
  source-prefix-mask prefix-length;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure intrusion detection service session limits for individual destination subnets or source subnets rather than individual addresses. This applies session limits to an aggregation of all sessions from or to an individual subnet of the specified length.

For example, if you configure a value of 24 for **destination-prefix-mask**, then sessions to 10.1.1.2 and 10.1.1.3 are counted as sessions to the 10.1.1/24 subnet.

Options

destination-prefix-ipv6-mask *prefix-length*—Prefix length for destination IPv6 address subnets.

Range: 0 through 128

destination-prefix-mask *prefix-length*—Prefix length for destination IPv4 address subnets.

Range: 0 through 32

source-prefix-ipv6-mask *prefix-length*—Prefix length for source IPv6 address subnets.

Range: 0 through 128

source-prefix-mask *prefix-length*—Prefix length for source IPv4 address subnets.

Range: 0 through 32

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services | 276](#)**alarm-without-drop (IDS Screen Next Gen Services)****Syntax**

```
alarm-without-drop;
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure the IDS screen to log an alarm for an offending packet, but not drop the packet. The screen skips the rest of the screen checks. The packet is not counted as a dropped packet.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services | 276](#)

white-list

Syntax

```
white-list name {
    destination-address [address];
    source-address [address];
}
```

Hierarchy Level

```
[edit security screen ids-option screen-name tcp syn-flood]
[edit security screen ids-option screen-name udp flood]
[edit tenants tenant-name security screen]
[edit services screen ids-option screen-name limit-session by-destination by-protocol tcp]
[edit services screen ids-option screen-name limit-session by-destination by-protocol udp]
[edit services screen ids-option screen-name limit-session by-source by-protocol tcp]
[edit services screen ids-option screen-name limit-session by-source by-protocol udp]
```

Release Information

Statement introduced in Junos OS Release 12.1.

Support for UDP flood screen allowlist introduced in Junos OS Release 17.4.

tenant option added in Junos OS Release 18.3R1.

Support for UDP and TCP flood screen allowlists added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480 and MX960 routers.

Description

Configure a list of IP addresses that are exempt from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process. This list of exempt addresses is called an allowlist.

You can also use this statement to configure an allowlist of IP addresses that bypass UDP flood detection.

NOTE: This statement is not supported to create UDP flood screen allowlists on SRX5400, SRX5600, and SRX5800 devices

Both IPv4 and IPv6 allowlists are supported. Addresses in the list must be all IPv4 or all IPv6. Each allowlist can have up to 32 IP address prefixes.

Options

- **name**—The name of the allowlist.

- **destination-address *address***—Destination IP address or an address prefix. You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.
- **source-address *address***—Source IP address or an address prefix. You can configure multiple addresses or address prefixes separated by spaces and enclosed in square brackets.

Required Privilege Level

security—To view this statement in the configuration.

security-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Attack Detection and Prevention Overview

Example: Configuring Multiple Screening Options

allow-overlapping-pools (NAT Next Gen Services)

Syntax

```
allow-overlapping-pools;
```

Hierarchy Level

```
[edit services nat]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify that NAT source or destination pools can have IP addresses that overlap with IP addresses in pools used in other service sets. However, pools that configure port-block allocation must not overlap with other pools.

Required Privilege Level

interface—To view this statement in the configuration.

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

application (NAT Next Gen Services)

Syntax

```
application [application-name]
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],  
[edit services nat source rule-set rule-set rule rule-name match]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify one or more application protocols to which the NAT rule applies. The number of applications must not exceed 3072.

Required Privilege Level

interface—To view this statement in the configuration.

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

application-profile (Services CoS Next Gen Services)

Syntax

```

application-profile name {
  ftp {
    data {
      dscp dscp;
      forwarding-class forwarding-class;
    }
  }
  sip {
    video {
      dscp dscp;
      forwarding-class forwarding-class;
    }
    voice {
      dscp dscp;
      forwarding-class forwarding-class;
    }
  }
}

```

Hierarchy Level

[edit services cos]

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure CoS actions for FTP and SIP traffic. The application profile can then be used in CoS rule actions. This enables you to apply a certain DSCP, or forwarding-class to a set of L7 flows.

Options

profile-name—Name of the application profile.

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

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Class of Service for Services PICs \(Next Gen Services\)](#) | 255

application-protocol

Syntax

```
application-protocol protocol-name;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

login options introduced in Junos OS Release 7.4.

ip option introduced in Junos OS Release 8.2.

ike-esp-nat option introduced in Junos OS Release 17.1.

ras option introduced in Junos OS Release 17.1.

Description

Identify the application protocol name. Application protocols are also called application layer gateways (ALGs).

Options

protocol-name—Name of the protocol. The following protocols are supported:

bootp—Bootstrap protocol

dce-rpc—DCE RPC

dce-rpc-portmap—DCE RPC portmap

dns—Domain Name Service

exec—Remote Execution Protocol

ftp—File Transfer Protocol

h323—H.323

icmp—ICMP

iiop—Internet Inter-ORB Protocol

ike-esp-nat—IKE ALG

ip—IP

login—Login

netbios—NetBIOS

netshow—NetShow

pptp—Point-to-Point Tunneling Protocol

ras—Gatekeeper RAS for H323

realaudio—RealAudio

rpc—RPC

rpc-portmap—RPC portmap

rtsp—Real Time Streaming Protocol

shell—Shell

sip—Session Initiation Protocol

snmp—SNMP

sqlnet—SQLNet

talk—Talk Program

tftp—Trivial File Transfer Protocol

traceroute—Traceroute

winframe—WinFrame

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

ALG Descriptions

[Configuring Application Sets | 421](#)

Configuring Application Properties

[Examples: Configuring Application Protocols | 437](#)

[Verifying the Output of ALG Sessions | 438](#)

application-set

Syntax

```
application-set application-set-name {  
    application application-name;  
}
```

Hierarchy Level

[edit [applications](#)]

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Configure one or more applications to include in an application set.

Options

application-set-name—Identifier of an application set.

Required Privilege Level

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

ALG Descriptions

[Configuring Application Sets](#) | 421

Configuring Application Properties

[Examples: Configuring Application Protocols](#) | 437

[Verifying the Output of ALG Sessions](#) | 438

applications (Services ALGs)

Syntax

```
applications { ... }
```

Hierarchy Level

```
[edit]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Define the applications used in services.

Required Privilege Level

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

RELATED DOCUMENTATION

<i>ALG Descriptions</i>
Configuring Application Sets 421
<i>Configuring Application Properties</i>
Examples: Configuring Application Protocols 437
Verifying the Output of ALG Sessions 438

automatic (Source NAT Next Gen Services)

Syntax

```
automatic (random-allocation | round-robin);
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure automatic port assignment for source NAT with port translation, except for deterministic NAT. Automatic port assignment uses the port range 1024 through 65535. Specify either random allocation or round-robin allocation. Random allocation randomly assigns a port from the range 1024 through 65535 for each port translation. Round robin allocation first assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

Options

random-allocation—Randomly assigns a port from the range 1024 through 65535 for each port translation.

round-robin—First assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

Required Privilege Level

interface—To view this statement in the configuration.

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

bad-option (IDS Screen Next Gen Services)

Syntax

```
bad-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop any packet with incorrectly formatted IPv4 options or IPv6 extension headers. Incorrectly formatted IPv4 options or IPv6 extension headers can cause unpredictable issues, depending on the IP stack implementation of routers and the target.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

block-allocation (Source NAT Next Gen Services)

Syntax

```
block-allocation {
  active-block-timeout timeout-interval;
  block-size block-size;
  interim-logging-interval timeout-interval;
  maximum-blocks-per-host maximum-block-number
  log disable | enable
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Allocate a block of ports for each subscriber to use for source NAT with port translation, except for deterministic NAT. New requests for NAT ports for the subscriber are served from the active block. With port block allocation, we generate one syslog log per set of ports allocated for a subscriber. This reduces the number of logs, making it easier to track subscribers.

Options

active-block-timeout *timeout-interval*—The interval, in seconds, for which the block is active. After the timeout, a new block is allocated, even if ports are available in the active block. If you set the timeout to 0, port blocks are filled completely before a new port block is allocated, and the last port block remains active indefinitely.

Range: 0 through 86,400

Default: 0

block-size *block-size*—Number of ports in a block.

Range: 1 through 64,512

Default: 128

interim-logging-interval *timeout-interval*—The interval, in seconds, at which to send interim system logs for active port blocks and for inactive port blocks with live sessions. This increases the reliability of system logs, which are UDP-based and can get lost in the network.

Range: 1800 through 86,400

Default: 0 (interim logs are disabled)

maximum-blocks-per-host *maximum-block-number*—The maximum number of blocks that can be allocated to a subscriber address.

Range: 1 through 512

Default: 8

log disable—Disable logs for port block allocation. Logs are enabled by default.

Required Privilege Level

interface—To view this statement in the configuration.

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

block-frag (IDS Screen Next Gen Services)

Syntax

```
block-frag;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop fragmented IP packets. IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the target receives these packets, the results can range from processing the packets incorrectly to crashing the entire system.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

by-destination (IDS Screen Next Gen Services)

Syntax

```
by-destination {
  by-protocol {
    icmp {
      maximum-sessions number;
      packets-rate number;
      session-rate number;
    }
    tcp {
      maximum-sessions number;
      packet-rate number;
      session-rate number;
    }
    udp {
      maximum-sessions number;
      packet-rate number;
      session-rate number;
    }
  }
  maximum-sessions number;
  packet-rate number;
  session-rate number;
  ;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name limit-session]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure session limits for individual destination addresses or for individual destination subnets. This protects against network probing attacks and network flooding attacks. You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol. When a session limit is exceeded for a destination, packets to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination subnets rather than individual addresses, include the **aggregations** statement at the `[edit services screen ids-option screen-name]` hierarchy level.

Options

maximum-sessions *number*—Specify the maximum number of concurrent sessions allowed for an individual destination address or subnet.

packet-rate *number*—Specify the maximum number of packets per second allowed for an individual destination address or subnet.

session-rate *number*—Specify the maximum number of connections per second allowed for an individual destination address or subnet.

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 Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

bypass-traffic-on-exceeding-flow-limits

Syntax

```
bypass-traffic-on-exceeding-flow-limits;
```

Hierarchy Level

```
[edit] services service-set service-set-name service-set-options
```

Release Information

Statement introduced in Junos OS Release 10.1.

Statement introduced in Junos OS Release 19.3R2 on MX240, MX480 and MX960 routers using the MX-SPC3 services card..

Description

[bypass-traffic-on-exceeding-flow-limits](#)[bypass-traffic-on-exceeding-flow-limits](#)[bypass-traffic-on-exceeding-flow-limits](#)

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Service Sets to be Applied to Services Interfaces*

by-protocol (IDS Screen Next Gen Services)

Syntax

```
by-protocol {
  icmp {
    maximum-sessions number;
    packet-rate number;
    session-rate number;
  }
  tcp {
    maximum-sessions number;
    packet-rate number;
    session-rate number;
  }
  udp {
    maximum-sessions number;
    packet-rate number;
    session-rate number;
  }
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name limit-session by-destination],
[edit services screen ids-option screen-name limit-session by-source]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure session limits for individual destination or source addresses, or for individual destination or source subnets, for the specified protocol. This protects against network probing attacks and network flooding attacks. When a session limit is exceeded for a source or destination for the protocol, packets from the source or to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination or source subnets rather than individual addresses, include the **aggregations** statement at the **[edit services screen ids-option *screen-name*]** hierarchy level.

Options

icmp—Apply session limits to ICMP packets.

maximum-sessions *number*—Specify the maximum number of concurrent ICMP sessions allowed for individual destination or source addresses, or for individual destination or source subnets.

packet-rate *number*—Specify the maximum number of ICMP packets per second allowed for individual destination or source addresses, or for individual destination or source subnets.

session-rate *number*—Specify the maximum number of ICMP connections per second allowed for individual destination or source addresses, or for individual destination or source subnets.

tcp—Apply session limits to TCP packets.

maximum-sessions *number*—Specify the maximum number of concurrent TCP sessions allowed for individual destination or source addresses, or for individual destination or source subnets.

packet-rate *number*—Specify the maximum number of TCP packets per second allowed for individual destination or source addresses, or for individual destination or source subnets.

session-rate *number*—Specify the maximum number of TCP connections per second allowed for individual destination or source addresses, or for individual destination or source subnets.

udp—Apply session limits to UDP packets.

maximum-sessions *number*—Specify the maximum number of concurrent UDP sessions allowed for individual destination or source addresses, or for individual destination or source subnets.

packet-rate *number*—Specify the maximum number of UDP packets per second allowed for individual destination or source addresses, or for individual destination or source subnets.

session-rate *number*—Specify the maximum number of UDP connections per second allowed for individual destination or source addresses, or for individual destination or source subnets.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

by-source (IDS Screen Next Gen Services)

Syntax

```
by-source {
  by-protocol {
    icmp {
      maximum-sessions number;
      packet-rate number;
      session-rate number;
    }
    tcp {
      maximum-sessions number;
      packet-rate number;
      session-rate number;
    }
    udp {
      maximum-sessions number;
      packet-rate number;
      session-rate number;
    }
  }
  maximum-sessions number;
  packet-rate number;
  session-rate number;
  ;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name limit-session]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure session limits for individual source addresses or for individual source subnets. This protects against network probing attacks and network flooding attacks. You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol. When a session limit is exceeded for a source, packets from the source are dropped until the session limit is no longer exceeded.

To specify limits for source subnets rather than individual addresses, include the **aggregations** statement at the **[edit services screen ids-option *screen-name*]** hierarchy level.

Options

maximum-sessions *number*—Specify the maximum number of concurrent sessions allowed for an individual source address or subnet.

packet-rate *number*—Specify the maximum number of packets per second allowed for an individual source address or subnet.

session-rate *number*—Specify the maximum number of connections per second allowed for an individual source address or subnet.

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 Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

category (System Logging)

Syntax

```
category category, category....category;
```

Hierarchy Level

```
[edit services service-set service-set-name syslog stream]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services on the MX-SPC3 services card.

Description

Specify the categories for which you want to collect logs.

Options

all—All events are logged

content-security—Content security events are logged

fw-auth—Fw-auth events are logged

screen —Screen events are logged

alg—ALG events are logged

nat—NAT events are logged

flow—Flow events are logged

sctp—Sctp events are logged

gtp—Gtp events are logged

ipsec—Ipsec events are logged

idp—Idp events are logged

rtlog—Rtlog events are logged

pst-ds-lite—Pst-ds-lite events are logged

appqos—Appqos events are logged

secintel—Secintel events are logged

aamw—AAMW events are logged

sfw—Stateful Firewall events are logged

session —Session open and close events are logged

session-open—Session open events are logged

session-close—Session close events are logged

urlf—DNS request filtering events are logged

ha—Stateful High-Availability open and close events are logged

ha-open—Stateful High-Availability open events are logged

ha-close—Stateful High-Availability close events are logged

pcp—PCP logs

Required Privilege Level

system—To view this statement in the configuration.

child-inactivity-timeout

Syntax

```
child-inactivity-timeout seconds;
```

Hierarchy Level

```
[edit applications application ike-esp-nat]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

For an IKE ALG application, configure the ESP session (IPsec data traffic) idle timeout. If no IPsec data traffic is passed on the ESP session in this time, the session is deleted.

The IKE ALG enables the passing of IKEv1 and IPsec packets through NAPT-44 and NAT64 rules between IPsec peers that are not NAT-T compliant.

Options

seconds—Number of seconds.

Default: 800 seconds

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

ALG Descriptions

[Configuring Application Sets | 421](#)

Configuring Application Properties

clat-prefix (Source NAT Next Gen Services)

Syntax

```
clat-prefix clat-prefix;
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the customer-side translator (CLAT) IPv6 source prefix, which is used for 464XLAT.

464XLAT lets an IPv4 client with a private IP address connect to an IPv4 host over an IPv6 network. The CLAT translates IPv4 source addresses to IPv6 by embedding the IPv4 source address in this IPv6 source prefix. The CLAT then sends the packets over an IPv6 network to the MX Series router, which acts as a provider-side translator (PLAT). The MX translates the embedded IPv4 private IP address to a public IPv4 address.

Required Privilege Level

interface—To view this statement in the configuration.

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

clear-dont-fragment-bit (NAT Next Gen Services)

Syntax

```
set clear-dont-fragment-bit;
```

Hierarchy Level

```
[edit services nat natv6v4]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify that the don't fragment (DF) bit for IPv4 packet headers is cleared when the packet length is less than 1280 bytes. Use this statement when configuring stateful NAT64, deterministic NAPT64, and 464XLAT. This prevents unnecessary creation of an IPv6 fragmentation header when translating IPv4 packets that are less than 1280 bytes.

Required Privilege Level

interface—To view this statement in the configuration.

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

close-timeout

Syntax

```
close-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]  
[edit services service-set service-set-name service-set-options tcp-session]
```

Release Information

Statement introduced in Junos OS Release 12.3.

Support for Next Gen Services added in Junos OS Release 19.3R2 on MX Series MX240, MX480 and MX960 using MX-SPC3 services card.

Description

Configure the timeout period for Transmission Control Protocol (TCP) session tear-down.

Options

seconds—Timeout period.

Default: 1 second

Range: 2 through 300 seconds

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Default Timeout Settings for Services Interfaces*

cos-rule-sets (Service Set Next Gen Services)

Syntax

```
cos-rule-sets [cos-rule-set-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the services CoS rule set to apply to the service set. The service set processes the rules in the order they appear in the rule set.

The service set that the CoS rule set is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

cos-rule-set-name—Name of the services CoS rule set.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Class of Service for Services PICs \(Next Gen Services\)](#) | 255

cos-rules (Service Set Next Gen Services)

Syntax

```
cos-rules [cos-rule-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the CoS rules to apply to the service set. You can configure multiple rules.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

cos-rule-name—CoS rule name.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Class of Service for Services PICs \(Next Gen Services\) | 255](#)

cpu-load-threshold

Syntax

```
cpu-load-threshold percentage;
```

Hierarchy Level

```
[edit interfaces interface-name services-options session-limit]
```

Release Information

Statement introduced in Release 13.2.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Regulate the usage of CPU resources on services cards. When the CPU usage exceeds the configured value (percentage of the total available CPU resources), the system reduces the rate of new sessions so that the existing sessions are not affected by low CPU availability. The CPU utilization is constantly monitored, and if the CPU usage remains above the configured **cpu-load-threshold** value for a continuous period of 5 seconds, Junos OS reduces the session rate value configured at **edit interfaces *interface-name* services-options session-limit *rate*** by 10%. This is repeated until the CPU utilization comes down to the configured limit.

Options

percentage—Percentage of total available CPU resources.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

cpu-throttle (Next Gen Services)

Syntax

```
cpu-throttle {
  percentage percent;
}
```

Hierarchy Level

```
[edit services screen]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the services card CPU utilization percentage that triggers the installation of a dynamic filter on the PFEs of the line cards for suspicious activity. The dynamic filter drops the suspicious traffic.

In addition to this threshold, at least one of the following conditions is required to trigger the installation of a dynamic filter:

- The packet rate from an individual source address or to an individual destination address must exceed four times the configured **packet-rate** at the **[edit services screen ids-option *screen-name* limit-session by-source]** or **[edit services screen ids-option *screen-name* limit-session by-destination]** hierarchy level.
- The connection rate from an individual source address or to an individual destination address must exceed four times the configured **session-rate** at the **[edit services screen ids-option *screen-name* limit-session by-source]** or **[edit services screen ids-option *screen-name* limit-session by-destination]** hierarchy level.

Dynamic filters are not created from IDS screens that use subnet aggregation.

The dynamic filter drops the suspicious traffic at the PFE, without the traffic being processed by the IDS screen. When the packet or connection rate no longer exceeds four times the limit in the IDS screen, the dynamic filter is removed.

Options

percentage *percent*—The CPU utilization percentage.

Range: 1 through 100

Default: 90

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

data (FTP)

Syntax

```
data {
  dscp (alias | bits);
  forwarding-class class-name;
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name ftp]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Set the appropriate **dscp** and **forwarding-class** value for FTP data.

Default

By default, the system will not alter the DSCP or forwarding class for FTP data traffic.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring CoS Rules on Services PICs](#)

[video \(Application Profile\)](#) | 749

[voice \(Application Profile\)](#) | 754

description (Security Policies Next Gen Services)

Syntax

```
description description;
```

Hierarchy Level

```
[edit security ike policy policy-name],  
[edit security ike proposal proposal-name],  
[edit security ipsec policy policy-name],  
[edit security ipsec proposal proposal-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Enter descriptive text for an IKE policy, an IPsec policy, an IKE proposal, or an IPsec proposal.

Options

description—Descriptive text.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

destination-address (NAT Next Gen Services)

Syntax

```
destination-address (address | any | any-ipv4 | any-ipv6);
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],  
[edit services nat source rule-set rule-set rule rule-name match]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the destination address that the packet must match for the NAT rule to take effect.

Options

address—A specific address that must be matched.

any—Any unicast destination address results in a match.

any-ipv4—Any IPv4 destination address results in a match.

any-ipv6—Any IPv6 destination address results in a match.

Required Privilege Level

interface—To view this statement in the configuration.

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

destination-address-name (NAT Next Gen Services)

Syntax

```
destination-address-name address-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],  
[edit services nat source rule-set rule-set rule rule-name match]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the name of the range of destination addresses that the packet must match for the NAT rule to take effect. The range of addresses is configured with the **address** statement at the **[edit services address-book global]** hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

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

destination-prefix (Destination NAT Next Gen Services)

Syntax

```
destination-prefix destination-prefix;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name then destination-nat]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the IPv6 prefix that is used to embed an IPv4 destination address in an IPv6 address. The **destination-prefix** statement is used in Stateful NAT64 and 464XLAT translations.

Required Privilege Level

interface—To view this statement in the configuration.

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

deterministic (Source NAT Next Gen Services)

Syntax

```
deterministic {
  block-size block-size;
  host {
    address address;
  }
  include-boundary-addresses;
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure deterministic NAT to ensure that the original internal source IPv4 or IPv6 address and port always map to the same post-NAT IPv4 address and block of ports. In addition, the reverse mapping of a given translated external IPv4 address and port are always mapped to the same internal IP address.

This eliminates the need for address translation logging.

Options

block-size *block-size*—The number of ports in the port block.

Range: 1 to 64,512

Default: 256

host address *address*—The first usable pre-NAT subscriber address, which is used to perform the deterministic NAT mapping.

include-boundary-addresses—Include the translation of the lowest and highest IPv4 addresses (the network and broadcast addresses) in the source address range of a NAT rule. This does not apply to IPv6 source addresses.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Deterministic NAT for Next Gen Services | 114](#)**deterministic-nat-configuration-log-interval (Source NAT Next Gen Services)****Syntax**

```
deterministic-nat-configuration-log-interval seconds;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure the interval at which the syslog is generated for the deterministic NAT configuration.

Options

deterministic-nat-configuration-log-interval *seconds*—Number of seconds in the interval.

Range: 1800 through 86400

Default: 1800

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Deterministic NAT for Next Gen Services | 114](#)

disable-global-timeout-override

Syntax

```
disable-global-timeout-override;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]  
[services service-set service-set-name service-set-options]
```

Release Information

Statement introduced in Junos OS Release 10.0.

Support added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

Description

Disallow overriding a global inactivity or session timeout.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Defining an Application Identification*

dns-filter

Syntax

```
dns-filter {
  database-file filename;
  dns-resp-ttl seconds;
  dns-server [ ip-address ];
  hash-key key-string;
  hash-method hash-method-name;
  statistics-log-timer minutes;
  wilddcarding-level level;
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name],
[edit services web-filter profile profile-name dns-filter-template template-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added for Next Gen Services on MX Series routers MX240, MX480 and MX960 with MX-SPC3 services cards in Junos OS Release 19.3R2.

Description

Configure the settings for filtering DNS requests for disallowed website domains. Filtering can result in either:

- Blocking access to the site by sending the client a DNS response that includes an IP address or domain name of a sinkhole server instead of the disallowed domain.
- Logging the DNS request and allowing access.

Settings at the **[edit services web-filter profile *profile-name* dns-filter-template *template-name*]** hierarchy level override the corresponding settings at the **[edit services web-filter profile *profile-name*]** hierarchy level.

Options

database-file *filename*—Name of the domain filter database file to use when filtering DNS requests.

dns-resp-ttl *seconds*— Number of seconds to live while sending the DNS response after taking the DNS sinkhole action.

Default: 1800

Range: 0 through 86,400

dns-server [*ip-address*]—(Optional) IP addresses (IPv4 or IPv6) for up to three specific DNS servers. DNS filtering examines only DNS requests that are destined for those DNS servers.

hash-key *key-string*—Hash key that you used to create the hashed domain name in the domain filter database file.

hash-method *hash-method-name*—Hash method that you used to create the hashed domain name in the domain filter database file. The only supported hash method is **hmac-sha2-256**.

statistics-log-timer *minutes*—Number of minutes in the interval for logging statistics for DNS requests and for sinkhole actions performed for each customer IP address.

Default: 5

Range: 0 through 60

wildcarding-level *level*—Level of subdomains that are searched for a match. A value of 0 indicates that subdomains are not searched.

For example, if you set the **wildcarding-level** to 4 and the database file includes an entry for **example.com**, the following comparisons are made for a DNS request that arrives with the domain **198.51.100.0.example.com**:

- **198.51.100.0.example.com**: no match
- **51.100.0.example.com**: no match for one level down
- **100.0.example.com**: no match for two levels down
- **0.example.com**: no match for three levels down
- **example.com**: match for four levels down

Range: 0 through 10

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [DNS Request Filtering for Disallowed Website Domains](#) | 311

dns-filter-template

Syntax

```
dns-filter-template template-name {
  client-interfaces [ client-interface-name ];
  client-routing-instance client-routing-instance-name;
  dns-filter {
    database-file filename;
    dns-resp-ttl seconds;
    dns-server [ ip-address ];
    hash-key key-string;
    hash-method hash-method-name;
    statistics-log-timer minutes;
    wildcarding-level level;
  }
  server-interfaces [ server-interface-name ];
  server-routing-instance server-routing-instance-name;
  term term-name {
    from {
      src-ip-prefix [ source-prefix ];
    }
    then {
      accept;
      dns-sinkhole;
    }
  }
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure filtering of DNS requests for disallowed website domains for requests on specific uplink and downlink logical interfaces or routing instances, or for requests from specific source IP address prefixes. The DNS filter template overrides the corresponding settings at the DNS profile level. You can configure up to 32 DNS filter templates in a profile.

Filtering can result in either:

- Blocking access to the site by sending the client a DNS response that includes an IP address or domain name of a sinkhole server instead of the disallowed domain.
- Logging the DNS request and allowing access.

Options

accept—Accept DNS requests for DNS filtering.

client-interfaces [*client-interface-name*]—(Optional) Client-facing (uplink) logical interfaces on which the DNS filter template settings are applied.

client-routing-instance *client-routing-instance-name*—(Optional) Client-facing (uplink) routing instance on which the DNS filter template settings are applied.

dns-filter-template *template-name*—Name of the DNS filter template.

dns-sinkhole—Perform the sinkhole action identified in the domain filter database for disallowed DNS requests.

server-interfaces [*server-interface-name*]—(Optional) Server-facing logical interfaces (downlink) on which the DNS filter template settings are applied.

server-routing-instance *server-routing-instance-name*—(Optional) Server-facing (downlink) routing instance on which the DNS filter template settings are applied.

NOTE: If you configure the client and server interfaces or the client and server routing instances, implicit filters are installed on the interfaces or routing instances to direct DNS traffic to the MS-MPC for DNS filtering. If you configure neither the client and server interfaces nor the routing instances, you must provide a way to direct DNS traffic to the MS-MPC (for example, via routes).

src-ip-prefix [*source-prefix*]—(Optional) Source IP address prefixes of DNS requests you want to filter.

You can configure a maximum of 64 prefixes in a term. If you do not specify any source prefixes, then all DNS requests are filtered.

term *term-name*—Name for a term. You can configure a maximum of 64 terms in a template.

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

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [DNS Request Filtering for Disallowed Website Domains](#) | 311

drop-member-traffic (Aggregated Multiservices)

Syntax

```
drop-member-traffic {
  rejoin-timeout rejoin-timeout;
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify whether the broadband gateway should drop traffic to a services PIC when it fails.

For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration is valid only when two or more services PICs have failed.

The remaining statement is explained separately. See [CLI Explorer](#).

Default

If this statement is not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[member-failure-options \(Aggregated Multiservices\)](#) | 596

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

dscp (Services CoS)

Syntax

```
dscp (alias | bits);
```

Hierarchy Level

```
[edit services cos application-profile profile-name (ftp | sip) (data | video | voice)],  
[edit services cos rule rule-name term term-name then],  
[edit services cos rule rule-name term term-name then reverse]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Define the Differentiated Services code point (DSCP) mapping that is applied to the packets. Change the DSCP (or TOS) on the packet to the specified value. Any conformant bit string can be specified, but only the default alias can be used.

Options

alias—Name assigned to a set of CoS markers.

bits—Mapping value in the packet header.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Actions in CoS Rules.

Configuring CoS Rules on Services PICs

ds-lite

Syntax

```
ds-lite ds-lite-software-concentrator {
  auto-update-mtu;
  flow-limit flow-limit | session-limit-per-prefix session-limit-per-prefix;
  mtu-v6 bytes;
  software-address software-address;
}
```

Hierarchy Level

```
[edit services software software-concentrator]
[edit services softwares software-types]
```

Release Information

Statement introduced in Junos OS Release 10.4.

auto-update-mtu option introduced in Junos OS Release 10.4.

copy-dscp option introduced in Junos OS Release 11.2.

mtu-v6 option introduced in Junos OS Release 10.4.

software-address option introduced in Junos OS Release 10.4.

Support for DS-Lite at the **[edit services softwares **software-types**]** added in Junos OS release 20.2R1 for Next Gen Services on MX240, MX480 and MX960 routers.

Description

Configure settings for a DS-Lite concentrator used to process IPv4 packets encapsulated in IPv6.

The **ds-lite** statement is supported on MX Series routers with MS-DPCs and on M Series routers with MS-100, MS-400, and MS-500 line Multiservices PICs. Starting in Junos OS release 17.4R1, DS-Lite is supported on MX Series routers with MS-MPCs and MS-MICs.

Options

bytes—Maximum transmission unit (MTU), in bytes, for encapsulating IPv4 packets into IPv6. If the final length is greater than the configured value, the IPv6 packet is fragmented. This option is supported on MX Series routers equipped with MS-DPCs. Starting in Junos OS release 18.1R1, this option is also supported on MX Series routers with MS-MPCs or MS-MICs.

ds-lite-software-concentrator—Name applied to a DS-Lite software concentrator.

auto-update-mtu—This option is not currently supported.

copy-dscp—Copy DSCP information to IPv4 headers during decapsulation.

flow-limit—Maximum number of IPv4 flows per software.

Range: 0 through 16384 flows

Range: 0 through 9192 bytes

session-limit-per-prefix—Maximum number of sessions per B4 subnet prefix. This option is supported on MX Series routers equipped with MS-DPCs. Starting in Junos OS Release 18.2R1, this option is also supported on MS-MPCs and MS-MICs.

Range: 0 through 16384 sessions

software-address—Address of the DS-Lite software concentrator.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring a DS-Lite Software Concentrator*

ei-mapping-timeout (Source NAT Next Gen Services)

Syntax

```
ei-mapping-timeout ei-mapping-timeout;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the timeout period for endpoint independent translations that use the NAT pool. Mappings that are inactive for this amount of time are dropped.

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

Options

ei-mapping-timeout *ei-mapping-timeout*—The timeout period in seconds.

Range: 120 through 86,400

Default: 300 (timeout period for endpoint independent translations is set by **mapping-timeout** value at the **[edit services nat source pool *nat-pool-name*]** hierarchy level)

Required Privilege Level

interface—To view this statement in the configuration.

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

enable-asymmetric-traffic-processing (Service Set Next Gen Services)

Syntax

```
enable-asymmetric-traffic-processing;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Enable the service set to handle unidirectional traffic.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

enable-rejoin (Aggregated Multiservices)

Syntax

```
enable-rejoin;
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options redistribute-all-traffic]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Enable the failed member to rejoin the aggregated Multiservices (AMS) interface after the member comes back online.

For many-to-one (N:1) high availability (HA) for service applications like Network Address Translation (NAT), this configuration allows the failed members to rejoin the pool of active members automatically.

Default

If you do not configure this option, then the failed members do not automatically rejoin the **ams** interface even after coming back online.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[redistribute-all-traffic \(Aggregated Multiservices\)](#) | 642

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

enable-subscriber-analysis (Services Options VMS Interfaces)

Syntax

```
enable-subscriber-analysis;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Enable the creation of subscribers if the following are not configured, but you want subscribers to be created:

- NAT
- The **max-sessions-per-subscriber** statement at the **[edit services service-set *service-set-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

[How to Configure Services Interfaces for Next Gen Services](#) | 60

event-rate (Next Gen Services Service-Set Local System Logging)

Syntax

```
event-rate rate-per-second;
```

Hierarchy Level

```
[edit services services-set name syslog]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Rate at which log messages are sent per second to the local file.

Options

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

file (Next Gen Services Global System Logging)

Syntax

```
file <filename> <files files> <match match> <size size> <(world-readable | no-world-readable)>;
```

Hierarchy Level

```
[edit services rtlog traceoptions]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Trace file information

Options

All other options are explained separately.

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

files (Next Gen Services Global System Logging)

Syntax

```
files files;
```

Hierarchy Level

```
[edit services rtlog traceoptions file filename]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Maximum number of trace files

Options

files—Maximum number of trace files

Default: 3

Range: 2 through 1000

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging](#) | 71

[Enabling Global System Logging for Next Gen Services](#) | 73

[Configuring System Logging to One or More Remote Servers for Next Gen Services](#) | 76

[Configuring Local System Logging for Next Gen Services](#) | 74

filename (Next Gen Services Global System Logging)

Syntax

```
filename;
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Name of file in which to write trace information

Options

filename—Name of file in which to write trace information

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

filtering-type (Source NAT Next Gen Services)

Syntax

```
filtering-type {  
  endpoint-independent {  
    prefix-list [allowed-host] except [denied-host];  
  }  
}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify prefix lists that contain prefixes of hosts that are allowed to establish inbound connections using endpoint-independent mapping, and prefix lists for hosts that are not allowed to establish inbound connections. (Prefix lists are configured at the **[edit policy-options]** hierarchy level.)

Options

[allowed-host]—Names of the prefix lists for hosts that are allowed to establish connections.

except [denied-host]—Names of prefix lists for hosts that are not allowed to establish connections.

Required Privilege Level

interface—To view this statement in the configuration.

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

fin-no-ack (IDS Screen Next Gen Services)

Syntax

```
fin-no-ack;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop any packet with a FIN flag set and without the ACK flag set. The TPC FIN No Ack attack can allow the attacker to identify the operating system of the target or to identify open ports on the target.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

flag (Next Gen Services Global System Logging)

Syntax

```
flag name;
```

Hierarchy Level

```
[edit services rtlog traceoptions]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

List of things to include in trace.

Options

name—

Values:

- all—Enable all interface trace flags. event —Trace interface events.
- cache—Enable interface flags for Web filtering cache maintained on the routing table.
- enhanced—Enable interface flags for processing through Enhanced Web Filtering.
- ipc—Trace interface IPC messages.
- media—Trace interface media changes.
- critical—Trace critical events.
- major—Trace major events

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

format (Next Gen Services Service-Set Remote System Logging)

Syntax

```
format format;
```

Hierarchy Level

```
edit services service-set name syslog stream stream-name
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services MX-SPC3 services card.

Description

Specify the file format for the log messages being sent to the remote server.

Options

The file format can be one of the following:

binary—Binary syslog defined by Juniper Networks. Requires Juniper Networks decoders on the server side to decode the logs.

sd-syslog—Structured syslog (defined by RFC5424)

syslog—Traditional syslog (defined by RFC5424)

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

forwarding-class (Services PIC Classifiers)

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit services cos application-profile profile-name (ftp | sip) (data | video | voice)],  
[edit services cos rule rule-name term term-name then],  
[edit services cos rule rule-name term term-name then reflexive; | revert; | reverse {}]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Define the forwarding class to which packets are assigned.

Options

class-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring CoS Rules on Services PICs*

forwarding-class (Services PIC Classifiers)

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit services cos application-profile profile-name (ftp | sip) (data | video | voice)],  
[edit services cos rule rule-name term term-name then],  
[edit services cos rule rule-name term term-name then reverse]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Assign the packets to the specified forwarding class.

Options

class-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Actions in CoS Rules.*

forwarding-class (Services PIC Classifiers)

Syntax

```
forwarding-class class-name;
```

Hierarchy Level

```
[edit services cos application-profile profile-name (ftp | sip) (data | video | voice)],  
[edit services cos rule rule-name term term-name then],  
[edit services cos rule rule-name term term-name then reflexive; | revert; | reverse {}]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Define the forwarding class to which packets are assigned.

Options

class-name—Name of the target application.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring CoS Rules on Services PICs*

fragment (IDS Screen Next Gen Services)

Syntax

```
fragment;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop ICMP packets that are IP fragments. These are considered suspicious packets because ICMP packets are usually short. When the target receives these packets, the results can range from processing packets incorrectly to crashing the entire system.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

fragment-limit

Syntax

```
fragment-limit number-of-fragments;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]  
[edit security flow]
```

Release Information

Statement introduced in Junos OS Release 12.1.

Statement added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

Description

Configure the maximum number of fragments permitted in a packet before the packet is dropped.

Options

number-of-fragments—Maximum number of fragments permitted.

Range: 1 to 250 fragments.

Default: 250 fragments.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces*

ftp (Services CoS Next Gen Services)

Syntax

```
ftp {
  data {
    dscp (alias | bits);
    forwarding-class class-name;
  }
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure CoS actions for FTP traffic in an application profile. The application profile can then be used in CoS rule actions.

Options

dscp (*alias* | *bits*)—Either a code point alias or a DSCP bit value to apply to the FTP packets.

forwarding-class *class-name*—Forwarding class name to apply to the FTP packets. The choices are:

- 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

[Class of Service for Services PICs \(Next Gen Services\)](#) | 255

gate-timeout

Syntax

```
gate-timeout seconds;
```

Hierarchy Level

```
[edit applications application ike-esp-nat]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Description

For an IKE ALG application, configure the length of time that can pass after IKE establishes the security association between the IPsec client and server and before the ESP traffic starts in both directions. If the ESP traffic has not started before this timeout value, the ESP gates are deleted and the ESP traffic is blocked.

The IKE ALG enables the passing of IKEv1 and IPsec packets through NAPT-44 and NAT64 rules between IPsec peers that are not NAT-T compliant.

Options

seconds—Number of seconds.

Default: 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

ALG Descriptions

[Configuring Application Sets | 421](#)

Configuring Application Properties

global-dns-stats-log-timer

Syntax

```
global-dns-stats-log-timer minutes;
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure the interval for logging per-client statistics for filtering of DNS requests for disallowed website domains.

Options

minutes—The number of minutes in the logging interval.

Default: 5

Range: 0 through 60

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[DNS Request Filtering for Disallowed Website Domains](#) | 311

group (Traffic Load Balancer)

Syntax

```
group group-name {
  health-check-interface-subunit health-check-interface-subunit;
  network-monitoring-profile [profile-name1, <profile-name2>];
  real-service-rejoin-options no-auto-rejoin;
  real-services [server-list];
  <routing-instance routing-instance>;
}
```

Hierarchy Level

```
[edit services traffic-load-balance instance instance-name]
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure a group of servers as a pool for next-hop session distribution.

Options

group-name—Use the specified string identifier for a group of servers to which sessions are distributed using the server distribution table in conjunction with the session distribution API.

group health-check-interface-subunit health-check-interface-subunit—Use the specified subunit of the ms- interface used for health checking.

network-monitoring-profile profile-name1—Name of the network monitoring profile used to monitor the health of servers in the group.

network-monitoring-profile profile-name2—(Optional) Name of a second network monitoring profile used to monitor the health of servers in the group.

real-services server-list—Use the specified list of individual servers to which sessions are distributed using the server distribution table in conjunction with the session distribution API.

real-services-rejoin-options no-auto-rejoin—Disable the default behavior that allows a server to rejoin the group automatically when it comes up.

routing-instance routing-instance—(Optional) Use the specified routing instance if the default **inet.0** is not used.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Traffic Load Balancer Overview | 287](#)

[Configuring TLB | 296](#)

hash-keys (Interfaces)

Syntax

```
hash-keys {
  egress-key (source-ip | destination-ip);
  ingress-key (source-ip | destination-ip);
  ipv6-source-prefix-length ipv6-source-prefix-length;
}
```

Hierarchy Level

```
[edit interfaces unit unit-name load-balancing-options]
```

Release Information

Statement introduced in Junos OS Release 11.4.

ipv6-source-prefix-length option introduced in Junos OS Release 18.2R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card. The **ipv6-source-prefix-length** option is not supported for Next Gen Services.

Description

Configure the hash keys used for load balancing in aggregated multiservices (AMS) for next-hop style services. The hash keys supported in the ingress and egress direction are the source IP address and destination IP address.

Hash keys are used to define the load-balancing behavior among the various members in the AMS. For example, if **hash-keys** is configured as **source-ip**, then the hashing is performed based on the source IP address of the packet, so that all packets with the same source IP address land on the same member. When you use **ingress-key** and **egress-key**, you must configure hash keys to take the traffic direction into consideration. For example, if you configure **hash-keys** as **source-ip** in the ingress direction, then you must configure **hash-keys** as **destination-ip** in the egress direction. This is required to ensure that the packets of the same flow reach the same member of the AMS group.

If you are configuring an AMS interface used in a service set for DS-Lite,

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

Options

egress-key destination-ip—Use the destination IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction.

egress-key source-ip—Use the source IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the egress flow direction.

ingress-key destination-ip—Use the destination IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the ingress flow direction.

ingress-key source-ip—Use the source IP address of the flow to compute the hash used in load balancing. Configure the hash keys to be used in the ingress flow direction.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Load Balancing on AMS Infrastructure](#) | 361

header-integrity-check (Next Gen Services)

Syntax

```
header-integrity-check {  
    enable-all;  
}
```

Hierarchy Level

```
[edit services service-set service-set service-set-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Drop packets that have packet header anomalies. These anomalies include:

- Not an IP packet
- Not an IPv4 packet or an IPv6 packet
- TTL error (TTL is 0)
- Bad source/destination IP
- IP checksum error
- Protocol error
- TCP port zero
- TCP header length error (less than 20 bytes)
- TCP SEQNUM is zero and no flags are set
- TCP SEQNUM is zero and flags are set
- No TCP flags are set
- TCP FIN with no Ack
- TCP FIN & Reset
- TCP SYN & (FIN or URG or RESET)
- UDP port zero
- UDP header length error
- ICMP header length error (not within 48-576 bytes)

- ICMP packet error length
- ICMP large packet (1024)

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

high-availability-options (Aggregated Multiservices)

Syntax

```
high-availability-options {
  (many-to-one | one-to-one) {
    preferred-backup preferred-backup;
  }
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure the high availability options for the aggregated multiservices (AMS) interface. For service applications, if only the load-balancing feature is being used, then this configuration is optional.

For many-to-one (N:1) high availability support for service applications like Network Address Translation (NAT), the preferred backup services PIC, in hot standby mode, backs up one or more (N) active services PICs.

NOTE: In both cases, if one of the active services PICs goes down, then the backup replaces it as the active PIC. When the failed PIC comes back up, it becomes the new backup. This is called *floating backup*.

One-to-one (1:1) high availability support associates a single backup interface with a single active interface. 1:1 configuration is supported only on the MS-MPC and MX-SPC3. In 1:1 (stateful) configurations, synchronization causes the active and back up PICs to synchronize traffic states and data structures, preventing data loss during a failover event. Stateful synchronization is required for IPsec high availability support. For IPsec connections, AMS supports 1:1 configuration only.

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

[load-balancing-options | 575](#)

[Understanding Aggregated Multiservices Interfaces](#)

[Example: Configuring an Aggregated Multiservices Interface \(AMS\)](#)

host (Next Gen Services Service-Set Remote System Logging)

Syntax

```
host host-ip-address;
```

Hierarchy Level

```
edit services service-set name syslog stream stream-name
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Specify the IP address of syslog server to receive log messages.

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

host-address-base (Source NAT Next Gen Services)

Syntax

```
host-address-base ip-address;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure static mapping of the source address.

For static NAT that is performed on the services card, configure a one-to-one static shifting of a range of original source addresses to the range of addresses in the source pool by specifying the base address of the original source address range.

For example, if the host address base is 198.51.100.30 and the NAT pool uses the range 203.0.113.10 to 203.0.113.20, then 198.51.100.30 translates to 203.0.113.10, 198.51.100.31 translates to 203.0.113.11, and so on.

Options

host-address-base *ip-address*—The IP address used as the host address base.

Required Privilege Level

interface—To view this statement in the configuration.

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

inactivity-timeout

Syntax

```
inactivity-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]  
[edit services service-set-name service-set-options]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for MX-SPC3 services card on MX240, MX480 and MX960 routers.

Description

Configure the inactivity timeout period for established flows. The timeout value configured in the application protocol definition overrides this value.

Options

seconds—Timeout period.

Default: 30 seconds

Range: 4 through 86,400 seconds

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*

inactivity-asymm-tcp-timeout (Service Set Next Gen Services)

Syntax

```
inactivity-asymm-tcp-timeout seconds;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options tcp-session]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure the number of seconds that a unidirectional TCP session can be inactive before it is closed. Valid settings: 4 through 86400 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

icmp (IDS Screen Next Gen Services)

Syntax

```
icmp {  
    fragment;  
    icmpv6-malformed;  
    large;  
    ping-death;  
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure ICMP intrusion detection service options.

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

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

icmp-type

Syntax

```
icmp-type value;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

ICMP packet type value.

Options

value—The ICMP type value, such as **echo** or **echo-reply**. For a complete list, see *Configuring the ICMP Code and Type*.

Required Privilege Level

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

RELATED DOCUMENTATION

<i>ALG Descriptions</i>	
Configuring Application Sets	421
<i>Configuring the ICMP Code and Type</i>	
Examples: Configuring Application Protocols	437
Verifying the Output of ALG Sessions	438

icmpv6-malformed (IDS Screen Next Gen Services)

Syntax

```
icmpv6-malformed;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop malformed ICMPv6 packets, which might cause damage to the device and network. Examples of malformed IPv6 packets are packets that are too big (message type 2), that have the next header set to routing (43), or that have a routing header set to hop-by hop.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

ip (IDS Screen Next Gen Services)

Syntax

```
ip {
  bad-option;
  block-frag;
  ipv6-extension-header {
    AH-header;
    ESP-header;
    fragment-header;
    hop-by-hop-header {
      CALIPSO-option;
      jumbo-payload-option;
      quick-start-option;
      router-alert-option;
      RPL-option;
      SFM-DPD-option;
      user-defined-option-type <type-low> to <type-high>;
    }
    mobility-header;
    routing-header;
  }
  loose-source-route-option;
  record-route-option;
  security-option;
  source-route-option;
  stream-option;
  strict-source-route-option;
  tear-drop;
  timestamp-option;
  unknown-protocol;
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R1 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure protection against suspicious IP packet attacks.

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 Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

ipv6-extension-header (IDS Screen Next Gen Services)

Syntax

```

ipv6-extension-header {
  AH-header;
  ESP-header;
  fragment-header;
  hop-by-hop-header {
    CALIPSO-option;
    jumbo-payload-option;
    quick-start-option;
    router-alert-option;
    RPL-option;
    SFM-DPD-option;
    user-defined-option-type <type-low> to <type-high>;
  }
  mobility-header;
  routing-header;
}

```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IP packets that have the configured IPv6 extension header values.

Options

ah-header—Authentication Header extension header

esp-header—Encapsulating Security Payload extension header

fragment-header—Fragment Header extension header

hop-by-hop-header—The specified Hop-by-Hop option:

CALIPSO-option—Common Architecture Label IPv6 Security Option

jumbo-payload-option—IPv6 jumbo payload option

quick-start-option—IPv6 quick start option

router-alert-option—IPv6 router alert option

RPL-option—Routing Protocol for Low-Power and Lossy Networks option

SFM-DPD-option—Simplified Multicast Forwarding IPv6 Duplicate Packet Detection option

user-defined-option-type *type-low* to *type-high*—A range of header types

Range: 1 through 255.

mobility-header—Mobility Header extension header

routing-header—Routing Header extension header

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

limit-session (IDS Screen Next Gen Services)

Syntax

```
limit-session {  
  by-destination{  
    by-protocol {  
      icmp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      tcp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      udp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
    }  
    maximum-sessions number;  
    packet-rate number;  
    session-rate number;  
  }  
  by-source {  
    by-protocol {  
      icmp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      tcp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      udp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
    }  
  }  
}
```

```

    maximum-sessions number;
    packet-rate number;
    session-rate number;
  }
}

```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure session limits for individual destination or source addresses, or for individual destination or source subnets. This protects against network probing attacks and network flooding attacks. You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol. When a session limit is exceeded for a source or destination, packets from the source or to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination or source subnets rather than individual addresses, include the **aggregations** statement at the `[edit services screen ids-option screen-name]` 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

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

inline-services (PIC level)

Syntax

```
inline-services {  
    bandwidth (1g | 10g | 20g | 30g | 40g | 100g);  
}
```

Hierarchy Level

```
[edit chassis fpc slot-number pic number]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Enable inline services on PICs residing on MPCs and optionally specify a bandwidth for traffic on the inline service interface.

NOTE: For an MPC, such as MPC2, always configure inline-services at the **[chassis fpc slot-number pic number]** hierarchy level. Do not configure inline services for a service card such as MS-MPC.

The remaining statement is explained separately. Search for a statement in [CLI Explorer](#) or click a linked statement in the Syntax section for details.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Enabling Inline Service Interfaces

Configuring an L2TP LNS with Inline Service Interfaces

ipv6-extension-header (IDS Screen Next Gen Services)

Syntax

```

ipv6-extension-header {
    AH-header;
    ESP-header;
    fragment-header;
    hop-by-hop-header {
        CALIPSO-option;
        jumbo-payload-option;
        quick-start-option;
        router-alert-option;
        RPL-option;
        SFM-DPD-option;
        user-defined-option-type <type-low> to <type-high>;
    }
    mobility-header;
    routing-header;
}

```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IP packets that have the configured IPv6 extension header values.

Options

ah-header—Authentication Header extension header

esp-header—Encapsulating Security Payload extension header

fragment-header—Fragment Header extension header

hop-by-hop-header—The specified Hop-by-Hop option:

CALIPSO-option—Common Architecture Label IPv6 Security Option

jumbo-payload-option—IPv6 jumbo payload option

quick-start-option—IPv6 quick start option

router-alert-option—IPv6 router alert option

RPL-option—Routing Protocol for Low-Power and Lossy Networks option

SFM-DPD-option—Simplified Multicast Forwarding IPv6 Duplicate Packet Detection option

user-defined-option-type *type-low* to *type-high*—A range of header types

Range: 1 through 255.

mobility-header—Mobility Header extension header

routing-header—Routing Header extension header

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

instance (Traffic Load Balancer)

Syntax

```
instance instance-name {
  client-interface client-interface;
  client-vrf client-vrf;
  group group-name {
    health-check-interface-subunit health-check-interface-subunit;
    network-monitoring-profile profile-name;
    real-service-rejoin-options no-auto-rejoin;
    real-services [ server-list ];
    <routing-instance routing-instance>;
  }
  interface interface-name;
  real-service real-service {
    address server-ip-address;
    admin-down;
  }
  server-inet-bypass-filter server-inet-bypass-filter ;
  server-inet6-bypass-filter server-inet6-bypass-filter ;
  server-interface server-interface;
  server-vrf server-vrf-name;
  virtual-service virtual-service-name {
    address virtual-ip-address;
    group group-name;
    load-balance-method {
      hash {
        hash-key method;
      }
      random;
    }
    mode (layer2-direct-server-return | direct-server-return | translated);
    <routing-instance routing-instance-name>;
    <routing-metric route-metric>;
    server-interface server-interface;
    service service-name {
      protocol (udp | tcp);
      server-listening-port port;
      virtual-port virtual-port;
    }
  }
}
```

Hierarchy Level

```
[edit services traffic-load-balance]
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure a Traffic Load Balancer instance.

Options

client-interface *client-interface*—For translated mode, client interface where the implicit filter is installed to direct the traffic in the forward direction.

client-vrf *client-vrf*—Use the specified name of the routing instance in which the data traffic in the reverse direction is routed to the clients.

instance *instance-name*—Identifier (text string) for a TLB configuration.

server-inet-bypass-filter *server-inet-bypass-filter*—Name of the firewall filter from which the terms are referenced and added to the server-side implicit filters. This enables the operator to bypass reverse (RIP to VIP) translation of IPv4 traffic.

server-inet6-bypass-filter *server-inet6-bypass-filter*—Name of the firewall filter from which the terms are referenced and added to the server-side implicit filters. This enables the operator to bypass reverse (RIP to VIP) translation of IPv6 traffic.

server-interface *server-interface*—For translated mode, specifies the server interfaces where the server filters are implicitly installed to direct the return traffic to the load balancing next hop.

server-vrf *server-vrf-name*—The routing instance in which the data traffic in the forward direction is routed to the 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

[Traffic Load Balancer Overview](#) | 287

[Configuring TLB](#) | 296

interface-service (Services Interfaces)

Syntax

```
interface-service {  
  load-balancing-options {  
    hash-keys {  
      egress-key (destination-ip | source-ip);  
      ingress-key (destination-ip | source-ip);  
    }  
  }  
  service-interface name;  
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the device name for the interface service Physical Interface Card (PIC).

Options

service-interface *name*—Name of the service device associated with the interface-wide service set.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Service Sets to be Applied to Services Interfaces*

land (IDS Screen Next Gen Services)

Syntax

```
land;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop SYN packets that have the same source and destination address or port, which protects against land attacks. In a land attack, the target using up its resources as it repeatedly replies to itself.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

large (IDS Screen Next Gen Services)

Syntax

```
large;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop any ICMP frame with an IP length greater than 1024 bytes, which protects against ICMP large packet attacks.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

limit-session (IDS Screen Next Gen Services)

Syntax

```
limit-session {  
  by-destination{  
    by-protocol {  
      icmp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      tcp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      udp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
    }  
    maximum-sessions number;  
    packet-rate number;  
    session-rate number;  
  }  
  by-source {  
    by-protocol {  
      icmp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      tcp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
      udp {  
        maximum-sessions number;  
        packet-rate number;  
        session-rate number;  
      }  
    }  
  }  
}
```

```

    maximum-sessions number;
    packet-rate number;
    session-rate number;
  }
}

```

Hierarchy Level

```
[edit services screen ids-option screen-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure session limits for individual destination or source addresses, or for individual destination or source subnets. This protects against network probing attacks and network flooding attacks. You can specify limits for specific protocols (ICMP, TCP, and UDP), or specify limits independent of a protocol. When a session limit is exceeded for a source or destination, packets from the source or to the destination are dropped until the session limit is no longer exceeded.

To specify limits for destination or source subnets rather than individual addresses, include the **aggregations** statement at the `[edit services screen ids-option screen-name]` 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

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

load-balancing-options (Aggregated Multiservices)

Syntax

```
load-balancing-options {
  high-availability-options {
    (many-to-one | one-to-one) {
      preferred-backup preferred-backup;
    }
  }
  member-failure-options {
    drop-member-traffic {
      rejoin-timeout rejoin-timeout;
    }
    redistribute-all-traffic {
      enable-rejoin;
    }
  }
  hash-keys {
    egress-key (destination-ip | source-ip);
    ingress-key (destination-ip | source-ip);
  }
  member-interface interface-name;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure the high availability (HA) options for the aggregated multiservices (AMS) interface.

Many-to-one (N:1) high availability mode for service applications like Network Address Translation (NAT) is supported. In the case of N:1 high availability mode, one services PIC is the backup (in hot standby mode) for one or more (N) active services PICs. If one of the active services PICs goes down, then the backup replaces it as the active services PIC. When the failed PIC comes back online, it becomes the new backup. This is called *floating backup mode*. In an N:1 (stateless) configuration, traffic states and data structures are not synchronized between active PICs and the backup PIC.

You can also configure a one-to-one (1:1) high availability mode. In the 1:1 configuration, a single interface is configured as the backup for another single active interface. If the active interface goes down, the backup interface replaces it as the active interface. A 1:1 (stateful) configuration synchronizes traffic states and data structures between the active services PIC and the backup services PIC. This is required for IPsec connections. One-to-one high availability is supported on the MS-MPC but it is not supported for MX-SPC3 in this release.

Load-balancing might not be uniform among member interfaces in certain network deployments. The variance can be because of a misconfiguration, which causes the traffic itself not to be sufficiently randomly distributed, causing the hash keys to be ineffective (for example, the hash key is destination IP but all sessions have only source IP address). The variation can be within the expected range and the load balancing depends on the IP addresses chosen. The hash calculation performs a checksum on several bits of the IP address and not only on the last few lower significant bits of the IP address. In such a scenario, the load-balancing ratio can change, for instance, if the source IP address is changed from 20.0.0.0/24 to 20.0.1.0/24.

The distribution of traffic across member interfaces of an AMS interface is static load-balancing. Flows are load balanced based on a packet hash on parameters such as source IP or destination IP. Load-balancing effectiveness depends on the IP address or protocol diversity. For example, if the hash key is destination IP and all packets have the same destination, then all flows are directed to the same member. This is flow-level load balancing and not per packet. As a result, traffic between a pair of addresses may be 10,000 pps, whereas another pair of addresses may have 1 pps. The load of the former is not distributed among members. High availability is limited to stateless HA. When a backup interface takes over as an active interface, all flows are reestablished (for example, packets may undergo NAT processing differently after failover).

With a stateful firewall, static NAT as basic-nat44 or destination-nat44, and dynamic NAT as nat64, napt-44, dynamic-nat44, and with application layer gateways (ALGs) configured, NAT hairpinning is not supported. Input direction for rule match to be applied is supported only for dynamic NAT types (NAT64, NAT44, and dynamic-NAT44). Service-set policies need to have input or input-output direction only. Flows on all active members are reset when the number of actives changes. The resetting of flows can be avoided at the cost of failed-member's traffic loss using certain options.

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

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

local-category (Next Gen Services Service-Set Local System Logging)

Syntax

```
local-category category, category....category;
```

Hierarchy Level

```
[edit services service-set name syslog
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Specify the category for which you want to collect local logs.

Options

all—All events are logged

content-security—Content security events are logged

fw-auth—Fw-auth events are logged

screen —Screen events are logged

alg—Alg events are logged

nat—NAT events are logged

flow—Flow events are logged

sctp—Sctp events are logged

gtp—Gtp events are logged

ipsec—Ipsec events are logged

idp—Idp events are logged

rtlog—Rtlog events are logged

pst-ds-lite—Pst-ds-lite events are logged

appqos—Appqos events are logged

secintel—Secintel events are logged

aamw—AAMW events are logged

sfw—Stateful Firewall events are logged

session —Session open and close events are logged

session-open—Session open events are logged

session-close—Session close events are logged

urlf—DNS request filtering events are logged

ha—Stateful High-Availability open and close events are logged

ha-open—Stateful High-Availability open events are logged

ha-close—Stateful High-Availability close events are logged

pcp—PCP logs

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

local-log-tag (Next Gen Services Service-Set System Logging)

Syntax

```
local-log-tag tag-stamp;
```

Hierarchy Level

```
[edit services service-set name syslog  
edit services service-set name syslog stream stream-name
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Each log message is stamped with this tag.

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

loose-source-route-option (IDS Screen Next Gen Services)

Syntax

```
loose-source-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IPv4 packets that have the IP option of 3 (Loose Source Routing).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

many-to-one (Aggregated Multiservices)

Syntax

```
many-to-one {
  preferred-backup preferred-backup;
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options high-availability-options]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure the many-to-one (N:1) preferred backup for the aggregated multiservices (AMS) interface.

NOTE: The preferred backup must be one of the member interfaces (mams-) that have already been configured at the `[edit interfaces interface-name load-balancing-options]` hierarchy level. Even in the case of mobile control plane redundancy, which is one-to-one (1:1), the initial preferred backup is configured at this hierarchy level.

Options

preferred-backup *preferred-backup*—Use the specified interface as the preferred backup member interface.

The member interface format is mams-*a/b/0*, where *a* is the FPC slot number and *b* is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[high-availability-options \(Aggregated Multiservices\)](#) | 551

Understanding Aggregated Multiservices Interfaces

| *Example: Configuring an Aggregated Multiservices Interface (AMS)*

map-e

Syntax

```
map-e name {
  disable-auto-route;
  ea-bits-len ea-bits-len;
  ipv4-prefix ipv4-prefix;
  mape-prefix mape-prefix;
  mtu-v6 mtu-v6;
  psid-length psid-length;
  psid-offset psid-offset;
  software-address software-address;
  v4-reassembly;
  v6-reassembly;
  version-03;
}
```

Hierarchy Level

```
[edit services software software-concentrator]
[edit services softwares software-types]
```

Release Information

Statement introduced in Junos OS Release 18.2R1 for MX Series Routers with MPC and MIC interfaces. Support added in Junos OS release 20.2R1 at MAP-E for Next Gen Services on MX240, MX480, and MX960 routers.

Description

Configure Mapping of Address and port – Encapsulation (MAP-E) as an inline service on MX Series routers that use MPC and MIC interfaces. MAP-E is an automatic tunneling mechanism that encapsulates IPv4 packets within an IPv6 address. The IPv4 packets are carried in an IPV4-over-IPV6 tunnel from the MAP-E Customer Edge (CE) devices to the MAP-E Provider Edge (PE) devices (also called as Border Relay (BR) devices) through an IPV6 routing topology, where they are de-tunneled for further processing.

Options

disable-auto-route—Disable auto-routes and enable static routes to facilitate ECMP load balancing.

NOTE: When you enable the **disable-auto-route** option, you must configure static routes.

name—Name of the MAP-E software concentrator.

ea-bits-len—Configure rule for Embedded Address (EA) length for the MAP-E domain.

NOTE:

- If **v4-prefix-len** is 0 then **ea-bits-len** must be non-zero, and vice versa.
- It is possible that **ea-bits-len** is equal to 0, but **psid-len** is non-zero.
- If the sum of **v4-prefix-len** and **ea-bits-len** is less than 32, then the **psid-len** must be equal to the difference between 32 and the sum total of **v4-prefix-len** and **ea-bits-len**.

Range: 0 through 48

ipv4-prefix—Configure rule for IPv4 prefix and length of the MAP-E domain.

Range: 0 through 32

map-e-prefix—Configure rule for IPV6 prefix and length for the MAP-E domain. The MAP-E IPv4 and IPv6 prefix must be unique per software concentrator.

mtu-v6—(Optional) Specify the Maximum transmission unit (MTU) for the MAP-E software tunnel.

Default: 9192

Range: 1280 through 9192

psid-length—Configure Port Set ID (PSID) length value for the MAP-E domain.

NOTE:

- If the sum of **v4-prefix-len** and **ea-bits-len** is less than 32, then the **psid-len** must be equal to the difference between 32 and the sum total of **v4-prefix-len** and **ea-bits-len**.

Range: 0 through 16

psid-offset—(Optional) Configure PSID offset value for the MAP-E domain.

Default: 4

Range: 0 through 16

softwire-address—Specify the Border Relay device unicast IPv6 address as the softwire concentrator IPv6 address.

v4-reassembly | v6-reassembly—(Optional) Enable IPv4 and IPv6 reassembly for MAP-E.

version-03—(Optional) Configure version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03 (expires on July 28, 2013), *Mapping of Address and Port with Encapsulation (MAP)* and the latest available version.

Required Privilege Level

system

mapping-timeout (Source NAT Next Gen Services)

Syntax

```
mapping-timeout mapping-timeout;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the timeout period for address-pooling paired mappings that use the specified NAT pool. Mappings that are inactive for this amount of time are dropped.

If you do not configure **ei-mapping-timeout** for endpoint independent translations, then the **mapping-timeout** value is used for endpoint independent translations.

Options

mapping-timeout mapping-timeout—Length of timeout period in seconds.

Range: 120 through 86,400

Default: 300

Required Privilege Level

interface—To view this statement in the configuration.

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

mapping-type (Source NAT Next Gen Services)

Syntax

```
mapping-type {  
    address-pooling-paired;  
    endpoint-independent;  
}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure the source NAT mapping type.

Options

endpoint-independent —Mapping to ensure that the same external address and port are assigned to all connections from a given host.

address-pooling-paired —Mapping to ensure assignment of the same external IP address for all sessions originating from the same internal host.

Required Privilege Level

interface—To view this statement in the configuration.

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

match (Next Gen Services Global System Logging)

Syntax

```
match match;
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Regular expression for lines to be logged

Options

match—Regular expression for lines to be logged

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

match (Services CoS Next Gen Services)

Syntax

```
match {
  application [ application-names ];
  destination-address address;
  destination-address-range low minimum-value high maximum-value;
  destination-port port-number;
  destination-prefix-list list-name;
  source-address address;
  source-address-range low minimum-value high maximum-value;
  source-prefix-list list-name;
}
```

Hierarchy Level

```
[edit services cos rule rule-name policy policy-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure the matching conditions for a policy in a services CoS rule. Matching conditions include packet source and destination addresses and packet applications. Packets that are processed by a service set and that match the conditions are assigned the Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignments specified in the policy.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

application [*application-names*]—One or more port-based applications.

destination-address *address*—Destination address of the packet.

destination-address-range low *minimum-value* high *maximum-value*—Range of destination addresses of the packet.

minimum-value—Lower boundary of address range.

maximum-value—Upper boundary of address range.

destination-port *port-number*—Destination port number of the packet.

source-address *address*—Source address of the packet.

source-address-range *low minimum-value high maximum-value*—Range of source addresses of the packet.

minimum-value—Lower boundary of address range.

maximum-value—Upper boundary of address range.

source-prefix-list *list-name*—Name of a prefix list for matching the source address prefix.

You configure the prefix list by using the **prefix-list** statement at the **[edit policy-options]** hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Class of Service for Services PICs \(Next Gen Services\)](#) | 255

match (Stateful Firewall Rule Next Gen Services)

Syntax

```
match {
  application [application-name];
  destination-address (address | any);
  destination-address-excluded address;
  source-address (address | any);
  source-address-excluded address;
}
```

Hierarchy Level

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the matching properties for a stateful firewall rule policy. When a flow matches these properties, the policy actions are applied to the flow.

Options

application [*application-name*]*—*One or more application protocols of flows to which the stateful firewall policy applies. The application protocol definition is configured at the **[edit applications]** hierarchy level.

destination-address (*address* | **any**)*—*The destination address of the flows to which the stateful firewall rule policy applies. The option **any** matches all destination addresses.

destination-address-excluded *address**—*The destination address of the flows to which the stateful firewall rule policy does not apply.

source-address (*address* | **any**)*—*The source address of the flows to which the stateful firewall rule policy applies. The option **any** matches all source addresses.

source-address-excluded *address**—*The source address of the flows to which the stateful firewall rule policy does not apply.

Required Privilege Level

interface*—*To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Stateful Firewalls for Next Gen Services](#) | 267

match-direction (NAT Next Gen Services)

Syntax

Hierarchy Level

```
[edit services nat source rule-set rule-set],  
[edit services nat destination rule-set rule-set]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Required Privilege Level

interface—To view this statement in the configuration.

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

match-rules-on-reverse-flow (Next Gen Services)

Syntax

```
match-rules-on-reverse-flow;
```

Hierarchy Level

```
[edit services service-set service-set-name cos-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure the service set to create a CoS session even if a packet is first received in the reverse direction of the matching direction of the CoS rule. The CoS rule values are then applied as soon as a packet in the correct match direction is received.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Class of Service for Services PICs \(Next Gen Services\)](#) | 255

max-session-setup-rate (Service Set)

Syntax

```
max-session-setup-rate (number | numberk);
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 17.1R1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Set the maximum number of session setups allowed per second for the service set. After this setup rate is reached, any additional session setup attempts are dropped. If you do not include the **max-session-setup-rate** statement, the session setup rate is not limited.

Options

max-session-setup-rate *number*—Use the specified maximum number of session setups per second.

Range: 1 through 429,496,729

Default: 0 (The session setup rate is not limited.)

numberk—Maximum number of sessions, expressed in thousands. Starting in Junos OS Release 18.4R1, 1k=1000. Prior to Junos OS Release 18.4R1, 1k=1024.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Service Set Limitations](#)

max-sessions-per-subscriber (Service Set Next Gen Services)

Syntax

```
max-sessions-per-subscriber session-number;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Set the maximum number of sessions allowed from a single subscriber.

Options

session-number—Maximum number of sessions.

NOTE: There is no default value. You must configure a value for the configuration to take effect.

Range: 1 through 32000

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

maximum

Syntax

```
maximum number;
```

Hierarchy Level

```
[edit interfaces interface-name services-options session-limit]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the maximum number of sessions allowed simultaneously on services cards. If you specify the maximum number of sessions to be zero, it indicates that the configuration is not effective. You must specify a value higher than zero for the maximum number of sessions.

Options

number—Maximum number of sessions.

Range: 1 through 4,294,967,295

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

member-failure-options (Aggregated Multiservices)

Syntax

```
member-failure-options {
  drop-member-traffic {
    rejoin-timeout rejoin-timeout;
  }
  redistribute-all-traffic {
    enable-rejoin;
  }
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure the possible behavior for the aggregated Multiservices (AMS) interface in case of failure of more than one active member.

NOTE: The **drop-member-traffic** configuration and the **redistribute-all-traffic** configuration are mutually exclusive.

[Table 54 on page 597](#) displays the behavior of the member interface after the failure of the first services PIC. [Table 55 on page 597](#) displays the behavior of the member interface after the failure of two services PICs.

NOTE: The AMS infrastructure has been designed to handle one failure automatically. However, in the unlikely event that more than one services PIC fails, the AMS infrastructure provides configuration options to minimize the impact on existing traffic flows.

Table 54: Behavior of Member Interface After One Multiservices PIC Fails

High Availability Mode	Member Interface Behavior
Many-to-one (N:1) high availability support for service applications	Automatically handled by the AMS infrastructure

Table 55: Behavior of Member Interface After Two Multiservices PICs Fail

High Availability Mode	Configuration	rejoin-timeout	Behavior when member rejoins before rejoin-timeout expires	Behavior when member rejoins after rejoin-timeout expires
Many-to-one (N:1) high availability support for service applications	drop-member-traffic	Configured	<p>The existing traffic for the second failed member will <i>not</i> be redistributed to the other members.</p> <p>The first member to rejoin becomes an active member. The second member to rejoin becomes the backup. This behavior is handled automatically by the AMS infrastructure.</p>	<p>The existing traffic for the second failed member will <i>not</i> be redistributed to the other members.</p> <p>The first member will rejoin the AMS automatically. However, the other members who are rejoining will be moved to the discard state.</p>
Many-to-one (N:1) high availability support for service applications	redistribute-all-traffic	Not applicable	<p>Before rejoin, the traffic is redistributed to existing active members.</p> <p>After a failed member rejoins, the traffic is load-balanced afresh. This may impact existing traffic flows.</p>	

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

Default

If **member-failure-options** are not configured, then the default behavior is to drop member traffic with a rejoin timeout of 120 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[load-balancing-options \(Aggregated Multiservices\)](#) | 575

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

member-interface (Aggregated Multiservices)

Syntax

```
member-interface interface-name;
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the member interfaces for the aggregated multiservices (AMS) interface. You can configure multiple interfaces by specifying each interface in a separate statement.

Starting with Junos OS Release 16.2, an AMS interface can have up to 32 member interfaces. In Junos OS Release 16.1 and earlier, an AMS interface can have a maximum of 24 member interfaces. If you configure more than 24 member interfaces, you must set the *pic-boot-timeout* value to 240 or 300 seconds at the **[edit interfaces *interface-name* multiservice-options]** hierarchy level for every services PIC interface on the MX Series router.

For high availability service applications like Network Address Translation (NAT) that support many-to-one (N:1) redundancy, you can specify two or more interfaces.

On an MS-MPC, you can configure one-to-one (1:1) redundancy. In a 1:1 (stateful) configuration, a single backup interface provides redundancy for a single active interface. A 1:1 configuration is required for IPsec. 1:1 redundancy is not supported on the MX-SPC3 in this release.

NOTE: The member interfaces that you specify must be members of aggregated multiservices interfaces (mams-).

Options

interface-name—Name of the member interface. The member interface format is mams-*a*/*b*/0, where *a* is the FPC slot number and *b* is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Understanding Aggregated Multiservices Interfaces for Next Gen Services](#) | 353

[Configuring Aggregated Multiservices Interfaces](#) | 359

[load-balancing-options \(Aggregated Multiservices\)](#) | 575

mode (Next Gen Services Service-Set System Logging)

Syntax

```
mode {  
    event ;  
    stream stream-name;  
}
```

Hierarchy Level

```
[edit services services-set name syslog]
```

Release Information

Support introduced in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Mode in which the system message logger sends messages

Options

event—Send messages to a file on the local routing engine

stream—Send messages to one or more remote log servers. Each remote server requires its own stream.

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

name (Next Gen Services Global System Logging)

Syntax

```
name;
```

Hierarchy Level

```
[edit services rtlog traceoptions flag]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Specify what to flag in the trace information.

Options

all—Everything

configuration—Reading of configuration

hpl—Trace HPL logging

report—Trace report

source—Communication with security log forwarder

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

nat-options (Next Gen Services)

Syntax

```
nat-options {  
  nptv6 {  
    icmpv6-error-messages;  
  }  
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Send ICMP error messages if NPTv6 address translation fails.

Required Privilege Level

interface—To view this statement in the configuration.

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

nat-rule-sets (Service Set Next Gen Services)

Syntax

```
nat-rule-sets rule-set-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the NAT rules set included in the service set.

Required Privilege Level

interface—To view this statement in the configuration.

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

next-hop-service

Syntax

```
next-hop-service {
  inside-service-interface interface-name.unit-number;
  outside-service-interface interface-name.unit-number;
  outside-service-interface-type interface-type;
  service-interface-pool name;
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

service-interface-pool option added in Junos OS Release 9.3.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify interface names or a service interface pool for the forwarding next-hop service set. You cannot specify both a service interface pool and an inside or outside interface.

Options

inside-service-interface *interface-name.unit-number*—Name and logical unit number of the service interface associated with the service set applied inside the network.

outside-service-interface *interface-name.unit-number*—Name and logical unit number of the service interface associated with the service set applied outside the network.

outside-service-interface-type *interface-type*—Identifies the interface type of the service interface associated with the service set applied outside the network. For inline IP reassembly, set the interface type to local.

service-interface-pool *name*—Name of the pool of logical interfaces configured at the **[edit services service-interface-pools pool *pool-name*]** hierarchy level. You can configure a service interface pool only if the service set has a PGCP rule configured. The service set cannot contain any other type of rule.

NOTE: **service-interface-pool** is not applicable for IP reassembly configuration on L2TP.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Service Sets to be Applied to Services Interfaces*

no-remote-trace (Next Gen Services Global System Logging)**Syntax**

```
no-remote-trace;
```

Hierarchy Level

```
[edit services rtlog traceoptions]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Disable remote tracing

Required Privilege Level

system

RELATED DOCUMENTATION

| [Understanding Next Gen Services CGNAT Global System Logging](#) | **71**

| [Enabling Global System Logging for Next Gen Services](#) | **73**

| [Configuring System Logging to One or More Remote Servers for Next Gen Services](#) | **76**

| [Configuring Local System Logging for Next Gen Services](#) | **74**

no-translation (Source NAT Next Gen Services)

Syntax

```
no-translation;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Disable port translation for NAT. By default, port translation is enabled for NAT.

Required Privilege Level

interface—To view this statement in the configuration.

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

no-world-readable (Next Gen Services Global System Logging)

Syntax

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Don't allow any user to read the log file

Options

no-world-readable—Don't allow any user to read the log file

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

off (Destination NAT Next Gen Services)

Syntax

```
off;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set-name rule rule-name then destination-nat]
```

Description

Tun off destination address translation for the rule. Use this statement when configuring port forwarding without destination address translation.

Required Privilege Level

interface—To view this statement in the configuration.

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

open-timeout

Syntax

```
open-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

```
[edit services service-set service-set-name service-set-options tcp-session]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure a timeout period for Transmission Control Protocol (TCP) session establishment.

Options

seconds—Timeout period.

Default: 5 seconds

Range: 4 through 224 seconds

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*

pcp-rules

Syntax

```
pcp-rules rule-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 13.2R1.

Description

Specify the PCP rule to apply to the service set. A PCP rule assigns the PCP server that handles selected traffic.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC. Starting in Junos OS Release 20.1R1, PCP is also supported for Next Gen Services.

Options

rule-name—The PCP rule to apply to the service set.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

ping-death (IDS Screen Next Gen Services)

Syntax

```
ping-death;
```

Hierarchy Level

```
[edit services screen ids-option screen-name icmp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop oversized and irregular ICMP packets, which protects against the ping of death attack. In the ping of death attack, the attacker sends the target ping packets whose IP datagram length (ip_len) exceeds the maximum legal length (65,535 bytes) for IP packets, and the packets are fragmented. When the target attempts to reassemble the IP packets, a buffer overflow might occur, resulting in system crashing, freezing, and restarting.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

policy (Services CoS Next Gen Services)

Syntax

```
policy policy-name {
  match {
    application [ application-names ];
    destination-address address;
    destination-address-range low minimum-value high maximum-value;
    destination-port port-number;
    destination-prefix-list list-name;
    source-address address;
    source-address-range low minimum-value high maximum-value;
    source-prefix-list list-name;
  }
  then {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
    reflexive; | revert; | reverse {
      application-profile profile-name;
      dscp (alias | bits);
      forwarding-class class-name;
    }
  }
}
```

Hierarchy Level

```
[edit services cos rule rule-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a policy in a services CoS rule. The policy specifies Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignment for packets that are processed by a service set. The policy identifies the matching conditions for packet source and destination addresses and for packet applications, and the actions to take on those packets. A CoS rule can include multiple policies.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

policy-name—Name of the policy.

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

| [Class of Service for Services PICs \(Next Gen Services\)](#) | 255

policy (Stateful Firewall Rules Next Gen Services)

Syntax

```
policy policy-name {
  match {
    application [application-name];
    destination-address (address | any);
    destination-address-excluded address;
    source-address (address | any);
    source-address-excluded address;
  }
  then {
    count;
    deny;
    permit;
    reject;
  }
}
```

Hierarchy Level

```
[edit services policies stateful-firewall-rule rule-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure one or more policies in a stateful firewall rule. Each policy identifies the matching conditions for a flow, and whether or not to allow the flow. Once a policy in the rule matches a flow, that policy is applied and no other policies in the rule are processed.

Options

policy-name—Name of the policy.

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 Stateful Firewalls for Next Gen Services](#) | 267

pool (Destination NAT Next Gen Services)

Syntax

```
pool nat-pool-name{  
    address address-prefix;  
}
```

Hierarchy Level

```
[edit services nat destination]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a set of addresses used for Network Address Translation (NAT) of destination addresses.

Options

nat-pool-name—Name of the NAT pool.

If you are configuring twice NAT, do not use the same name that you use for the source pool.

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.

pool (Source NAT Next Gen Services)

Syntax

```

pool nat-pool-name {
  address address-prefix | address address-prefix to address address-prefix;
  address-pooling {
  }
  ei-mapping-timeout ei-mapping-timeout;
  host-address-base ip-address;
  mapping-timeout mapping-timeout;
  pool-utilization-alarm {
    clear-threshold value;
    raise-threshold value;
  }
  port {
    automatic (random-allocation | round-robin);
    block-allocation {
      active-block-timeout timeout-interval;
      block-size block-size;
      interim-logging-interval timeout-interval;
      maximum-blocks-per-host maximum-block-number
    }
    deterministic {
      block-size block-size;
      host {
        address address;
      }
      include-boundary-addresses;
    }
    deterministic-nat-configuration-log-interval seconds;
    no-translation;
    preserve-range;
    preserve-parity;
    range {
      port-low to port-high;
      (random-allocation | round-robin);
    }
  }
}

```

Hierarchy Level

[edit services nat source]

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a set of addresses (or prefixes), address ranges, and ports used for Network Address Translation (NAT) of source addresses.

Options

nat-pool-name—Name of the NAT pool.

If you are configuring twice NAT, do not use the same name that you use for the destination pool.

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.

pool (NAT Rule Next Gen Services)

Syntax

```
pool nat-pool-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name then source-nat],  
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the name of the NAT pool that contains the addresses or subnets to which addresses are translated.

Required Privilege Level

interface—To view this statement in the configuration.

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

pool-default-port-range (Source NAT Next Gen Services)

Syntax

```
pool-default-port-range port-low to port-high;
```

Hierarchy Level

```
[edit services nat source]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a global default port range for NAT pools that use port translation. This port range is used when a NAT pool does not specify a port range and does not specify automatic port assignment.

Options

port-low—The lower end of the port range.

port-high—The upper end of the port range.

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.

pool-utilization-alarm (Source NAT Next Gen Services)

Syntax

```
pool-utilization-alarm {  
    clear-threshold value;  
    raise-threshold value;  
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Define the NAT pool utilization level that triggers SNMP traps and the pool utilization level that clears SNMP traps. For pools that use port-block allocation, the utilization is based on the number of ports that are used; for pools that do not use port-block allocation, the utilization is based on the number of addresses that are used.

If you do not configure **pool-utilization-alarm**, traps are not created.

Options

clear-threshold *value*—NAT pool utilization percentage that clears the trap.

Range: 40 through 100

Default: 0 (traps are not created)

raise-threshold *value*—NAT pool utilization percentage that triggers the trap.

Range: 50 through 100

Default: There is not default value. Traps are not raised if you do not configure a value.

Required Privilege Level

interface—To view this statement in the configuration.

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

port (Source NAT Next Gen Services)

Syntax

```
port {
  automatic (random-allocation | round-robin);
  block-allocation {
    active-block-timeout timeout-interval;
    block-size block-size;
    interim-logging-interval timeout-interval;
    maximum-blocks-per-host maximum-block-number
  }
  deterministic {
    block-size block-size;
    host {
      address address;
    }
    include-boundary-addresses;
  }
  deterministic-nat-configuration-log-interval seconds;
  no-translation;
  preserve-range;
  preserve-parity;
  range {
    port-low to port-high;
    (random-allocation | round-robin);
  }
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure port assignment for a source NAT pool.

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.

port-forwarding (Destination NAT Next Gen Services)

Syntax

```
port-forwarding map-name {  
    destined-port port-id translated-port port-id;  
}
```

Hierarchy Level

```
[edit services nat destination]
```

Description

Configure a port forwarding map, which translates the original destination port of a packet to a different port. This translation is a static, one-to-one mapping.

Port forwarding allows a packet to reach a host within a masqueraded, typically private, network, based on the port number on which the packet was received from the originating host. An example of this type of destination is the host of a public HTTP server within a private network.

Options

map-name—Name of the port forwarding map.

destined-port port-id—Original destination port number.

translated-port port-id—Port number to which the original port is mapped.

Required Privilege Level

interface—To view this statement in the configuration.

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

port-forwarding-mappings (Destination NAT Rule Next Gen Services)

Syntax

```
port-forwarding-mappings map-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set-name rule rule-name then]
```

Description

Specify the name of the port-forwarding map that the NAT rule uses to translate the original destination port of a packet to a different port.

Required Privilege Level

interface—To view this statement in the configuration.

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

port-round-robin (Source NAT Next Gen Services)

Syntax

```
port-round-robin {  
    disable;  
}
```

Hierarchy Level

```
[edit services nat source]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Disable round-robin port allocation for any NAT pools that do not specify an **automatic (random-allocation | round-robin)** setting at the **[edit services nat source pool nat-pool-name port]** hierarchy level. The **automatic (random-allocation | round-robin)** setting for a pool overrides the **port-round-robin disable** setting.

Required Privilege Level

interface—To view this statement in the configuration.

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

ports-per-session

Syntax

```
ports-per-session ports;
```

Hierarchy Level

```
[edit services nat pool nat-pool-name pgcp]
```

Release Information

Statement introduced in Junos OS Release 8.4.

Description

Configure the number of ports required to support Real-Time Transport Protocol (RTP), Real-Time Control Protocol (RTCP), Real-Time Streaming Protocol (RTSP), and forward error correction (FEC) for voice and video flows on the Multiservices PIC.

Options

number-of-ports—Number of ports to enable: 2 or 4 for combined voice and video services.

Default: 2

Required Privilege Level

interface—To view this statement in the configuration.

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

preserve-parity (Source NAT Next Gen Services)

Syntax

```
preserve-parity;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Assign a port with the same parity (even or odd) as the incoming source port. This feature is not available if you configure port-block allocation, and is not available for deterministic NAT.

Required Privilege Level

interface—To view this statement in the configuration.

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

preserve-range (Source NAT Next Gen Services)

Syntax

```
preserve-range;
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

For source NAT with port translation, except for deterministic NAT, assign a port within the same range as the incoming port—either 0 through 1023 or 1024 through 65,535. This feature is not available if you configure port block allocation.

Required Privilege Level

interface—To view this statement in the configuration.

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

profile (Traffic Load Balancer)

Syntax

```
profile profile-name {
  custom {
    cmd priority {
      default-real-service-status (down | up);
      expect (ascii | binary) receive-string;
      port port;
      real-service-action (down | up);
      send (ascii | binary) send-string;
    }
    protocol (tcp | udp);
  }
  failure-retries number-of-retries;
  http {
    host hostname;
    method (get | option);
    port http-port-number;
    url url;
  }
  icmp;
  probe-interval interval;
  recovery-retries number-of-recovery-retries;
  ssl-hello {
    port port;
    ssl-version;
  }
  tcp {
    port tcp-port-number;
  }
}
```

Hierarchy Level

```
[edit services network-monitoring]
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure a monitoring profile that can be used for health-checking a group of TLB servers.

Options

custom—Use custom probes for server health checking.

cmd *priority*—Use the specified command priority to send for a custom probe.

Values: 1 or 2

default-real-service-status (*down* | *up*)—Assign a server status for when the probe times out. The **up** value is used when the server or the intermediate network nodes are only expected to send a negative response to a probe.

Default: down

expect (*ascii* | *binary*) *receive-string* —Use the specified *ascii* or *binary* string as an expected probe response.

Range: 1 through 512 characters

port *port*—Use the specified port for custom probes.

protocol (*tcp* | *udp*)—Use the selected protocol for custom probes.

real-service-action (*down* | *up*)—Assign a server status for when the expected response to the probe is received.

Default: down

send (*ascii* | *binary*) *send-string* —Send the specified *ascii* or *binary* string as a probe.

Range: 1 through 512 characters

failure-retries *number-of-retries*—Use the specified number of probes that are sent after which the real server is tagged as down.

Default: 5

http—Use HTTP probes for server health checking.

host *hostname*—Use the specified hostname for HTTP probes for server health checks.

method (*get* | *option*)—Use the *get* or *option* HTTP method for server health checks.

port *http-port-number*—Use the specified port number for HTTP probes.

url *url*—Use the specified URL for HTTP probes. Maximum length is 128 bytes.

icmp—Use ICMP probes for server health checking.

probe-interval *interval*—Use the specified interval of time, in seconds, at which health check probes are sent.

Default: 5

profile-name—Identifier for the network monitoring profile.

recovery-retries *number-of-recovery-retries*—Use the specified number of successful probe attempts after which the server is declared up.

Default: 5

ssl-hello—Use a **Client Hello** for server health checks

port *port*—Use the specified port number for **Client Hello** server health checks.

ssl-version—SSL version.

Default: 3

tcp—Use TCP probes for server health checks.

port *tcp-port-number*—Use the specified port number for TCP probes.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Traffic Load Balancer Overview](#) | 287

[Configuring TLB](#) | 296

profile (Web Filter)

Syntax

```

profile profile-name {
  dns-filter {
    database-file filename;
    dns-resp-ttl seconds;
    dns-server [ ip-address ];
    hash-key key-string;
    hash-method hash-method-name;
    statistics-log-timer minutes;
    wildcarding-level level;
  }
  dns-filter-template template-name {
    client-interfaces [ client-interface-name ];
    client-routing-instance client-routing-instance-name;
    dns-filter {
      database-file filename;
      dns-resp-ttl seconds;
      dns-server [ ip-address ];
      hash-key key-string;
      hash-method hash-method-name;
      statistics-log-timer minutes;
      wildcarding-level level;
    }
    server-interfaces [ server-interface-name ];
    server-routing-instance server-routing-instance-name;
    term term-name {
      from {
        src-ip-prefix [ source-prefix ];
      }
      then {
        accept;
        dns-sinkhole;
      }
    }
  }
}

global-dns-stats-log-timer minutes;
url-filter-database filename;
(url-filter-template | template) template-name {
  client-interfaces [ client-interface-name1 client-interface-name2 ];
  disable-url-filtering;
  dns-resolution-interval minutes;
  dns-resolution-rate seconds;
}

```

```

dns-retries number;
dns-routing-instance dns-routing-instance-name;
dns-server [ ip-address1 ip-address2 ip-address3 ];
dns-source-interface loopback-interface-name;
dns-routing-instance dns-routing-instance-name;
routing-instance routing-instance-name;
server-interfaces [ server-interface-name1 server-interface-name2 ];
term term-name {
    from {
        src-ip-prefix [prefix1 prefix2];
        dest-port [port1 port2];
    }
    then {
        accept;
        custom-page custom-page;
        http-status-code http-status-code;
        redirect-url redirect-url;
        tcp-reset;
    }
}
url-filter-database filename
}

```

Hierarchy Level (starting in Junos OS Release 18.3R1)

```
[edit services web-filter]
```

Hierarchy Level (before Junos OS Release 18.3R1)

```
[edit services url-filter]
```

Release Information

Statement introduced in Junos OS Release 17.2.

dns-filter, **dns-filter-templates**, **global-dns-stats-log-timer**, and **url-filter-template** options introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Define URL filter profile or DNS filter profile.

A URL filter profile is for filtering access to disallowed URLs. A URL filter profile includes a general database setting and templates. The template settings apply to specific interfaces or to access from specific source IP address prefixes, and override the database setting at the profile level.

A DNS filter profile is used to filter DNS requests for disallowed website domains. A DNS filter profile includes general DNS filtering settings and up to 32 templates. The template settings apply to DNS requests on specific interfaces or to DNS requests from specific source IP address prefixes, and override the corresponding settings at the profile level. You can configure up to eight DNS filter profiles.

NOTE: For URL filtering, use the **url-filter-template** option starting in Junos OS Release 18.3R1 and use the **template** option in Junos OS Releases before 18.3R1.

Options

profile-name—Name of the filter profile.

url-filter-database filename—Specify the filename of the URL filter database. This option is mandatory.

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

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[DNS Request Filtering for Disallowed Website Domains | 311](#)

[Configuring URL Filtering | 335](#)

protocol (Applications)

Syntax

```
protocol type;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Networking protocol type or number.

Options

type—Networking protocol type. The following text values are supported:

ah

egp

esp

gre

icmp

icmp6

igmp

ipip

ospf

pim

rsvp

tcp

udp

NOTE: IP version 6 (IPv6) is not supported as a network protocol in application definitions.

Required Privilege Level

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

RELATED DOCUMENTATION

<i>ALG Descriptions</i>
Configuring Application Sets 421
<i>Configuring Application Properties</i>
Examples: Configuring Application Protocols 437
Verifying the Output of ALG Sessions 438

range (Source NAT Next Gen Services)

Syntax

```
range {
    port-low to port-high;
    (random-allocation | round-robin);
}
```

Hierarchy Level

```
[edit services nat source pool nat-pool-name port]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

To configure a range of ports to assign to a pool, specify the low and high values for the port. If you do not configure automatic port assignment, you must configure a range of ports. This statement applies to source NAT with port translation, but not to deterministic NAT.

If you specify a range, ports are selected a round-robin fashion. If you specify a range of ports to assign, the automatic statement is ignored.

Options

port-low—Lowest port number.

port-high—Highest port number.

random-allocation—Randomly assigns a port from the range 1024 through 65535 for each port translation.

round-robin—First assigns port 1024, and uses the next higher port for each successive port assignment. Round robin allocation is the default.

Required Privilege Level

interface—To view this statement in the configuration.

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

rate

Syntax

```
rate new-sessions-per-second;
```

Hierarchy Level

```
[edit interfaces interface-name services-options session-limit]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the maximum number of new sessions allowed per second on services cards.

Options

rate *new-sessions-per-second*—Specify the maximum number of new sessions allowed per second.

Range: 0, which indicates no limit, or greater.

Required Privilege Level

interface—To view this statement in the configuration.

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

real-service (Traffic Load Balancer)

Syntax

```
real-service real-service-name {  
    address server-ip-address;  
    admin-down;  
}
```

Hierarchy Level

```
[edit services traffic-load-balance instance instance-name]
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure a traffic load balancer server.

Options

admin-down—Set a server's status to Down.

real-service-name—Identifier for a server to which sessions can be distributed using the server distribution table in conjunction with the session distribution API.

server-ip-address—IP address for the server.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Traffic Load Balancer Overview](#) | 287

[Configuring TLB](#) | 296

reassembly-timeout

Syntax

```
reassembly-timeout seconds;
```

Hierarchy Level

```
[edit interfaces interface-name services-options]  
[edit security flow]
```

Release Information

Statement introduced in Junos OS Release 12.1.

Statement added in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

Description

The maximum acceptable time, in seconds, from the receipt of the first and latest fragments in a packet. When the number is exceeded, the packet is dropped.

Options

seconds—Maximum seconds allowed.

Range: 1 to 60 seconds.

Default: 4 seconds.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring Fragmentation Control for MS-DPC and MS-PIC Service Interfaces*

record-route-option (IDS Screen Next Gen Services)

Syntax

```
record-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IPv4 packets that have the IP option of 7 (Record Route).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

redistribute-all-traffic (Aggregated Multiservices)

Syntax

```
redistribute-all-traffic {
  enable-rejoin;
}
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Enable the option to redistribute traffic of a failed active member to the other active members.

For many-to-one (N:1) high availability support for Network Address Translation (NAT), the traffic for the failed member is automatically redistributed to the other active members.

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

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

[member-failure-options \(Aggregated Multiservices\)](#) | 596

redundancy-event (Services Redundancy Daemon)

Syntax

```
redundancy-event event-name {
  monitor {
    <link-down interface-name>
    <peer {
      (mastership-acquire | mastership-release);
    }>
    <process routing abort>;
    <process routing restart>;
  }
}
```

Hierarchy Level

```
[edit services event-options]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure events monitored to trigger change of mastership and routing using inter-chassis redundancy.

Options

event-name—Alphanumeric name for a monitored event.

link-down interface-name—Name of an interface, link, or link aggregation, to monitor.

peer mastership-acquire—(Optional) Monitor mastership acquisition peer events.

peer mastership-release—(Optional) Monitor mastership release peer events.

process routing abort—(Optional, and only applies to Next Gen Services) Monitor process routing daemon (rpd) abort requests.

process routing restart—(Optional) Monitor process routing daemon (rpd) restart requests.

Required Privilege Level

maintenance—To view this statement in the configuration.

maintenance-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Inter-Chassis Services Redundancy for Next Gen Services | 391](#)

Configuring the Service Redundancy Daemon

redundancy-options (Aggregated Multiservices)

Syntax

```
redundancy-options {
  primary mams-a/b/0;
  secondary mams-a/b/0;
}
```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 17.2 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure warm standby for an aggregated multiservices (AMS) interface. Specify a primary and a secondary (backup) member services interface for the AMS interface. The primary interface is the service interface that you want to back up, and it is the active interface unless it fails. The secondary interface is the backup interface, and does not handle any traffic unless the primary interface fails. You can use the same services interface as the backup in multiple warm standby AMS interfaces.

You cannot use both the **redundancy-options** and the **load-balancing-options** statements in the same AMS interface.

Options

primary mams-a/b/0—Name of the primary services interface, where *a* is the FPC slot number and *b* is the PIC slot number.

secondary mams-a/b/0—Name of the secondary (backup) services interface, where *a* is the FPC slot number and *b* is the PIC slot number.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Warm Standby for Services Interfaces](#) | 365

redundancy-options (Stateful Synchronization)

Syntax

```

redundancy-options {
  redundancy-local {
    data-address address;
  }
  redundancy-peer {
    ipaddress address;
  }
  replication-threshold seconds;
  routing-instance instance-name;
  apply-groups (apply-groups-except | redundancy-local | redundancy-peer)
  replication-options (apply-groups | apply-groups-except | mtu | replication-threshold | replication-threshold
    routing-instance )
}

```

Hierarchy Level

```
[edit interfaces interface-name]
```

Release Information

Statement introduced in Junos OS Release 13.3.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card (interfaces of type **vms-x/y/z**).

Description

Specify the primary and secondary (backup) adaptive services PIC interfaces.

Options

data-address *address*—Internal IP address of the local redundant PIC.

ipaddress *address*—Internal IP address of the remote redundant PIC.

instance-name—Name of the routing instance to apply to the HA synchronization traffic between the high availability pair.

seconds—Length of time that the flow remains active for replication.

Default: 180 seconds

apply-groups *apply-groups-except*—Specify the groups from which NOT to inherit the configuration.

apply-groups *redundancy-local*—Specify information for the local peer.

apply-groups *redundancy-peer*—Specify information for peer.

replication-options *apply-groups*—Specify groups from which to inherit the configuration.

replication-options *apply-groups-except*—Specify the groups from which NOT to inherit the configuration.

replication-options *mtu*—Specify the maximal packet size for the replicated data.

Range: 1500 through 8000 bytes

replication-options *replication-threshold*—Specify the duration for which flow should remain active for replication.

Range: 60 through 3600 seconds

replication-options *replication-threshold routing-instance*—Specify routing-instance for the HA traffic.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Inter-Chassis Stateful Synchronization for Long Lived NAT and Stateful Firewall Flows (MS-MPC, MS-MIC) (Release 16.1 and later)

Inter-Chassis High Availability for MS-MIC and MS-MPC (Release 15.1 and earlier)

redundancy-policy (Interchassis Services Redundancy)

Syntax

```

redundancy-policy policy-name {
  redundancy-events [event-list] {
    then {
      acquire-mastership;
      <add-static-route destination {
        (next-hop next-hop | receive);
        routing-instance routing-instance
      }>
      <broadcast-warning> ;
      <delete-static-route destination {
        routing-instance routing-instance;
      }>
      <(release-mastership | release-mastership-force);>
    }
  }
}

```

Hierarchy Level

[edit policy-options]

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the actions to be taken for redundancy events. These include acquiring or releasing mastership and adding or deleting static routes.

Options

acquire-mastership—Switch from standby to master role.

add-static-route *destination*—(Optional) Use the specified destination IP address and prefix for an added signal route.

broadcast-warning—(Optional) Switch status from Standby to Standby (Warned).

delete-static-route *destination*—(Optional) Use the specified destination IP address and prefix for a deleted signal route.

event-list—List of names of one or more monitored events that trigger the actions specified in this policy.

next-hop—Interface name for the next hop for an added signal route.

policy-name—Name of the redundancy policy.

receive—Use the added signal route as a receive route.

release-mastership—(Optional) Switch from master to standby role.

release-mastership-force—(Optional) Force switch from master to standby role.

routing-instance *routing-instance*—(Optional) Name of the vrf used for the added signal route.

Required Privilege Level

maintenance—To view this statement in the configuration.

maintenance-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Inter-Chassis Services Redundancy for Next Gen Services | 391](#)

Configuring the Service Redundancy Daemon

redundancy-set

Syntax

```
redundancy-set redundancy-set {
  healthcheck-timer-interval healthcheck-timer-interval;
  hold-time hold-time;
  keepalive keepalive;
  redundancy-group redundancy-group;
  redundancy-policy [redundancy-policy-list]
}
```

Hierarchy Level

```
[edit services]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the characteristics of a redundancy set.

Options

healthcheck-timer-interval *healthcheck-timer-interval*—Frequency of health check probes in seconds.

Range: 0 through 3600 seconds

hold-time—Maximum wait time for a health check response. When this time expires, the peer is considered down.

Range: 0 through 3600 seconds

keepalive—Frequency of srd hello messages in seconds.

Range: 1 through 60 seconds

redundancy-group—Redundancy group identifier. This must match a redundancy group ID in the ICCP configuration.

Range: 1 through 100

redundancy-policy-list—Names of one or more redundancy policies applied to the redundancy set.

redundancy-set—Redundancy set identifier.

Range: 1 through 100

Required Privilege Level

maintenance—To view this statement in the configuration.

maintenance-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring Inter-Chassis Services Redundancy for Next Gen Services | 391](#)

Configuring the Service Redundancy Daemon

redundancy-set-id (Service Set)

Syntax

```
redundancy-set-id redundancy-set;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the identifier of the redundancy set to use in the stateful synchronization of services for a service set.

Options

redundancy-set—Identifier for the redundancy set. The identifier can be a number from 1-100.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Inter-Chassis Services Redundancy for Next Gen Services | 391](#)

Configuring the Service Redundancy Daemon

rejoin-timeout (Aggregated Multiservices)

Syntax

```
rejoin-timeout rejoin-timeout;
```

Hierarchy Level

```
[edit interfaces interface-name load-balancing-options member-failure-options drop-member-traffic]
```

Release Information

Statement introduced in Junos OS Release 11.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure the time by when failed members (members in the **DISCARD** state) should rejoin the aggregated Multiservices (AMS) interface automatically. All members that do not rejoin by the configured time are moved to the **INACTIVE** state and the traffic meant for each of the members is dropped.

If multiple members fail around the same time, then they are held in the **DISCARD** state using a single timer. When the timer expires, all the failed members move to **INACTIVE** state at the same time.

Default

If you do not configure a value, the default value of 120 seconds is used.

Options

rejoin-timeout—Time, in seconds, by which a failed member must rejoin.

Default: 120 seconds

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

Example: Configuring an Aggregated Multiservices Interface (AMS)

[drop-member-traffic \(Aggregated Multiservices\)](#) | 521

rpc-program-number

Syntax

```
rpc-program-number number;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Remote procedure call (RPC) or Distributed Computing Environment (DCE) value.

Options

number—RPC or DCE program value.

Range: 100,000 through 400,000

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

<i>ALG Descriptions</i>
<i>Configuring an RPC Program Number</i>
Examples: Configuring Application Protocols 437
Verifying the Output of ALG Sessions 438

rtlog (Next Gen Services Global System Logging)

Syntax

```
rtlog {
  name {
    apply-groups group-names;
    apply-groups-except group-names;
    flag name;
    file filename,
    no-remote-trace;
  }
}
```

Hierarchy Level

[edit services]

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Enable global system logging for Next Gen Services.

traceoptions—Specify the options to include in the trace.

All other options are explained separately.

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

[traceoptions | 734](#)

rule (Destination NAT Next Gen Services)

Syntax

```
rule rule-name {
  match {
    application [application-name]
    destination-address (NAT Next Gen Services) (address | any-unicast);
    destination-address-name address-name;
    source-address (address | any-unicast);
    source-address-name address-name;
  }
}
then {
  destination-nat {
    destination-prefix destination-prefix;
    off;
    pool nat-pool-name;
  }
  port-forwarding-mappings map-name;
}
syslog;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a destination NAT rule, which translates the destination address of IP 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.

rule (Services CoS Next Gen Services)

Syntax

```
rule rule-name {
  match-direction (input | input-output | output);
  policy policy-name {
    match {
      application [ application-names ];
      destination-address address;
      destination-address-range low minimum-value high maximum-value;
      destination-port port-number;
      destination-prefix-list list-name;
      source-address address;
      source-address-range low minimum-value high maximum-value;
      source-prefix-list list-name;
    }
    then {
      application-profile profile-name;
      dscp (alias | bits);
      forwarding-class class-name;
      reflexive; | revert; | reverse {
        application-profile profile-name;
        dscp (alias | bits);
        forwarding-class class-name;
      }
    }
  }
}
```

Hierarchy Level

```
[edit services cos]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a services CoS rule, which specifies Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignment for packets that are processed by a service set. The CoS rule identifies the matching conditions for packet source and destination addresses and for packet applications, and the actions to take on those packets.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

match-direction (input | input-output | output)—The direction in which the rule is matched.

input—Apply the rule match on input. If the CoS rule is assigned to an interface service set, input means traffic entering the interface. If the CoS rule is assigned to a next-hop service set, input means traffic routed with the inside interface.

input-output—Apply the rule match in both directions.

output—Apply the rule match on output. If the CoS rule is assigned to an interface service set, input means traffic leaving the interface. If the CoS rule is assigned to a next-hop service set, output means traffic routed with the outside interface.

rule-name—Name of the CoS rule.

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

| [Class of Service for Services PICs \(Next Gen Services\)](#) | 255

rule (PCP)

Syntax

```
rule rule-name {
  match-direction (input | output);
  term term-name {
    from {
      application-sets set-name;
      applications [ application-name ];
      destination-address address <except>;
      destination-address-range high maximum-value low minimum-value <except>;
      destination-port high maximum-value low minimum-value;
      destination-prefix-list list-name <except>;
      source-address address <except>;
      source-address-range high maximum-value low minimum-value <except>;
      source-prefix-list list-name <except>;
    }
    then {
      pcg-server server-name;
    }
  }
}
```

Hierarchy Level

```
[edit services pcg]
```

Release Information

Statement introduced in Junos OS Release 13.2R1.

Description

Configure a rule to assign the port control protocol (PCP) server that handles selected traffic. PCP enables hosts to operate servers for a long time (as in the case of a webcam) or a short time (for example, while playing a game or on a phone call) when behind a NAT device, including when behind a carrier-grade NAT operated by their ISP. PCP enables applications to create mappings from an external IP address and port to an internal IP address and port.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC.

Options

rule-name—Rule name

The remaining statements are explained separately.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Port Control Protocol](#) | 188

rule (Source NAT Next Gen Services)

Syntax

```
rule rule-name {
  match {
    application [application-name]
    destination-address (NAT Next Gen Services) address;
    destination-address-name address-name;
    source-address (address | any-unicast);
    source-address-name address-name;
  }
  then {
    source-nat {
      clat-prefix clat-prefix;
      filtering-type {
        endpoint-independent {
          prefix-list [allowed-host] except [denied-host];
        }
      }
      mapping-type {
        endpoint-independent;
      }
      pool nat-pool-name;
      secure-nat-mapping {
        eif-flow-limit number-of-flows;
        mapping-refresh (inbound | inbound-outbound | outbound);
      }
    }
    syslog;
  }
}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a source NAT rule, which translates the source address of IP 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.

rule-set (Services CoS Next Gen Services)

Syntax

```
rule-set rule-set-name {
  [ rule rule-name ];
}
```

Hierarchy Level

```
[edit services cos]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a set of services CoS rules. You can then assign the rule set to a service set, which processes the rules in the order they appear. Once a rule matches the packet, the router performs the corresponding action, and no further rules are applied.

Options

rule *rule-name*—The name of each rule in the rule set.

rule-set-name—The name for the set of rules.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Class of Service for Services PICs \(Next Gen Services\)](#) | 255

rule-set (Softwires Next Gen Services)

Syntax

```
rule-set rule-set-name {
  match-direction (input | output);
  rule rule-name {
    then {
      ds-lite ds-lite-concentrator-name
      map-e map-e-concentrator-name
      v6rd v6rd-software-concentrator;
    }
  }
}
```

Hierarchy Level

[edit services softwires]

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a rule to apply a DS-Lite, MAP-E, or v6rd software concentrator to a flow.

Options

input—Apply the rule on the input side of the interface.

output—Apply the rule on the output side of the interface.

rule rule-name—Name of the rule.

rule-set rule-set-name—Name of the rule set that contains the rule.

ds-lite ds-lite-software-concentrator—Name of the software concentrator that the rule assigns to a flow.

map-e map-e-software-concentrator—Name of the software concentrator that the rule assigns to a flow.

v6rd v6rd-software-concentrator—Name of the software concentrator that the rule assigns to a flow.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[6rd Softwires in Next Gen Services | 165](#)

secure-nat-mapping (Source NAT Next Gen Services)

Syntax

```
secure-nat-mapping {  
    eif-flow-limit number-of-flows;  
    mapping-refresh (inbound | inbound-outbound | outbound);  
}
```

Hierarchy Level

```
[edit services nat source rule-set rule-set rule rule-name then source-nat]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

For endpoint-independent mapping, configure the maximum number of simultaneous inbound flows and the direction in which mappings are refreshed.

Options

eif-flow-limit *number-of-flows*—Maximum number of simultaneous inbound flows.

Range: 0 through 655334

mapping-refresh (inbound | inbound-outbound | outbound)—Direction in which mappings are refreshed.

Required Privilege Level

interface—To view this statement in the configuration.

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

security-intelligence

Syntax

```
authentication {
    auth-token auth-token;
    tls-profile tls-profile;
    traceoptions {
        no-remote-trace;
        file [ filename <files number> <size bytes> <match expression> <world-readable | no-world-readable>];
        flag [all | feed | ipc];
        level [all | error | info | notice | verbose | warning];
        no-remote-trace;
    }
    url url;
```

Hierarchy Level

```
[edit services]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers with Juniper Sky Advanced Threat Prevention (ATP).

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960. This support runs inline on the MPC card.

Description

You can configure security intelligence profiles and policies to work with security intelligence feeds, such as infected hosts and C&C. You then configure a firewall policy to include the security intelligence policy, for example, block outgoing requests to a C&C host.

Options

authentication—Configure authentication, such as an auth token or TLS profile, to commute with the feed server. This operation is performed by the ops script used to enroll your devices and is typically not required afterwards. If you have problems establishing a connection with the Juniper Sky ATP cloud server, we recommend that you rerun the ops script instead of manually entering all the CLI commands.

traceoptions—Set security intelligence trace options.

- **file**—Name of the file to receive the output of the tracing operation.
 - **files *number*** —Maximum number of trace files

Range: 2 through 1000
 - **match**— Regular expression for lines to be logged

- no-world-readable—Prevent any user from reading the log file
- size—Maximum size of each trace file

Range: 10240 through 1073741824

- world-readable—Allow any user to read the log file
- flag—Tracing operation to perform
 - all—All interface tracing operation
 - feed—Trace feed operation
 - ipc—Trace interface interprocess communication (IPC) module messages
- level—Level of debugging output
- no-remote-trace—Disable the remote trace

url *url-address*—Configure the URL of the feed server. This operation is performed by the ops script used to enroll your devices and is typically not required afterwards. If you have problems establishing a connection with the Juniper Sky ATP cloud server, we recommend that you rerun the ops script instead of manually entering all the CLI commands.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

security-intelligence-policy

Syntax

```
security-intelligence-policy {
  threat-level threat-level;
  threat-action {
    drop
    drop-and-log
    drop-and-sample
    drop-log-and-sample
    log
    log-and-sample
    sample
  }
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R1 on MX Series routers with Juniper Sky Advanced Threat Prevention (Juniper Sky ATP) .

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card

Description

Define the threat level and action for the Web filter profile. The packets are redirected at the Packet Forwarding Engine based on the configured threat-level action associated with the threat-level of the destination IP address.

Options

threat-level—Define the Web filtering threat level. The value ranges from 1 through 10

threat-action—Define the way the Packet Forwarding Engine processes packets in response to a threat. Only one action can be configured for each threat level that is defined. The default threat-action is **accept**.

- **drop**—Drop the packets and do not generate a log message.
- **drop-and-log**—Drop the packets and generate a log message.
- **drop-and-sample**—Drop and sample the packets.

- **drop-log-and-sample**—Drop, sample, and allow the packets, and generate a log message.
- **log**—Allow the packets and generate a log message.
- **log-and-sample**—Allow, sample the packets, and generate a log message.
- **sample**—Sample the packets.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [web-filter](#) | 755

security-option (IDS Screen Next Gen Services)

Syntax

```
security-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IPv4 packets that have the IP option of 2 (Security).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

server (pcp)

Syntax

```
server server-name {
    ipv4-address ipv4-address;
    ipv6-address ipv6-address;
    long-lifetime-error long-lifetime-error;
    mapping-lifetime-max mapping-lifetime-max;
    mapping-lifetime-min mapping-lifetime-min;
    max-mappings-per-client max-mappings-per-client;
    nat-options {
        pool pool-name ;
    }
    pcp-options {
        prefer-failure;
        third-party;
    }
    short-lifetime-error short-lifetime-error;
    softwire-concentrator softwire-concentrator-name;
}
```

Hierarchy Level

```
[edit services pcp]
```

Release Information

Statement introduced in Junos OS Release 13.2R1.

Description

Configure PCP server options. PCP enables hosts to operate servers for a long time (as in the case of a webcam) or a short time (for example, while playing a game or on a phone call) when behind a NAT device, including when behind a carrier-grade NAT operated by their ISP. PCP enables applications to create mappings from an external IP address and port to an internal IP address and port.

PCP is supported on the MS-DPC, MS-100, MS-400, and MS-500 MultiServices PICS. Starting in Junos OS Release 17.4R1, PCP is also supported on the MS-MPC and MS-MIC. Starting in Junos OS Release 20.1R1, PCP is also supported for Next Gen Services.

Options

ipv4-address—IPv4 address of the PCP server.

ipv6-address—IPv6 address of the PCP server.

long-lifetime-error—Time limit for generating long lifetime errors.

Default: 1800 seconds

Range: 900 through 18,000 seconds

mapping-lifetime-max—Maximum lifetime, in seconds, for PCP mapping. If the PCP client requests a lifetime less than the maximum configured, the server will assign the maximum lifetime and respond accordingly.

Default: 86,400 seconds

Range: 3600 through 4294667 seconds

mapping-lifetime-min—Minimum lifetime, in seconds, for PCP mapping. If a PCP client requests a lifetime less than the minimum configured, the server will assign a minimum lifetime and respond accordingly.

Default: 300 seconds

Range: 120 through 3600 seconds

max-mappings-per-client—Maximum number of PCP mappings that the PCP client can request.

Default: 32

Range: 1 through 32

pool-name—Name of the NAT pool to use for PCP mapping. You can identify multiple pools. If you do not specify a NAT pool for mapping, the Junos OS performs a partial rule match based on the source IP, source port, and protocol, and the Junos OS uses the NAT pool configured for the first matching rule to allocate mappings for PCP.

prefer-failure—Generate an error message when the PCP client requests a specific IP address or port that is not available, rather than assigning another available address from the NAT pool.

short-lifetime-error—Time limit for generating short lifetime errors.

Default: 30 seconds

Range: 15 through 300 seconds

software-concentrator-name—Software concentrator name whose software-address is used in creating PCP mappings. The PCP server address must be the same as the software-concentrator address.

third-party—Enable third-party requests by the PCP client.

The other 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

service-domain

Syntax

```
service-domain (inside | outside);
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the service interface domain. If you specify this interface using the **next-hop-service** statement at the **[edit services service-set service-set-name]** hierarchy level, the interface domain must match that specified with the **inside-service-interface** and **outside-service-interface** statements.

Options

inside—Interface used within the network.

outside—Interface used outside the network.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| *Configuring the Address and Domain for Services Interfaces*

service-interface (Services Interfaces)

Syntax

```
service-interface interface-name;
```

Hierarchy Level

```
[edit services service-set service-set-name interface-service]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the name for the services interface associated with an interface-wide service set.

Options

interface-name—Identifier of the service interface.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring Service Sets to be Applied to Services Interfaces

Applying Services to Subscriber-Aware Traffic with a Service Set

services-options (Next Gen Services Interfaces)

Syntax

```

services-options {
  enable-subscriber-analysis
  fragment-limit;
  jflow-log {
    message-rate-limit messages-per-second;
  }
  session-limit {
    maximum number;
    rate new-sessions-per-second;
    cpu-load-threshold percentage;
  }
  flow
    traceoptions {
      file {
        filename;
        files number;
        match regular-expression;
        size maximum-file-size;
        (world-readable | no-world-readable);
      }
      flag flag;
      no-remote-trace;
      packet-filter filter-name {
        conn-tag session-conn
        destination-port port-identifier;
        destination-prefix address;
        interface interface-name;
        protocol protocol-identifier;
        source-port port-identifier;
        source-prefix address;
      }
      rate-limit messages-per-second;
      trace-level (brief | detail | error);
    }
  }
}

```

Hierarchy Level

[edit interfaces *interfaces-name*]

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series.

Support introduced in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480 and MX960 routers for the **flow** configuration statement.

Description

Define the service options to be applied on the virtual multi-service (VMS) interface.

This statement is supported only on the MX-SPC3 Services Card.

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

Options

file—Configure the trace file options.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.

files *number*—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000 files

Default: 10 files

match *regular-expression*—Refine the output to include lines that contain the regular expression.

size *maximum-file-size*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: **x K** to specify KB, **x m** to specify MB, or **x g** to specify GB

Range: 0 KB through 1 GB

Default: 128 KB

world-readable | no-world-readable—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

flag—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

all—Trace with all flags enabled

basic-datapath—Trace basic packet flow activity

fragmentation—Trace IP fragmentation and reassembly events

high-availability—Trace flow high-availability information

host-traffic—Trace flow host traffic information

multicast—Trace multicast flow information

route—Trace route lookup information

session—Trace session creation and deletion events

session-scan—Trace session scan information

tcp-basic—Trace TCP packet flow information

tunnel—Trace tunnel information

no-remote-trace—Set remote tracing as disabled.

packet-filter *filter-name*—Packet filter to enable during the tracing operation. Configure the filtering options.

destination-port *port-identifier*—Match TCP/UDP destination port

destination-prefix *address*—Destination IP address prefix

interface *interface-name*—Logical interface

protocol *protocol-identifier*—Match IP protocol type

source-port *port-identifier*—Match TCP/UDP source port

source-prefix *address*—Source IP address prefix

rate-limit *messages-per-second*—Limit the incoming rate of trace messages.

trace-level—Set the level for trace logging. This option is available only when the flag is set.

brief—Trace key flow information, such as message types sent between SPU and central point, policy match, and packet drop reasons.

detail—Trace extensive flow information, such as detailed information about sessions and fragments. Detail is the default level.

error—Trace error information, such as system failure, unknown message type, and packet drop.

fragment-limit—Specify the maximum number of fragments to be supported for the PIC. This overrides the value specified, if any, in the **set security flow fragment-limit** statement.

reassembly-timeout—Specify the reassembly timeout value for all fragmentation packets for the PIC. This overrides the value specified, if any, in the **set security flow reassembly-timeout** statement

Required Privilege Level

interface—To view this statement in the configuration.

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

service-set (Interfaces)

Syntax

```
service-set service-set-name;
```

Hierarchy Level

```
[edit interfaces interface-name unit logical-unit-number family inet service (input | output)],  
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family inet service (input  
| output)]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Define one or more service sets to be applied to an interface. If you define multiple service sets, the router software evaluates the filters in the order in which they appear in the configuration.

Options

service-set-name—Name of the service set.

Required Privilege Level

System—To view this statement in the configuration.

System-control—To add this statement to the configuration.

RELATED DOCUMENTATION

Guidelines for Configuring Service Filters

service-set (Services)

Syntax

```

service-set service-set-name {
    allow-multicast;
    captive-portal-content-delivery-profile;
    cos-options {
        match-rules-on-reverse-flow;
    }
    cos-rules [cos-rule-name];
    extension-service service-name {
        provider-specific-rules-configuration;
    }
    (ids-rules rule-name | ids-rule-sets rule-set-name);
    interface-service {
        load-balancing-options {
            hash-keys {
                egress-key (destination-ip | source-ip);
                ingress-key (destination-ip | source-ip);
            }
        }
        service-interface interface-name;
    }
    ipsec-vpn-options {
        anti-replay-window-size bits;
        clear-dont-fragment-bit;
        ike-access-profile profile-name;
        local-gateway address;
        no-anti-replay;
        no-certificate-chain-in-ike;
        passive-mode-tunneling;
        trusted-ca [ ca-profile-names ];
        tunnel-mtu bytes;
        udp-encapsulation {
            <udp-dest-port destination-port>;
        }
    }
    ip-reassembly-rules rule-name;
    (ipsec-vpn-rules rule-name | ipsec-vpn-rule-sets rule-set-name);
    max-flows number;
    max-drop-flows {
        ingress ingress-flows;
        egress egress-flows;
    }
}

```



```

max-session-setup-rate max-setup-rate;
nat-options {
    land-attack-check (ip-only | ip-port);
    max-sessions-per-subscriber session-number;
    stateful-nat64{
        clear-dont-fragment-bit;
    }
}
(nat-rules rule-name | nat-rule-sets rule-set-name);
next-hop-service {
    inside-service-interface interface-name.unit-number;
    outside-service-interface interface-name.unit-number;
    outside-service-interface-type local;
    service-interface-pool name;
}
pcp-rules rule-name;
(pgcp-rules rule-name | pgcp-rule-sets rule-set-name);
(ptsp-rules rule-name | ptsp-rule-sets rule-set-name);
service-set-options {
    bypass-traffic-on-exceeding-flow-limits;
    bypass-traffic-on-pic-failure;
    disable-session-open-syslog;
    enable-asymmetric-traffic-processing;
    header-integrity-check;
    routing-engine-services;
    support-uni-directional-traffic;
}
snmp-trap-thresholds{
    flows high high-threshold | low low-threshold;
    nat-address-port high-threshold | low low-threshold;
}
}
software-options {
    dslite-ipv6-prefix-length dslite-ipv6-prefix-length;
}
}
(software-rules rule-name | software-rule-sets rule-set-name);
(stateful-firewall-rules rule-name | stateful-firewall-rule-sets rule-set-name);

```

```

syslog {
  host hostname {
    class {
      alg-logs;
      deterministic-nat-configuration-log;
      ids-logs;
      nat-logs;
      packet-logs;
      pcp-logs;
      session-logs <open | close>;
      stateful-firewall-logs ;
    }
    services severity-level;
    facility-override facility-name;
    interface-service prefix-value;
    port port-number;
    services severity-level;
  }
}
(web-filter-profile | url-filter-profile) profile-name;
}

```

Hierarchy Level

[edit services]

Release Information

Statement introduced before Junos OS Release 7.4.

pcp-rules option added in Junos OS Release 13.2R1.

pgcp-rules and **pgcp-rule-sets** options added in Junos OS Release 8.4.

server-set-options option added in Junos OS Release 10.1.

ptsp-rules and **ptsp-rule-sets** options added in Junos OS Release 10.2.

software-rules and **clear-rule-sets** options added in Junos OS Release 10.4.

software-options option added in Junos OS Release 14.1.

url-filter-profile option added in Junos OS Release 17.2R1.

match-rules-on-reverse-flow option added in Junos OS Release 16.1R5 and 17.4R1

web-filter-profile option added in Junos OS Release 18.3R1.

Support added in Junos 20.2R1 for Next Gen Services NAT PT feature.

Description

Define the service set.

NOTE: Use the **web-filter-profile** option starting in Junos OS Release 18.3R1 and use the **url-filter-profile** option in Junos OS Releases before 18.3R1.

Options

service-set-name—Name of the service set. You can include special characters, such as a forward slash (/), colon (:), or a period (.).

Range: Up to 64 alphanumeric characters.

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

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| *Understanding Service Sets*

service-set-options (Next Gen Services Services)

Syntax

```
service-set-options {
  bypass-traffic-on-exceeding-flow-limits;
  disable-global-timeout-override;
  disable-session-open-syslog ;
  enable-asymmetric-traffic-processing;
  inactivity-non-tcp-timeout ;
  max-sessions-per-subscriber
  session-limit;
  session-timeout;
  tcp-session {
    inactivity-asymm-tcp-timeout;
    inactivity-tcp-timeout ;
    open-timeout ;
    tcp-fast-open ;
    tcp-mss ;
    tcp-non-syn ;
    tcp-tickles ;
  }
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Specify the service set options to apply to a service set.

disable-session-open-syslog—Disable session open information from being collected in system logs.

inactivity-non-tcp-timeout —Specify the inactivity timeout period for non-TCP established sessions.

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 Service Sets to be Applied to Services Interfaces

Configuring APPID Support for Unidirectional Traffic

session-limit

Syntax

```
session-limit {  
    maximum number;  
    rate new-sessions-per-second;  
    cpu-load-threshold percentage;  
}
```

Hierarchy Level

```
[edit interfaces interface-name services-options]
```

Release Information

Statement introduced in Junos OS Release 9.6.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Restrict the maximum number of sessions and the session rate on services cards.

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

session-limit (Service Set Next Gen Services)

Syntax

```
session-limit {  
    maximum number;  
}
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the maximum number of sessions allowed simultaneously on the service set. If you specify the maximum number of sessions to be zero, it indicates that the configuration is not effective. You must specify a value higher than zero for the maximum number of sessions.

Options

number—Maximum number of sessions.

Range: 1 through 4,294,967,295

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

session-timeout (Service Set Next Gen Services)

Syntax

```
session-timeout seconds;
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series.

Description

Define session lifetime for the service set in seconds. The session is closed after this amount of time, even if traffic is running on the session.

Options

seconds—Duration of session.

Range: 4 through 86,400

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

severity (Next Gen Services Service-Set Remote System Logging)

Syntax

```
severity severity;
```

Hierarchy Level

```
edit services service-set name syslog stream stream-name
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Specify the level of severity for the stream.

You can set the following severity levels:

- ANY — Includes all severity levels
- ALERT — Action must be taken immediately
- CRITICAL — Critical conditions
- EMERGENCY — System is unusable
- ERROR — Error conditions
- WARNING — Warning conditions
- NOTICE — Normal but significant condition
- INFO — Informational
- DEBUG — Debug-level messages

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

sip (Services CoS Next Gen Services)

Syntax

```
sip {
  data {
    dscp (alias | bits);
    forwarding-class class-name;
  }
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure CoS actions for SIP traffic in an application profile. The application profile can then be used in CoS rule actions.

Options

dscp (*alias* | *bits*)—Either a code point alias or a DSCP bit value to apply to the SIP packets.

forwarding-class *class-name*—Forwarding class name to apply to the SIP packets. The choices are:

- 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

[Class of Service for Services PICs \(Next Gen Services\)](#) | 255

size (Next Gen Services Global System Logging)

Syntax

```
size size;
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Maximum trace file size

Options

size—Maximum trace file size

Default: 128k

Range: through

Required Privilege Level

system

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging	71
Enabling Global System Logging for Next Gen Services	73
Configuring System Logging to One or More Remote Servers for Next Gen Services	76
Configuring Local System Logging for Next Gen Services	74

snmp-command

Syntax

```
snmp-command command;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

SNMP command format.

Options

command—Supported commands are SNMP **get**, **get-next**, **set**, and **trap**.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

ALG Descriptions

Configuring an SNMP Command for Packet Matching

[Examples: Configuring Application Protocols](#) | [437](#)

[Verifying the Output of ALG Sessions](#) | [438](#)

snmp-trap-thresholds (Next Gen Services)

Syntax

```
snmp-trap-thresholds {  
    flow high percent low percent;  
    nat-address-port high percent low percent;  
    session high percent low percent;  
}
```

Hierarchy Level

```
[edit services service-set]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Define snmp traps for Next Gen Services service sets.

Options

session—Specify the low and high session threshold limits for generating SNMP traps.

The default for high = 90%.

The default for low = 70%.

Required Privilege Level

system

software-name (Next Gen Services)

Syntax

```
software-name v6rd-software-concentrator {
  ipv4-prefix ipv4-prefix;
  mtu-v4 number-of-bytes;
  software-concentrator address;
  software-type v6rd;
  v6rd-prefix v6rd-prefix
}
```

Hierarchy Level

```
[edit services softwires]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a 6rd software concentrator. A 6rd software allows an IPv6 end user to send traffic over an IPv4 network to reach an IPv6 network. The software concentrator decapsulates IPv6 packets that were encapsulated in IPv4 packets by a software initiator at the customer edge WAN, and forwards the packets for IPv6 routing.

Options

ipv4-prefix *ipv4-prefix*—IPv4 prefix of the customer edge (CE) network.

mtu-v4 *number-of-bytes*—The size, in bytes, of the maximum transmission unit for IPv6 packets encapsulated in IPv4. Compute this as the maximum expected IPv4 packet size plus 20. Packets that are larger than the configured value are dropped.

Range: 576 through 9192

software-concentrator *address*—IPv4 address of a software concentrator. This is an IPv4 address independent of any interface and on a different prefix.

software-name *v6rd-software-concentrator*—Name of the software concentrator.

software-type *v6rd*—Sets software concentrator type to 6rd.

v6rd-prefix *v6rd-prefix*—IPv6 prefix for the 6rd domain.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [6rd Softwires in Next Gen Services](#) | 165

softwires (Next Gen Services)

Syntax

```

softwires {
  rule-set name {
    match-direction (input | output);
    rule name {
      then {
        (ds-lite ds-lite | map-e map-e | v6rd v6rd);
      }
    }
  }
  software-name name {
  }
  software-types {
  }
  traceoptions {
    file <filename> <files files> <match match> <size size> <(world-readable | no-world-readable)>;
    flag name;
    no-remote-trace;
  }
}

```

Hierarchy Level

[edit services]

Release Information

Statement introduced in Junos OS Release 20.2 for Next Gen Services.

Description

Configure software feature

Required Privilege Level

system

software-name (Next Gen Services)

Syntax

```
software-name v6rd-software-concentrator {
  ipv4-prefix ipv4-prefix;
  mtu-v4 number-of-bytes;
  software-concentrator address;
  software-type v6rd;
  v6rd-prefix v6rd-prefix
}
```

Hierarchy Level

```
[edit services softwires]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure a 6rd software concentrator. A 6rd software allows an IPv6 end user to send traffic over an IPv4 network to reach an IPv6 network. The software concentrator decapsulates IPv6 packets that were encapsulated in IPv4 packets by a software initiator at the customer edge WAN, and forwards the packets for IPv6 routing.

Options

ipv4-prefix *ipv4-prefix*—IPv4 prefix of the customer edge (CE) network.

mtu-v4 *number-of-bytes*—The size, in bytes, of the maximum transmission unit for IPv6 packets encapsulated in IPv4. Compute this as the maximum expected IPv4 packet size plus 20. Packets that are larger than the configured value are dropped.

Range: 576 through 9192

software-concentrator *address*—IPv4 address of a software concentrator. This is an IPv4 address independent of any interface and on a different prefix.

software-name *v6rd-software-concentrator*—Name of the software concentrator.

software-type *v6rd*—Sets software concentrator type to 6rd.

v6rd-prefix *v6rd-prefix*—IPv6 prefix for the 6rd domain.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [6rd Softwires in Next Gen Services](#) | 165

software-options

Syntax

```
software-options {  
  dslite-ipv6-prefix-length dslite-ipv6-prefix-length ;  
}
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 14.1.

Support added in Junos OS 20.2R1 for Next Gen Services on MX240, MX480, and MX960 routers.

Description

Specify the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions.

This feature is supported on MX Series routers equipped with MS-DPCs. Starting in Junos OS Release 18.2R1, this option is also supported on MS-MPCs and MS-MICs.

Options

dslite-ipv6-prefix-length—Subnet prefix representing the size of the subnet subject to session limitation.

Values: 56, 64, 96, 128

Default: 0—no limitation.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[DS-Lite Per Subnet Limitation Overview](#) | 180

software-types (Next Gen Services)

Syntax

```
software-types {
  ds-lite ds-lite-software-concentrator {
    auto-update-mtu;
    flow-limit flow-limit | session-limit-per-prefix session-limit-per-prefix;
    mtu-v6 bytes;
    software-address address;
  }
  map-e
  v6rd v6rd-software-concentrator {
    ipv4-prefix ipv4-prefix;
    v6rd-prefix ipv6-prefix;
    mtu-v4 mtu-v4;
  }
}
```

Hierarchy Level

[edit services softwires]

Release Information

Statement introduced in Junos OS Release 20.2 for Next Gen Services on MX240, MX480 and MX960.

Description

Configure ds-lite, 6rd and MAP-E software objects.

Options

ds-lite —Specify options for DS-Lite softwires.

v6rd —Specify options for v6rd softwires.

map-e — Specify options for map-e softwires.

The following options are available for each type of software:

auto-update-mtu—This option is not currently supported.

copy-dscp —Copy DSCP information to IPv4 headers during decapsulation.

flow-limit —Maximum number of IPv4 flows per software.

ipv4-prefix —IPv4 prefix of the customer edge (CE) network

mtu-v4—Maximum transmission unit (MTU), in bytes (576 through 9192), for IPv6 packets encapsulated into IPv4. If the final length is greater than the configured value, the IPv4 packet is dropped. This option is mandatory except for DS-Lite softwires since it depends on other network parameters under administrator control.

mtu-v6—Maximum transmission unit when encapsulating IPv4 packets into IPv6. If the final length is greater than the MTU, the IPv6 packet is fragmented. This option is mandatory since it depends on other network parameters under administrator control.

session-limit-per-prefix—Maximum number of sessions per B4 subnet prefix.

softwire-concentrator —Specify the IP address of the softwire concentrator.

softwire-type—Sets softwire concentrator type to 6rd.

Values: v6rd

v6rd-prefix—IPv6 prefix for the 6rd domain.

For map-e softwires:

name—Name of the MAP-E softwire domain name.

br-address—Specify the Border Relay (BR) device unicast IPv6 address as the softwire concentrator IPv6 address.

version—3(Optional) Configure version number to distinguish between currently supported version of the Internet draft draft-ietf-softwire-map-03 (expires on July 28, 2013), *Mapping of Address and Port with Encapsulation (MAP)* and the latest available version.

rule —Specify the name of Map-E the rule.

Options for MAP-E rules:

v4-reassembly | v6-reassembly—(Optional) Enable IPv4 and IPv6 reassembly for MAP-E.

disable-auto-route—Disable auto-routes and enable static routes to facilitate ECMP load balancing.

NOTE: When you enable the **disable-auto-route** option, you must configure static routes.

ipv4-prefix—Configure rule for IPv4 prefix of the MAP-E domain.

ipv6-prefix—Configure rule for IPv6 prefix of the MAP-E domain.

ea-bits-length—Configure rule for Embedded Address (EA) length for the MAP-E domain.

Range: 0 through 48

psid-length—Configure Port Set ID (PSID) length value for the MAP-E domain.

NOTE:

- If the sum of **v4-prefix-len** and **ea-bits-len** is less than 32, then the **psid-len** must be equal to the difference between 32 and the sum total of **v4-prefix-len** and **ea-bits-len**.

Range: 0 through 16

psid-offset—(Optional) Configure PSID offset value for the MAP-E domain.

Default: 4

Range: 0 through 16

mtu-v6—(Optional) Specify the Maximum transmission unit (MTU) for the MAP-E softwire tunnel.

Default: 9192

Range: 1280 through 9192

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

softwires-rule-set (Service Set Next Gen Services)**Syntax**

```
softwires-rule-set software-rule-set-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the software rule-set that contains the rule to be used with the service set.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[6rd Softwires in Next Gen Services](#) | 165

source-address (Next Gen Services Service-Set Remote System Logging)

Syntax

```
source-address address;
```

Hierarchy Level

```
edit services service-set name syslog
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Specify the IP address of the source for Next Gen Services system log messages.

BEST PRACTICE: The syslog source address can be any arbitrary IP address. It does not have to be an IP address that is assigned to the device. Rather, this IP address is used on the syslog collector to identify the syslog source. The best practice is to configure the source address as the IP address of the interface that the traffic is sent out on.

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

[stream | 711](#)

source-address (NAT Next Gen Services)

Syntax

```
source-address (address | any-unicast);
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],  
[edit services nat source rule-set rule-set rule rule-name match]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the source address that the packet must match for the NAT rule to take effect.

Options

address—A specific address that must be matched.

any-unicast—Any unicast source address results in a match.

Required Privilege Level

interface—To view this statement in the configuration.

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

source-address-name (NAT Next Gen Services)

Syntax

```
source-address-name address-name;
```

Hierarchy Level

```
[edit services nat destination rule-set rule-set rule rule-name match],  
[edit services nat source rule-set rule-set rule rule-name match]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the name of the range of source addresses that the packet must match for the NAT rule to take effect. The range of addresses is configured with the **address** statement at the **[edit services address-book global]** hierarchy level.

Required Privilege Level

interface—To view this statement in the configuration.

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

source-port

Syntax

```
source-port port-number;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Source port identifier.

Options

port-value—Identifier for the port. For a complete list, see *Configuring Source and Destination Ports*.

Required Privilege Level

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

RELATED DOCUMENTATION

ALG Descriptions
Configuring Application Properties
Configuring Source and Destination Ports
Verifying the Output of ALG Sessions 438

source-route-option (IDS Screen Next Gen Services)

Syntax

```
source-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IPv4 packets that have either the IP option of 3 (Loose Source Routing) or the IP option of 9 (Strict Source Routing).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

stateful-firewall-rules (Service Set Next Gen Services)

Syntax

```
stateful-firewall-rules [rule-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the stateful firewall rules to be used with the service set. A stateful firewall rule is configured at the **[edit services policies]** 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 Stateful Firewalls for Next Gen Services](#) | 267

stateful-firewall-rule-set (Next Gen Services)

Syntax

```
stateful-firewall-rule-set {  
    stateful-firewall-rule [rule-name];  
}
```

Hierarchy Level

[edit services policies]

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify a set of stateful firewall rules, which are processed in the order in which they appear in the rule set configuration. Once a stateful firewall rule in the rule set matches a flow, that rule is applied and no other rules in the rule set are processed`.

Options

stateful-firewall-rule [rule-name]—Names of the stateful firewall rules that belong to the rule set. A stateful firewall rule is configured at the **[edit services policies]** 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 Stateful Firewalls for Next Gen Services](#) | 267

stateful-firewall-rule-sets (Service Set Next Gen Services)

Syntax

```
stateful-firewall-rule-sets [rule-set-name];
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the stateful firewall rule sets to be used with the service set. A stateful firewall rule set is configured at the **[edit services policies]** 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 Stateful Firewalls for Next Gen Services](#) | 267

stream (Next Gen Services Service-Set Remote System Logging)

Syntax

```
stream stream-name (severity debug | category screen | format sd-syslog | host);
```

Hierarchy Level

```
edit services service-set name syslog
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the name of the stream to the remote log server.

NOTE: Each remote server requires a unique stream name.

Options

severity debug—

category screen—

format sd-syslog—

host—

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

stream-option (IDS Screen Next Gen Services)

Syntax

```
stream-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IPv4 packets that have the IP option of 8 (Stream ID).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

strict-source-route-option (IDS Screen Next Gen Services)

Syntax

```
strict-source-route-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IPv4 packets that have the IP option of 9 (Strict Source Routing).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

syn-ack-ack-proxy (IDS Screen Next Gen Services)

Syntax

```
syn-ack-ack-proxy {  
    threshold number;  
}
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Configure the maximum number of connections from an IP address that can be opened without being completed. Once this threshold has been reached, further connection requests are rejected. In the SYN-ACK-ACK attack, the session table can fill up, resulting in the device rejecting legitimate connection requests.

Options

threshold *number*—Maximum number of uncompleted connections from any single IP address.

Range: 1 through 250,000

Default: 512

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

syn-fin (IDS Screen Next Gen Services)

Syntax

```
syn-fin;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop packets that have both the SYN and FIN flags set, which can cause unpredictable behavior.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

syn-frag (IDS Screen Next Gen Services)

Syntax

```
syn-frag;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop SYN packet fragments. In TCP SYN fragment attacks, the target caches SYN fragments, waiting for the remaining fragments to arrive so it can reassemble them and complete the connection. A flood of SYN fragments eventually fills the host's memory buffer, preventing valid traffic connections.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

syslog (Services CoS)

Syntax

```
syslog;
```

Hierarchy Level

```
[edit services cos rule rule-name term term-name then],  
[edit services cos rule rule-name term term-name then reverse]
```

Release Information

Statement introduced in Junos OS Release 8.1.

Description

Enable system logging. The system log information from the Multiservices and Services PICs is passed to the kernel for logging in the **/var/log** directory. This setting overrides any **syslog** statement setting included in the service set or interface default configuration.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring CoS Rules on Services PICs](#)

[Configuring Actions in CoS Rules](#)

syslog (Next Gen Services Service-Set System Logging)

Syntax

```
syslog ;
```

Hierarchy Level

```
[edit services service-set name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Configure the filename Next Gen Services system logs.

Options

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

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

tcp-no-flag (IDS Screen Next Gen Services)

Syntax

```
tcp-no-flag;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop TCP packets that have no flag fields set. A TCP no flag attack can cause unpredictable behavior on the target.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

tcp-session (Service Set Next Gen Services)

Syntax

```
tcp-session {
  inactivity-asymm-tcp-timeout;
  inactivity-tcp-timeout ;
  open-timeout ;
  tcp-fast-open ;
  tcp-mss ;
  tcp-non-syn ;
  tcp-tickles ;
}
```

Hierarchy Level

```
[edit services service-set service-set-name service-set-options]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series.

Description

Configure the TCP options for the service set.

Options

close-timeout—Timeout period for TCP session tear-down (2. through 300 seconds)

ignore-errors—Ignore anomalies or errors for TCP

inactivity-asymm-tcp-timeout—

tcp-tickles—Number of TCP keep-alive packets to be sent for bidirectional TCP flows

inactivity-tcp-timeout—Inactivity timeout period for TCP established sessions

open-timeout—Timeout period for TCP session establishment (seconds)

tcp-fast-open—Tcp-fast-Open enabled packets will be handled accordingly

tcp-mss—Enable the limit on TCP Max. Seg. Size in SYN packets

tcp-non-syn—Deny session creation on receiving first non SYN packet

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

.

tcp-tickles (Service Set Next Gen Services)

Syntax

```
tcp-tickles tcp-tickles;
```

Hierarchy Level

```
[edit service-set service-set-name service-set-optiontcp-session]
```

Release Information

Statement introduced in Junos OS Release 19.3R1 on MX Series.

Description

Define the maximum number of keep-alive messages sent before a TCP session is allowed to timeout.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

.

tear-drop (IDS Screen Next Gen Services)

Syntax

```
tear-drop;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop fragmented IP packets that overlap, which protects against teardrop attacks. In teardrop attacks, the target machine uses up its resources as it attempts to reassemble the packets, and then it can no longer process valid traffic.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

then (Services CoS Next Gen Services)

Syntax

```

then {
  application-profile profile-name;
  dscp (alias | bits);
  forwarding-class class-name;
  reflexive; | revert; | reverse {
    application-profile profile-name;
    dscp (alias | bits);
    forwarding-class class-name;
  }
}

```

Hierarchy Level

```
[edit services cos rule rule-name policy policy-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the Differentiated Services (DiffServ) code point (DSCP) marking and forwarding-class assignments for packets that are processed by a service set and that match the conditions of the policy in a services CoS rule.

The service set that the CoS rule is assigned to must include at least one stateful firewall rule or NAT rule, or CoS does not work. Only stateful firewall and NAT rules can be used with CoS rules in a service set.

Options

application-profile *profile-name*—The application profile that sets the CoS actions for FTP and SIP traffic.

dscp (*alias* | *bits*)—Either a code point alias or a DSCP bit value to apply to the packet.

forwarding-class *class-name*—Forwarding class name to apply to the packet. The choices are:

- assured-forwarding
- best-effort
- expedited-forwarding
- network-control

reflexive—Applies the CoS rule policy actions to flows in the reverse direction as well as to flows in the matching direction.

revert—Stores the DSCP and forwarding class of a packet that is received in the match direction of the rule and then applies that DSCP and forwarding class to packets that are received in the reverse direction of the same session.

reverse—Specifies actions to apply to flows in the reverse direction of the matching direction.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [Class of Service for Services PICs \(Next Gen Services\) | 255](#)

then (Stateful Firewall Rule Next Gen Services)

Syntax

```
then {  
    count;  
    deny;  
    permit;  
    reject;  
}
```

Hierarchy Level

```
[edit services policies stateful-firewall-rule rule-name policy policy-name]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Specify the actions for a stateful firewall rule policy. The policy actions are applied to flows that meet the policy's matching properties.

Options

count—Enables a count, in bytes or kilobytes, of all network traffic the policy allows to pass.

deny—Drop the packets.

permit—Accept the packets and send them to their destination.

reject—Drop the packets. For TCP traffic, send a TCP reset (RST) segment to the source host. For UDP traffic, send an ICMP **destination unreachable**, **port unreachable** message (type 3, code 3) to the source host.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Stateful Firewalls for Next Gen Services](#) | 267

timestamp-option (IDS Screen Next Gen Services)

Syntax

```
timestamp-option;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IPv4 packets that have the IP option of 4 (Internet timestamp).

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

traceoptions (Next Gen Services Service-Set Flow)

Syntax

```

traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    size maximum-file-size;
    (world-readable | no-world-readable);
  }
  flag flag;
  no-remote-trace;
  packet-filter filter-name {
    conn-tag session-conn
    destination-port port-identifier;
    destination-prefix address;
    interface interface-name;
    protocol protocol-identifier;
    source-port port-identifier;
    source-prefix address;
  }
  rate-limit messages-per-second;
  trace-level (brief | detail | error);
}

```

Hierarchy Level

[edit services service-set name flow]

Release Information

Statement introduced in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480 and MX960.

Description

Configure flow tracing options for a service-set.

Options

file—Configure the trace file options.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. By default, the name of the file is the name of the process being traced.

files *number*—Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed to ***trace-file.0***, then ***trace-file.1***, and so on, until the maximum number of trace files is reached. The oldest archived file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000 files

Default: 10 files

match *regular-expression*—Refine the output to include lines that contain the regular expression.

size *maximum-file-size*—Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: **x K** to specify KB, **x m** to specify MB, or **x g** to specify GB

Range: 0 KB through 1 GB

Default: 128 KB

world-readable | no-world-readable—By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

flag—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

all—Trace with all flags enabled

basic-datapath—Trace basic packet flow activity

fragmentation—Trace IP fragmentation and reassembly events

high-availability—Trace flow high-availability information

host-traffic—Trace flow host traffic information

multicast—Trace multicast flow information

route—Trace route lookup information

session—Trace session creation and deletion events

session-scan—Trace session scan information

tcp-basic—Trace TCP packet flow information

tunnel—Trace tunnel information

no-remote-trace—Set remote tracing as disabled.

packet-filter *filter-name*—Packet filter to enable during the tracing operation. Configure the filtering options.

destination-port *port-identifier*—Match TCP/UDP destination port

destination-prefix *address*—Destination IP address prefix

interface *interface-name*—Logical interface

protocol *protocol-identifier*—Match IP protocol type

source-port *port-identifier*—Match TCP/UDP source port

source-prefix *address*—Source IP address prefix

rate-limit *messages-per-second*—Limit the incoming rate of trace messages.

trace-level—Set the level for trace logging. This option is available only when the flag is set.

brief—Trace key flow information, such as message types sent between SPU and central point, policy match, and packet drop reasons.

detail—Trace extensive flow information, such as detailed information about sessions and fragments. Detail is the default level.

error—Trace error information, such as system failure, unknown message type, and packet drop.

Required Privilege Level

trace—To view this in the configuration.

trace-control—To add this to the configuration.

RELATED DOCUMENTATION

traceoptions (Traffic Load Balancer)

Syntax

```
traceoptions {
  file file-name <files number> <no-word-readable | world-readable> <size size>;
  flag flag;
  level (all | critical | error | info | notice | verbose | warning);
  monitor monitor-object-name {
    instance-name instance-name;
    virtual-svc-name virtual-service-name;
  }
  no-remote-trace;
}
```

Hierarchy Level

```
[edit services traffic-load-balance]
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

instance-name and **virtual-service-name** options added in Junos OS Release 16.1R6 and 18.2R1 on MX Series.

Support for Next Gen Services MX-SPC3 services card add in Junos OS Release 19.3R2.

Description

Configure tracing options for the traffic load balancer.

Options

file *file-name*—Name of the file to receive the output of the tracing operation.

files *number*—(Optional) Maximum number of trace files. When a trace file named trace-file reaches its maximum size, it is renamed trace-file.0, then trace-file.1, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Range: 2 through 1000 files

Default: 3 files

flag *flag*— Specify which operations you want to trace from [Table 56 on page 732](#). To specify more than one operation, include multiple flag statements.

Table 56: Trace Flags

Flag	Support on MS-MPC and MX-SPC3 Cards	Description
all	MS-MPC and MX-SPC3	Trace all operations.
all-real-services	MS-MPC and MX-SPC3	Trace all real services.
database	MS-MPC and MX-SPC3	Trace database events.
file-descriptor-queue	MS-MPC and MX-SPC3	Trace file descriptor queue events.
inter-thread	MS-MPC and MX-SPC3	Trace inter-thread communication events.
messages	MS-MPC and MX-SPC3	Trace normal events.
probe	MS-MPC and MX-SPC3	Trace probe events.
probe-infra	MS-MPC and MX-SPC3	Trace probe infra events.

instance-name *instance-name*—(Optional) Name of the TLB instance to monitor.

level—Use the specified level of tracing. You can specify any of the following levels:

- **all**—Match all levels.
- **error**—Match error conditions.
- **info**—Match informational messages.
- **notice**—Match conditions that must be handled specially.
- **verbose**—Match verbose messages.
- **warning**—Match warning messages.

These trace levels are available for both the MS-MPC and MX-SPC3 services cards unless otherwise specified.

monitor *monitor-object-name*—Name of a monitoring object that contains an instance name or virtual service name.

no-remote-trace—(Optional) Disable remote tracing.

no-world-readable—(Optional) Disable unrestricted file access.

For Next Gen Services on the MX-SPC3 services card, set the *monitor-object-name* to either:

group-name—Name of the group.

real-services-name—Name of the real service

size *size*—(Optional) Use the maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the *size* option.

Syntax: *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB.

Range: 10,240 through 1,073,741,824 bytes.

Default: 128 KB

virtual-svc-name *virtual-service-name*—(Optional) Name of the virtual service to monitor.

word-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

trace and interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Traffic Load Balancer Overview | 287](#)

[Configuring TLB | 296](#)

traceoptions (Next Gen Services Global System Logging)

Syntax

```
traceoptions {  
  apply-groups group-names;  
  apply-groups-except group-names;  
  flag name;  
  file filename,  
  no-remote-trace;  
}
```

Hierarchy Level

```
[edit services rtlog]
```

Release Information

Support introduced in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the trace information you want to include in the system log messages.

Options

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

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

traceoptions (Next Gen Services Softwires)

Syntax

```
traceoptions {
  file {
    filename;
    files number;
    match regular-expression;
    (no-world-readable | world-readable);
    size maximum-file-size;
  }
  flag (all | configuration | flow);
  no-remote-trace;
}
```

Hierarchy Level

[edit security softwires]

Release Information

Statement introduced before Release 12.1 of Junos OS.

Description

Configure softwire tracing options.

Options

- **file**—Configure trace file information.
 - **filename**—Name of the file to which to write the trace information.
 - **files number**—Maximum number of trace files.
Range: 2 through 1000 files
 - **match regular-expression**—Regular expression for lines to be logged.
 - **no-world-readable | world-readable**—Allow or deny any user to read the log file.
 - **size maximum-file-size**—Maximum trace file size.
Range: 10,240 to 1,073,741,824 bytes
- **flag**—Specify events to trace.
 - **all**—Trace all events
 - **configuration**—Trace configuration events

- **flow**—Trace flow events
- **no-remote-trace**—Disable remote tracing.

Required Privilege Level

trace—To view this statement in the configuration.

trace-control—To add this statement to the configuration.

RELATED DOCUMENTATION

traffic-load-balance (Traffic Load Balancer)

Syntax

```

traffic-load-balance {
  instance instance-name {
    client-interface client-interface;
    client-vrf client-vrf;
    group group-name {
      health-check-interface-subunit health-check-interface-subunit;
      network-monitoring-profile [profile-name1, <profile-name2>];
      real-service-rejoin-options no-auto-rejoin;
      real-services [server-list];
      <routing-instance routing-instance>;
    }
    interface interface-name;
    real-service real-service {
      address server-ip-address;
      admin-down;
    }
    server-inet-bypass-filter server-inet-bypass-filter ;
    server-inet6-bypass-filter server-inet6-bypass-filter ;
    server-interface server-interface;
    server-vrf server-vrf;
    traceoptions {
      file file-name <files number> <no-word-readable | world-readable> <size size>;
      flag flag;
      level (all | critical | error | info | notice | verbose | warning);
      monitor {
        instance-name instance-name;
        virtual-svc-name virtual-service-name;
      }
      no-remote-trace;
    }
    virtual-service virtual-service-name {
      address virtual-ip-address;
      group group-name;
      load-balance-method {
        hash {
          hash-key method;
        }
        random;
      }
      mode ( layer2-direct-server-return | direct-server-return | translated );
      <routing-instance routing-instance-name>;
    }
  }
}

```

```

    <routing-metric route-metric>;
    server-interface server-interface;
    service service-name {
        protocol (udp | tcp);
        server-listening-port port;
        virtual-port virtual-port;
    }
}
}
}

```

Hierarchy Level

[edit services]

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure traffic load balancer options.

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

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Traffic Load Balancer Overview | 287](#)

[Configuring TLB | 296](#)

transport (Next Gen Services Syslog Message Security)

Syntax

```
transport;
```

Hierarchy Level

```
[edit services service-set name syslog
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Specify the category for which you want to collect local logs.

Options

apply-groups—Groups from which to inherit configuration data

apply-groups-except—Don't inherit configuration data from these groups

protocol—Set security log transport protocol for the device. You can set the protocol to TCP, TLS or UDP

tcp-connections—Set tcp connection number per-stream (1..5)

tls-profile—If you are using the TLS protocol as the security log transport, specify the TLS profile name to use.

Required Privilege Level

system

RELATED DOCUMENTATION

[Understanding Next Gen Services CGNAT Global System Logging | 71](#)

[Enabling Global System Logging for Next Gen Services | 73](#)

[Configuring System Logging to One or More Remote Servers for Next Gen Services | 76](#)

[Configuring Local System Logging for Next Gen Services | 74](#)

ttl-threshold

Syntax

```
ttl-threshold number;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the traceroute time-to-live (TTL) threshold value. This value sets the acceptable level of network penetration for trace routing.

Options

number—TTL threshold value.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

ALG Descriptions

Configuring the TTL Threshold.

[Examples: Configuring Application Protocols](#) | 437

[Verifying the Output of ALG Sessions](#) | 438

unknown-protocol (IDS Screen Next Gen Services)

Syntax

```
unknown-protocol;
```

Hierarchy Level

```
[edit services screen ids-option screen-name ip]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop IP frames with protocol numbers greater than 137 for IPv4 and 139 for IPv6, which protects against IP unknown protocol attacks.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

url-filter

Syntax

```
url-filter {
  profile profile-name {
    template template-name {
      client-interfaces [ client-interface-name1 client-interface-name2 ];
      disable-url-filtering;
      dns-resolution-interval minutes;
      dns-resolution-rate seconds;
      dns-retries number;
      dns-routing-instance dns-routing-instance-name;
      dns-server [ ip-address1 ip-address2 ip-address3 ];
      dns-source-interface loopback-interface-name;
      routing-instance routing-instance-name;
      server-interfaces [ server-interface-name1 server-interface-name2 ];
      term term-name {
        from {
          src-ip-prefix [prefix1 prefix2];
          dest-port [port1 port2];
        }
        then {
          accept;
          custom-page custom-page;
          http-status-code http-status-code;
          redirect-url redirect-url;
          tcp-reset;
        }
      }
      url-filter-database filename
    }
    url-filter-database filename;
  }
}
```

Hierarchy Level

[edit services]

Release Information

Statement introduced in Junos OS Release 17.2.

Description

Configure URL filtering service.

NOTE: Starting in Junos OS Release 18.3R1, the **url-filter** statement is deprecated and has been replaced by the **web-filter** statement. The **url-filter** statement is supported for backward compatibility.

Options

url-filter-database *filename*—Specify the filename of the URL filter database. This option is mandatory.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring URL Filtering | 335](#)

[URL Filtering Overview | 332](#)

url-filter-profile

Syntax

```
url-filter-profile profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 17.2.

Description

Specify the URL filter profile that the service set uses. The URL filter profile specifies how to filter access to disallowed URLs, and is configured at the **[edit services url-filter]** hierarchy level.

NOTE: You must also configure the **next-hop-service** statement with this statement.

NOTE: Starting in Junos OS Release 18.3R1, the **url-filter-profile** statement is deprecated and has been replaced by the **web-filter-profile** statement. The **url-filter-profile** statement is supported for backward compatibility.

Options

profile-name—Name of the URL filter profile.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

[Configuring URL Filtering | 335](#)

[URL Filtering Overview | 332](#)

[url-filter | 742](#)

url-filter-template

Syntax

```
url-filter-template template-name {
  client-interfaces [ client-interface-name1 client-interface-name2 ];
  disable-url-filtering;
  dns-resolution-interval minutes;
  dns-resolution-rate seconds;
  dns-retries number;
  dns-routing-instance dns-routing-instance-name;
  dns-server [ ip-address1 ip-address2 ip-address3 ];
  dns-source-interface loopback-interface-name;
  routing-instance routing-instance-name;
  security-intelligence-policy
  server-interfaces [ server-interface-name1 server-interface-name2 ];
  term term-name {
    from {
      src-ip-prefix [ prefix1 prefix2 ];
      dest-port [ port1 port2 ];
    }
    then {
      accept;
      custom-page custom-page;
      http-status-code http-status-code;
      redirect-url redirect-url;
      tcp-reset;
    }
  }
  url-filter-database filename
}
```

Hierarchy Level

```
[edit services web-filter profile profile-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Statement introduced in Junos OS Release 20.1R1 for Next Gen Services on MX240, MX480, and MX960 routers.

Description

Configure a URL filter template.

Options

template-name—Name of the URL filter template.

client-interfaces [*client-interface-name1 client-interface-name2*]—The list of client-facing logical interfaces (uplink) on which the URL filtering is configured. This option is mandatory.

disable-url-filtering—Disables the filtering of HTTP traffic that contains an embedded IP address (for example, http://10.1.1.1) belonging to a disallowed domain name in the URL filter database.

dns-resolution-interval *minutes*—DNS resolution time interval in minutes.

Default: 1440

Range: 60 through 1440 minutes.

dns-resolution-rate *seconds*—Number of DNS queries per second sent out from the system before initiating further DNS queries.

Default: 50

Range: 50 through 100.

dns-retries *number*—Number of retries for a DNS query in case query fails or times out.

Default: 3

Range: 1 through 5.

dns-routing-instance *dns-routing-instance-name*—The VRF on which the DNS server is reachable. This option is mandatory. You can use the default routing instance inet.0 or a defined routing instance.

dns-server [*ip-address1 ip-address2 ip-address3*]—One or more IP (IPv4 or IPv6) addresses of DNS servers to which the DNS queries are sent out. This option is mandatory.

dns-source-interface *loopback-interface-name*—The loopback interface for which source IP address is picked for sending DNS queries. This option is mandatory.

routing-instance *routing-instance-name*—The VRF on which URL filtering feature is configured. This option is mandatory. You can use the default routing instance inet.0 or a defined routing instance.

server-interfaces [*server-interface-name1 server-interface-name2*]—Server-facing interfaces to which traffic is destined. This option is mandatory.

The list of server-facing logical interfaces (downlink) on which the URL filtering is configured. This option is mandatory.

url-filter-database *filename*—The filename of the URL filter database. The file should be placed in the `/var/db/url-filterd` directory, but indicate just the filename here and not the full path.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [Configuring URL Filtering | 335](#)

uuid**Syntax**

```
uuid hex-value;
```

Hierarchy Level

```
[edit applications application application-name]
```

Release Information

Statement introduced before Junos OS Release 7.4.

Description

Specify the Universal Unique Identifier (UUID) for DCE RPC objects.

Options

hex-value—Hexadecimal value.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

| [ALG Descriptions](#)

| [Configuring a Universal Unique Identifier](#)

| [Examples: Configuring Application Protocols | 437](#)

| [Verifying the Output of ALG Sessions | 438](#)

v6rd

Syntax

```
v6rd v6rd-software-concentrator {
    ipv4-prefix ipv4-prefix;
    v6rd-prefix ipv6-prefix;
    mtu-v4 mtu-v4;
    software-address ipv4-address;
}
```

Hierarchy Level

```
[edit services software software-concentrator]
[edit services softwares software-types]
```

Release Information

Statement introduced in Junos OS Release 10.4.

Support added in Junos OS release 20.2R1 for the v6rd concentrator at the **[edit services softwares software-types]** edit hierarchy for Next Gen Services on MX240, MX480, and MX860 routers.

Description

Configure settings for a 6rd concentrator used to process IPv6 packets encapsulated in IPv4 packets.

The **v6rd** statement is supported only on the MS-DPC, MS-100, MS-400, and MS-500 line cards. The **v6rd** statement is *not* supported on MS-MPCs and MS-MICs.

Options

ipv4-prefix—IPv4 prefix of the customer edge (CE) network

ipv6-prefix—IPv6 prefix of the 6rd domain.

mtu-v4— Maximum transmission unit (MTU), in bytes (576 through 9192), for IPv6 packets encapsulated into IPv4. If the final length is greater than the configured value, the IPv4 packet will be dropped.

address—IPv4 address of a software concentrator. This is an IPv4 address independent of any interface and on a different prefix.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

video (Application Profile)

Syntax

```
video {  
  dscp (alias | bits);  
  forwarding-class class-name;  
}
```

Hierarchy Level

```
[edit services cos application-profileprofile-name sip]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Set the appropriate **dscp** and **forwarding-class** values for SIP video traffic.

Default

By default, the system will not alter the DSCP or forwarding class for SIP video traffic.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[voice \(Application Profile\)](#) | 754

video (Application Profile)

Syntax

```
video {  
  dscp (alias | bits);  
  forwarding-class class-name;  
}
```

Hierarchy Level

```
[edit services cos application-profile profile-name sip]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Set the appropriate **dscp** and **forwarding-class** values for SIP video traffic.

Default

By default, the system will not alter the DSCP or forwarding class for SIP video traffic.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[voice \(Application Profile\)](#) | 754

virtual-service (Traffic Load Balancer)

Syntax

```
virtual-service virtual-service-name {
  address virtual-ip-address;
  group group-name;
  load-balance-method {
    hash {
      hash-key method;
    }
    random;
  }
  mode ( layer2-direct-server-return | direct-server-return | translated );
  <routing-instance routing-instance-name>;
  <routing-metric route-metric>;
  server-interface server-interface;
  service service-name {
    protocol (udp | tcp);
    server-listening-port port;
    virtual-port virtual-port;
  }
}
```

Hierarchy Level

```
[edit services traffic-load-balance instance instance-name]
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure a TLB virtual service.

Options

address *virtual-ip-address*—Address of the virtual service.

group *group-name*—Server group for the virtual service.

load-balance method hash hash-key *method*—Use a combination of these hash-key methods for the session distribution API:

dest-ip—Hash on destination IP address.

proto—Hash on protocol.

source-ip—Hash on source IP address.

load-balance-method random—Use randomizing algorithm for session distribution.

mode (layer2-direct-server-return | direct-server-return | translated)—Traffic load balancer mode of operation:

direct-server-return—Transparent mode Layer 3 direct server return.

layer2-direct-server-return—Transparent mode Layer 2 direct server return. Load balancing works by changing the Layer 2 MAC of packets; Layer 3 and higher level headers are not modified.

translated—The Packet Forwarding Engine performs stateless load balancing.

route-metric—(Optional) Route metric

Range: 1 through 255

routing-instance-name—(Optional) Routing instance for the virtual service. Default is **inet.0**.

server-interface server-interface—(Optional) The server-interface specified under the virtual-service, will be used instead of the values provided under the instance level.

service service-name—Translated mode details. Packets destined to this virtual ip-address + virtual-port + protocol will be load balanced to the appropriate server. The destination IP address and port are replaced by the real services IP address and the server-listening-port (configured here).

protocol (udp | tcp)—Protocol.

server-listening-port port—Port number.

virtual-port virtual-port—Virtual port number.

virtual-ip-address—Local address for the virtual service.

virtual-service-name—Identifier for the virtual service.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Traffic Load Balancer Overview | 287](#)[Configuring TLB | 296](#)

voice

Syntax

```
voice {  
    dscp (Services CoS) (alias | bits);  
    forwarding-class (Services PIC Classifiers) class-name;  
}
```

Hierarchy Level

```
[edit services (CoS) cos application-profile profile-name sip]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Set the appropriate **dscp** and **forwarding-class** values for SIP voice 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

[Configuring Application Profiles for Use as CoS Rule Actions](#)

voice (Application Profile)

Syntax

```
voice {  
    dscp (alias | bits);  
    forwarding-class class-name;  
}
```

Hierarchy Level

```
[edit services cos application-profileprofile-name sip]
```

Release Information

Statement introduced in Junos OS Release 9.3.

Description

Set the appropriate **dscp** and **forwarding-class** values for SIP voice traffic.

Default

By default, the system will not alter the DSCP or forwarding class for SIP voice traffic.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

Configuring CoS Rules on Services PICs

[video \(Application Profile\)](#) | **749**

web-filter

Syntax

```

web-filter {
  profile (Web Filter) profile-name {
    dns-filter {
      database-file filename;
      dns-resp-ttl seconds;
      dns-server [ ip-address ];
      hash-key key-string;
      hash-method hash-method-name;
      statistics-log-timer minutes;
      wildcarding-level level;
    }
    dns-filter-template template-name {
      client-interfaces [ client-interface-name ];
      client-routing-instance client-routing-instance-name;
      dns-filter {
        database-file filename;
        dns-resp-ttl seconds;
        dns-server [ ip-address ];
        hash-key key-string;
        hash-method hash-method-name;
        statistics-log-timer minutes;
        wildcarding-level level;
      }
      server-interfaces [ server-interface-name ];
      server-routing-instance server-routing-instance-name;
      term term-name {
        from {
          src-ip-prefix [ source-prefix ];
        }
        then {
          accept;
          dns-sinkhole;
        }
      }
    }
  }
  global-dns-stats-log-timer minutes;
  url-filter-database filename;
  url-filter-template template-name {
    client-interfaces [ client-interface-name1 client-interface-name2 ];
    disable-url-filtering;
    dns-resolution-interval minutes;
  }
}

```

```

    dns-resolution-rate seconds;
    dns-retries number;
    dns-routing-instance dns-routing-instance-name;
    dns-server [ ip-address1 ip-address2 ip-address3 ];
    dns-source-interface loopback-interface-name;
    dns-routing-instance dns-routing-instance-name;
    routing-instance routing-instance-name;
    server-interfaces [ server-interface-name1 server-interface-name2 ];
    term term-name {
        from {
            src-ip-prefix [prefix1 prefix2];
            dest-port [port1 port2];
        }
        then {
            accept;
            custom-page custom-page;
            http-status-code http-status-code;
            redirect-url redirect-url;
            tcp-reset;
        }
    }
    url-filter-database filename
}
}
}

```

Hierarchy Level

[edit services]

Release Information

Statement introduced in Junos OS Release 18.3R1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Configure filtering of DNS requests for disallowed website domains. Filtering can result in either:

- Blocking access to the site by sending the client a DNS response that includes an IP address or domain name of a sinkhole server instead of the disallowed domain.
- Logging the DNS request and allowing access.

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

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [DNS Request Filtering for Disallowed Website Domains](#) | 311

web-filter-profile

Syntax

```
web-filter-profile profile-name;
```

Hierarchy Level

```
[edit services service-set service-set-name]
```

Release Information

Statement introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Specify the DNS filter profile or the URL filter profile that the service set uses. The filter profile is configured at the **[edit services web-filter]** hierarchy level, and specifies how to filter DNS requests for disallowed website domains or how to filter access to disallowed URLs.

Options

profile-name—Name of the DNS filter profile.

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

RELATED DOCUMENTATION

| [DNS Request Filtering for Disallowed Website Domains](#) | 311

winnuke (IDS Screen Next Gen Services)

Syntax

```
winnuke;
```

Hierarchy Level

```
[edit services screen ids-option screen-name tcp]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 on MX Series routers (MX240, MX480 and MX960) running Next Gen Services with the MX-SPC3 services card.

Description

Identify and drop TCP segments that are destined for port 139 and have the urgent (URG) flag set, which provides protection against WinNuke attacks.

Required Privilege Level

interface—To view this statement in the configuration.

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

RELATED DOCUMENTATION

[Configuring Network Attack Protection With IDS Screens for Next Gen Services](#) | 276

world-readable (Next Gen Services Global System Logging)

Syntax

```
world-readable;
```

Hierarchy Level

```
[edit services rtlog traceoptions file]
```

Release Information

Statement introduced in Junos OS Release 19.3R2 for Next Gen Services.

Description

Allow any user to read the log file

Options

world-readable—Allow any user to read the log file

Required Privilege Level

system

RELATED DOCUMENTATION

Understanding Next Gen Services CGNAT Global System Logging	71
Enabling Global System Logging for Next Gen Services	73
Configuring System Logging to One or More Remote Servers for Next Gen Services	76
Configuring Local System Logging for Next Gen Services	74

14

PART

Operational Commands

Operational Commands | **761**

Operational Commands

IN THIS CHAPTER

- `clear log (Next Gen Services)` | 764
- `clear services alg statistics` | 765
- `clear services nat source mappings` | 766
- `clear services sessions` | 769
- `clear services sessions analysis` | 773
- `clear services stateful-firewall flows` | 774
- `clear services stateful-firewall sip-call` | 777
- `clear services stateful-firewall sip-register` | 780
- `clear services stateful-firewall statistics` | 783
- `clear services subscriber analysis` | 784
- `clear services web-filter statistics profile` | 785
- `request services web-filter update dns-filter-database` | 787
- `request services web-filter validate dns-filter-file-name` | 788
- `request system disable unified-services` | 789
- `request system enable unified-services` | 790
- `show interfaces load-balancing (Aggregated Multiservices)` | 791
- `show log` | 796
- `show services alg conversations` | 804
- `show services alg statistics` | 812
- `show services cos statistics (Next Gen Services)` | 829
- `show services inline software statistics` | 833
- `show services inline ip-reassembly statistics` | 838
- `show services nat destination pool` | 844
- `show services nat destination rule` | 846
- `show services nat destination summary` | 849
- `show services nat ipv6-multicast-interfaces` | 851
- `show services nat resource-usage source-pool` | 854
- `show services nat source deterministic` | 856

- [show services nat source mappings address-pooling-paired | 859](#)
- [show services nat source mappings endpoint-independent | 863](#)
- [show services nat source mappings pcg | 866](#)
- [show services nat source mappings summary | 868](#)
- [show services nat source pool | 870](#)
- [show services nat source port-block | 876](#)
- [show services nat source rule | 879](#)
- [show services nat source rule-application | 882](#)
- [show services nat source summary | 884](#)
- [show services pcg statistics | 886](#)
- [show services policies | 890](#)
- [show services policies detail | 893](#)
- [show services policies hit-count | 896](#)
- [show services policies interface | 897](#)
- [show services policies service-set | 898](#)
- [show services redundancy-group | 899](#)
- [show services screen ids-option \(Next Gen Services\) | 909](#)
- [show services screen-statistics service-set \(Next Gen Services\) | 911](#)
- [show services security-intelligence category summary | 916](#)
- [show services security-intelligence update status | 919](#)
- [show services service-sets cpu-usage | 920](#)
- [show services service-sets memory-usage | 922](#)
- [show services service-sets plug-ins | 925](#)
- [show services service-sets statistic screen-drops \(Next Gen Services\) | 926](#)
- [show services service-sets statistic screen-session-limit-counters \(Next Gen Services\) | 936](#)
- [show services service-sets statistics integrity-drops | 943](#)
- [show services service-sets statistics packet-drops | 949](#)
- [show services service-sets statistics syslog | 951](#)
- [show services service-sets statistics tcp | 959](#)
- [show services service-sets summary | 961](#)
- [show services sessions \(Next Gen Services\) | 963](#)
- [show services sessions \(Aggregated Multiservices\) | 977](#)
- [show services sessions analysis | 986](#)
- [show services sessions analysis \(USF\) | 991](#)

- [show services sessions count | 996](#)
- [show services sessions service-set | 997](#)
- [show services sessions service-set | 998](#)
- [show services sessions softwire | 1000](#)
- [show services sessions utilization | 1005](#)
- [show services softwire | 1006](#)
- [show services softwire flows | 1008](#)
- [show services softwire statistics | 1013](#)
- [show services stateful-firewall conversations | 1024](#)
- [show services stateful-firewall flow-analysis | 1029](#)
- [show services stateful-firewall flows | 1035](#)
- [show services stateful-firewall sip-call | 1042](#)
- [show services stateful-firewall sip-register | 1048](#)
- [show services stateful-firewall statistics | 1052](#)
- [show services stateful-firewall statistics application-protocol sip | 1063](#)
- [show services subscriber analysis | 1067](#)
- [show services tcp-log | 1070](#)
- [show services traffic-load-balance statistics | 1071](#)
- [show services web-filter dns-resolution profile | 1086](#)
- [show services web-filter dns-resolution-statistics profile template | 1090](#)
- [show services web-filter secintel-policy status | 1096](#)
- [show services web-filter statistics dns-filter-template | 1099](#)
- [show services web-filter statistics profile | 1102](#)
- [show system unified-services status | 1108](#)

clear log (Next Gen Services)

Syntax

clear log service-set | interface | file-name

Release Information

Command introduced in Junos OS Release 20.3R1.

Description

Clear log for service-set, interface, or file.

Options

service-set—Specify the name of the service-set for which you want to clear the log.

interface-name—Specify the name of the interface for which you want to clear the log.

file-name—Specify the file-name for which you want to clear the log.

Required Privilege Level

View

RELATED DOCUMENTATION

| *monitor start (JDM)*

List of Sample Output

[clear log on page 764](#)

Output Fields

This command produces no output.

Sample Output

clear log

```
user@host> clear log vms 1/0/0
```

clear services alg statistics

Syntax

```
clear services alg statistics
```

Release Information

Command introduced in Junos OS Release 10.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Clear ALG statistics for Junos OS extension-provider packages.

Options

application-profile—Clear all sessions for the application profile.

interface—Clear all sessions for the interface.

Required Privilege Level

view

clear services nat source mappings

Syntax

```
clear services nat source mappings
<app | eim | pcp>
subscriber private-ip [port port-num] [service-set service-set]
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Clear services NAT source mappings. After one mapping is cleared, all the port block alloation blocks referring to that mapping are released.

Options

app—Clear all APP mappings.

app subscriber private-ip [port port-num] [service-set service-set]—Clear one APP mapping by matching conditions

eim—Clear all EIM mappings.

eim subscriber private-ip [port port-num] [service-set service-set]—Clear one EIM mapping by matching conditions

pcp—Clear all PCP mappings.

Required Privilege Level

view

List of Sample Output

- [clear services nat source mappings eim on page 767](#)
- [clear service nat source mappings eim subscriber 2.1.1.1 on page 767](#)
- [clear services nat source mappings subscriber 2.1.1.1 port 1026 service-set ss1 on page 767](#)
- [clear services nat source mappings app on page 767](#)
- [clear services nat source mappings app subscriber 2.1.1.1 on page 768](#)
- [clear services nat source mappings app subscriber 2.1.1.1 port 1026 service-set ss1 on page 768](#)

Output Fields

[Table 57 on page 767](#) lists the output fields for the **clear services nat source mappings** command. Output fields are listed in the approximate order in which they appear.

Table 57: clear services nat source mappings Output Fields

Field Name	Field Description
NAT pool	Name of the NAT pool.
Mappings removed	Number of mappings removed.
Sessions removed	Number of sessions removed.

Sample Output

clear services nat source mappings eim

```
user@host> clear services nat source mappings eim
```

NAT pool	Mappings removed	Sessions removed	
Test-pool		1	0

clear service nat source mappings eim subscriber 2.1.1.1

```
user@host> clear service nat source mappings eim subscriber 2.1.1.1
```

NAT pool	Mappings removed	Sessions removed	
Test-pool		1	0

clear services nat source mappings subscriber 2.1.1.1 port 1026 service-set ss1

```
user@host> clear services nat source mappings subscriber 2.1.1.1 port 1026 service-set ss1
```

NAT pool	Mappings removed	Sessions removed	
Test-pool		1	0

clear services nat source mappings app

```
user@host> clear services nat source mappings app
```

NAT pool	Mappings removed	Sessions removed	
Test-pool		1	0

clear services nat source mappings app subscriber 2.1.1.1

user@host> **clear services nat source mappings app subscriber 2.1.1.1**

NAT pool	Mappings removed	Sessions removed
Test-pool		1
		0

clear services nat source mappings app subscriber 2.1.1.1 port 1026 service-set ss1

user@host> **clear services nat source mappings app subscriber 2.1.1.1 port 1026 service-set ss1**

NAT pool	Mappings removed	Sessions removed
Test-pool		1
		0

clear services sessions

Syntax

```
clear services sessions
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<ip-action>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced in Junos OS Release 13.1.

Description

Clear services sessions currently active on the embedded PIC or MIC. When you enter this command, the sessions are marked for deletion and are cleared thereafter. The time that is taken to clear the currently active sessions varies, depending on the scaled nature of the environment.

Options

none—Clear all sessions.

application-protocol *protocol*—(Optional) Clear sessions for one of the following application protocols:

- **bootp**—Bootstrap protocol
- **dce-rpc**—Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—Domain Name System protocol
- **exec**—Exec
- **ftp**—File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **ip**—IP
- **login**—Login

- **netbios**—NetBIOS
- **netshow**—NetShow
- **pptp**—Point-to-Point Tunneling Protocol
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **talk**—Talk Program
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Clear sessions for the specified destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Clear sessions for the specified destination prefix.

interface *interface-name*—(Optional) Clear sessions for the specified interface. On M Series and T Series routers, the *interface-name* can be **ms-fpc/ pic/ port** or **rspnumber**.

ip-action—(Optional) Clear **ip-action** entries generated by the router to log, drop, or block traffic based on previous matches. The IP action options and targets are configured at the **{edit security idp idp-policy policy-name rulebase-ips rule rule-name then}** hierarchy level.

protocol *protocol*—(Optional) Clear sessions for one of the following IP types:

- **number**—Numeric protocol value from 0 to 255
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol

- **icmp6**—Internet Control Message Protocol version 6
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-over-IP Encapsulation Protocol
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Transmission Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Clear sessions for the specified service set.

source-port *source-port*—(Optional) Clear sessions for the specified source port. The range of values is from 0 through 65535.

source-prefix *source-prefix*—(Optional) Clear sessions for the specified source prefix.

Required Privilege Level

clear

RELATED DOCUMENTATION

| [show services sessions](#)

List of Sample Output

[clear services sessions on page 772](#)

Output Fields

[Table 58 on page 771](#) lists the output fields for the **clear services sessions** command. Output fields are listed in the approximate order in which they appear.

Table 58: clear services sessions Output Fields

Field Name	Field Description
Interface	Name of an interface.
Service set	Name of the service set from which sessions are being cleared.

Table 58: clear services sessions Output Fields (*continued*)

Field Name	Field Description
Sessions marked for deletion	Number of sessions that are marked for deletion and are subsequently cleared.

Sample Output

clear services sessions

user@host>clear services sessions

Interface	Service set	Sessions marked for deletion
ms-0/0/0	sset	10

clear services sessions analysis

Syntax

```
clear services sessions analysis
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series MS-MPC.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Clear session statistics.

Options

interface *interface-name*—(Optional) Clear sessions statistics for the specified interface. The *interface-name* can be *vms-fpc/ pic/ port*.

Required Privilege Level

view

clear services stateful-firewall flows

Syntax

```
clear services stateful-firewall flows
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Clear stateful firewall flows. Issue this command to clear the stateful firewall flows for the specified option. The default option is "none", that is, to close all stateful firewall flows unless another option is specified.

Starting in Junos Release 14.1, the method for closing flows has changed. With the change, even for peak flows, the command prompt now returns to an active state after 30 seconds and the clear command completes in 90 to 120 seconds. In previous releases, closing peak flows could take as long as 4 minutes, after which the command prompt would return. Note too that during the first 30 seconds of issuing the command, the flows to be deleted remain visible in the **show services stateful-firewall flows** command output.

Options

none—Clear all stateful firewall flows.

destination-port *destination-port*—(Optional) Clear stateful firewall flows for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Clear stateful firewall flows for a particular destination prefix.

interface *interface-name*—(Optional) Clear stateful firewall flows for a particular interface. On M Series and T Series routers, the *interface-name* can be **ms-fpc/pic/port** or **rspnumber**.

protocol—(Optional) Clear stateful firewall flows for one of the following IP types:

- **number**—Numeric protocol value from 0 to 255.
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol

- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-over-IP Encapsulation Protocol
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Clear stateful firewall flows for a particular service set.

source-port *source-port*—(Optional) Clear stateful firewall flows for a particular source port. The range of values is from 0 through 65535.

source-prefix *source-prefix*—(Optional) Clear stateful firewall flows for a particular source prefix.

Required Privilege Level

view

RELATED DOCUMENTATION

[show services stateful-firewall flows](#) | 1035

List of Sample Output

[clear services stateful-firewall flows](#) on page 776

Output Fields

[Table 59 on page 775](#) lists the output fields for the **clear services stateful-firewall flows** command. Output fields are listed in the approximate order in which they appear.

Table 59: clear services stateful-firewall flows Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.

Table 59: clear services stateful-firewall flows Output Fields (*continued*)

Field Name	Field Description
Service set	Name of the service set from which flows are being cleared.
Conv removed	Number of conversations removed.

Sample Output

clear services stateful-firewall flows

user@host> **clear services stateful-firewall flows**

Interface	Service set	Conv removed
ms-0/3/0	svc_set_trust	0
ms-0/3/0	svc_set_untrust	0

clear services stateful-firewall sip-call

Syntax

```
clear services stateful-firewall sip-call
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced in Junos OS Release 7.4.

Description

Clear Session Initiation Protocol (SIP) call information in stateful firewall flows.

Options

none—Clear stateful firewall statistics for all interfaces and all service sets.

application-protocol—(Optional) Clear information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio

- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Clear information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Clear information for a particular destination prefix.

interface *interface-name*—(Optional) Clear information for a particular adaptive services interface. On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port* or *rspnumber*.

protocol—(Optional) Clear information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

- service-set** *service-set*—(Optional) Clear information for a particular service set.
- source-port** *source-port*—(Optional) Clear information for a particular source port. The range of values is 0 to 65535.
- source-prefix** *source-prefix*—(Optional) Clear information for a particular source prefix.

Required Privilege Level
view

RELATED DOCUMENTATION

| [show services stateful-firewall sip-call](#) | 1042

List of Sample Output
[clear services stateful-firewall sip-call on page 779](#)

Output Fields

[Table 60 on page 779](#) lists the output fields for the **clear services stateful-firewall sip-call** command. Output fields are listed in the approximate order in which they appear.

Table 60: clear services stateful-firewall sip-call Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of the service set from which flows are being cleared.
SIP calls removed	Number of SIP calls removed.

Sample Output

clear services stateful-firewall sip-call

user@host> clear services stateful-firewall sip-call

Interface	Service set	SIP calls removed
sp-0/3/0	test_sip_777	1

clear services stateful-firewall sip-register

Syntax

```
clear services stateful-firewall sip-register
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced in Junos OS Release 7.4.

Description

Clear Session Initiation Protocol (SIP) register information in stateful firewall flows.

Options

application-protocol—(Optional) Clear information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol

- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Clear information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Clear information for a particular destination prefix.

interface *interface*—(Optional) Clear information about a particular interface. On M Series and T Series routers, the *interface-name* can be **sp-fpc/pic/port** or **rspnumber**.

protocol—(Optional) Clear information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

- service-set *service-set***—(Optional) Clear information for a particular service set.
- source-port *source-port***—(Optional) Clear information for a particular source port. The range of values is 0 through 65535.
- source-prefix *source-prefix***—(Optional) Clear information for a particular source prefix.

Required Privilege Level
view

RELATED DOCUMENTATION

| [show services stateful-firewall sip-register](#) | 1048

List of Sample Output
[clear services stateful-firewall sip-register on page 782](#)

Output Fields

[Table 61 on page 782](#) lists the output fields for the **clear services stateful-firewall sip-register** command. Output fields are listed in the approximate order in which they appear.

Table 61: clear services stateful-firewall sip-register Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of the service set from which flows are being cleared.
SIP registration removed	Number of SIP registers removed.

Sample Output

clear services stateful-firewall sip-register
user@host> clear services stateful-firewall sip-register

Interface	Service set	SIP registration removed
sp-0/3/0	test_sip_777	1

clear services stateful-firewall statistics

Syntax

```
clear services stateful-firewall statistics
<interface interface-name>
<service-set service-set>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Clear stateful firewall statistics.

Options

none—Clear stateful firewall statistics for all interfaces and all service sets.

interface *interface-name*—(Optional) Clear stateful firewall statistics for the specified interface. On M Series and T Series routers, the *interface-name* can be *ms-fpc/pic/port* or *rspnumber*.

service-set *service-set*—(Optional) Clear stateful firewall statistics for the specified service set.

Required Privilege Level

view

RELATED DOCUMENTATION

[show services stateful-firewall statistics](#) | 1052

List of Sample Output

[clear services stateful-firewall statistics on page 783](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
clear services stateful-firewall statistics
```

```
user@host> clear services stateful-firewall statistics
```

clear services subscriber analysis

Syntax

```
clear services subscriber analysis
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series MS-MPC.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Clear information about the number of active subscribers on the services PIC.

Options

interface *interface-name*—(Optional) Display information about a particular interface.

Required Privilege Level

view

clear services web-filter statistics profile

Syntax

```
clear services web-filter statistics profile profile-name
<dns-filter-template template-name>
<fpc-slot fpc-slot pic-slot pic-slot>
<url-filter-template template-name>
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Description

Clear statistics for DNS request filtering or URL filtering for the specified filter profile.

Options

dns-filter-template *template-name*—(Optional) Name of the DNS filter template for which statistics are cleared.

fpc-slot *fpc-slot* pic-slot *pic-slot*—(Optional) Location of the services PIC for which statistics are cleared.

profile *profile-name*—Name of the filter profile for which statistics are cleared.

url-filter-template *template-name*—(Optional) Name of the URL filter template for which statistics are cleared.

Required Privilege Level

clear

RELATED DOCUMENTATION

[DNS Request Filtering for Disallowed Website Domains | 311](#)

[Configuring URL Filtering | 335](#)

List of Sample Output

[clear services web-filter statistics profile on page 786](#)

Output Fields

When you enter this command, the statistics for DNS request filtering are cleared. There is no specific output.

Sample Output

```
clear services web-filter statistics profile
```

```
user@host> clear services web-filter statistics profile profile1
```

request services web-filter update dns-filter-database

Syntax

```
request services web-filter update dns-filter-database filename
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

When you make changes to the domain filter database file, which is used in filtering DNS requests for disallowed domains, apply the changes.

Options

filename—File name of the database file.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[DNS Request Filtering for Disallowed Website Domains](#) | 311

request services web-filter validate dns-filter-file-name

Syntax

```
request services web-filter validate dns-filter-file-name filename hash-key key-string hash-method hash-method-name
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Validate the file format of the domain filter database file, which is used in filtering DNS requests for disallowed domains.

Options

filename—File name of the database file.

hash-method-name—Hash method you used to produce the hashed domain name values in the database file.

key-string—Hash key you used to produce the hashed domain name values in the database file.

Required Privilege Level

maintenance

RELATED DOCUMENTATION

[DNS Request Filtering for Disallowed Website Domains](#) | 311

request system disable unified-services

Syntax

```
request system disable unified-services
```

Release Information

Command introduced in Junos OS Release 19.3R1 on MX Series running Next Gen Services.

Description

Disable Next Gen Services services on the MX Series.

Before you disable Next Gen Services, delete any router configuration for services. This includes configuration under the **[edit services]** hierarchy, configuration for services interfaces, and any configuration that refers to services interfaces.

After you enter **request system enable unified-services**, reboot the chassis.

Required Privilege Level

RELATED DOCUMENTATION

[Enabling and Disabling Next Gen Services](#) | 67

List of Sample Output

[request system disable unified-services on page 789](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

request system disable unified-services

```
user@host> request system disable unified-services
```

```
Before disabling unified services, please move to baseline configuration.  
Are above conditions satisfied ? [yes,no]
```

request system enable unified-services

Syntax

```
request system enable unified-services
```

Release Information

Command introduced in Junos OS Release 19.3R1 on MX Series running Next Gen Services.

Description

Enable Next Gen Services services on the MX Series.

Before you enable Next Gen Services, delete any router configuration for services. This includes configuration under the **[edit services]** hierarchy, configuration for services interfaces, and any configuration that refers to services interfaces.

After you enter **request system enable unified-services**, reboot the chassis.

In Junos node slicing, you can enable unified services at guest network function (GNF), by using the CLI **request system enable unified-services** at GNF.

Required Privilege Level

RELATED DOCUMENTATION

[Enabling and Disabling Next Gen Services](#) | 67

List of Sample Output

[request system enable unified-services on page 790](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

```
request system enable unified-services
```

```
user@host> request system enable unified-services
```

```
Before enabling unified services, please move to baseline configuration.  
Are above conditions satisfied ? [yes,no]
```


show interfaces load-balancing (Aggregated Multiservices)

Syntax

```
show interfaces load-balancing
<detail>
<interface-name>
```

Release Information

Command introduced in Junos OS Release 11.4.

interface-name option added in Junos OS Release 16.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display information about the aggregated multiservices interface (AMS) as well as its individual member interfaces and the status of the replication state.

Options

none—Display standard information about status of all AMS interfaces.

detail—(Optional) Display detailed status of all AMS interfaces.

interface-name—(Optional) Name of the aggregated multiservices interface (**ams**). If this is omitted, then the information for all the aggregated multiservices interfaces, including those used in control plane redundancy and high availability (HA) for service applications, is displayed.

Required Privilege Level

view

RELATED DOCUMENTATION

Understanding Aggregated Multiservices Interfaces

[Understanding Aggregated Multiservices Interfaces for Next Gen Services | 353](#)

Example: Configuring an Aggregated Multiservices Interface (AMS)

List of Sample Output

[show interfaces load-balancing on page 793](#)

[show interfaces load-balancing detail on page 794](#)

[show interfaces load-balancing detail \(Specific Interface\) on page 794](#)

Output Fields

Table 62 on page 792 lists the output fields for the **show interfaces load-balancing** (aggregated multiservices interfaces) command. Output fields are listed in the approximate order in which they appear.

Table 62: Aggregated Multiservices show interfaces load-balancing Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the aggregated multiservices (AMS) interface.	detail none
State	Status of AMS interfaces: <ul style="list-style-type: none"> • Coming Up—Interface is becoming operational. • Members Seen—Member interfaces (mams) are available. • Up—Interface is configured and operational. • Wait for Members—Member interfaces (mams) are not available. • Wait Timer—Interface is waiting for member interfaces (mams) to come online. 	detail none
Last change	Time (in <i>hh:mm:ss [hours:minutes:seconds]</i> format) when the state last changed.	detail none
Members	Number of member interfaces (mams-).	none specified
Member count	Number of member PICs (mams) that are part of the aggregated interface.	detail none
HA Model	High availability (HA) model supported on the interface. <ul style="list-style-type: none"> • Many-to-One—The preferred backup Multiservices PIC, in hot standby mode, backs up one or more (N) active Multiservices PICs. • One-to-One—The preferred backup Multiservices PIC, in hot standby mode, backs up only one active Multiservices PIC. <p>NOTE: One-to-One is not supported on MX-SPC3 cards.</p>	detail none

Table 62: Aggregated Multiservices show interfaces load-balancing Output Fields (continued)

Field Name	Field Description	Level of Output
Members	<p>Information about the member interfaces:</p> <ul style="list-style-type: none"> • Interface—Name of the member interface. • Weight—Not applicable for the current release. • State—State of the member interface (mams-). <ul style="list-style-type: none"> • Active—Member is an active member. • Backup—Member is a backup. • Discard—Member has not yet rejoined the ams interface after failure. • Down—Member has not yet powered on. • Inactive—Member has failed to rejoin the ams interface within the configured rejoin-timeout. • Invalid—Multiservices PIC corresponding to the member interface has been configured but is not physically present in the chassis. 	detail
Sync-state	<p>Synchronization (sync) status of the control plane redundancy. The sync state is displayed only when the ams interface is Up.</p> <ul style="list-style-type: none"> • Interface—Name of the member interface. • Status—Synchronization status of the member interfaces. <ul style="list-style-type: none"> • In progress—The active member is currently synchronizing its state information with the backup member. • In sync—The active member has finished synchronizing its state information with the backup and the backup is ready to take over if the active member fails. • NA (Not applicable)—The backup member is not yet ready to synchronize with the active (primary) member. This condition may occur if the backup is still powered off or still booting. • Unknown—The daemons are still initializing and the state information is unavailable. 	detail

Sample Output

```
show interfaces load-balancing
```

```
user@host> show interfaces load-balancing
```

Interface	State	Last change	Members	HA Model
ams0	Up	00:10:02	4	Many-to-One

show interfaces load-balancing detail

```
user@host> show interfaces load-balancing detail
```

```
Load-balancing interfaces detail
Interface      : ams0
State          : Up
Last change    : 00:10:23
Member count   : 4
HA Model       : Many-to-One
Members        :
  Interface    Weight  State
  mams-4/0/0   10      Active
  mams-4/1/0   10      Active
  mams-5/0/0   10      Active
  mams-5/1/0   10      Backup
Sync-state     :
  Interface    Status
  mams-4/0/0   Unknown
  mams-4/1/0   Unknown
  mams-5/0/0   Unknown
```

show interfaces load-balancing detail (Specific Interface)

```
user@host> show interfaces load-balancing ams0 detail
```

```
Load-balancing interfaces detail
Interface      : ams0
State          : Up
Last change    : 00:11:28
Member count   : 4
HA Model       : Many-to-One
Members        :
  Interface    Weight  State
  mams-4/0/0   10      Active
  mams-4/1/0   10      Active
  mams-5/0/0   10      Active
  mams-5/1/0   10      Backup
```

Sync-state :

Interface	Status
mams-4/0/0	Unknown
mams-4/1/0	Unknown
mams-5/0/0	Unknown

show log

List of Syntax

[Syntax on page 796](#)

[Syntax \(QFX Series and OCX Series\) on page 796](#)

[Syntax \(TX Matrix Router\) on page 796](#)

Syntax

```
show log  
<filename | user <username>>
```

Syntax (QFX Series and OCX Series)

```
show log filename  
<device-type (device-id | device-alias)>
```

Syntax (TX Matrix Router)

```
show log  
<all-lcc | lcc number | scc>  
<filename | user <username>>
```

Release Information

Command introduced before Junos OS Release 7.4.

Command introduced in Junos OS Release 9.0 for EX Series switches.

Command introduced in Junos OS Release 11.1 for the QFX Series.

Option *device-type (device-id | device-alias)* is introduced in Junos OS Release 13.1 for the QFX Series.

Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Command introduced in Junos OS Release 20.3R1 for Next Gen Services on MX240, MX480, and MX960 routers.

Description

List log files, display log file contents, or display information about users who have logged in to the router or switch.

NOTE: On MX Series routers, modifying a configuration to replace a service interface with another service interface is treated as a catastrophic event. When you modify a configuration, the entire configuration associated with the service interface—including NAT pools, rules, and service sets—is deleted and then re-created for the newly specified service interface. If there are active sessions associated with the service interface that is being replaced, these sessions are deleted and the NAT pools are then released, which leads to the generation of the NAT_POOL_RELEASE system log messages. However, because NAT pools are already deleted as a result of the catastrophic configuration change and no longer exist, the NAT_POOL_RELEASE system log messages are not generated for the changed configuration.

Options

none—List all log files.

<all-lcc | lcc *number* | scc>—(Routing matrix only)(Optional) Display logging information about all T640 routers (or line-card chassis) or a specific T640 router (replace ***number*** with a value from 0 through 3) connected to a TX Matrix router. Or, display logging information about the TX Matrix router (or switch-card chassis).

device-type—(QFabric system only) (Optional) Display log messages for only one of the following device types:

- **director-device**—Display logs for Director devices.
- **infrastructure-device**—Display logs for the logical components of the QFabric system infrastructure, including the diagnostic Routing Engine, fabric control Routing Engine, fabric manager Routing Engine, and the default network Node group and its backup (NW-NG-0 and NW-NG-0-backup).
- **interconnect-device**—Display logs for Interconnect devices.
- **node-device**—Display logs for Node devices.

NOTE: If you specify the **device-type** optional parameter, you must also specify either the **device-id** or **device-alias** optional parameter.

(*device-id* | *device-alias*)—If a device type is specified, display logs for a device of that type. Specify either the device ID or the device alias (if configured).

filename—(Optional) Display the log messages in the specified log file. For the routing matrix, the filename must include the chassis information.

NOTE: The *filename* parameter is mandatory for the QFabric system. If you did not configure a syslog filename, specify the default filename of **messages**.

user <username>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include **username**, display logging information about the specified user.

Required Privilege Level

trace

RELATED DOCUMENTATION

| [syslog \(System\)](#)

List of Sample Output

[show log on page 798](#)

[show log filename on page 799](#)

[show log filename \(QFabric System\) on page 801](#)

[show log user on page 802](#)

[show log accepted-traffic \(SRX4600, SRX5400, SRX5600, and SRX5800\) on page 802](#)

Sample Output

show log

user@host> **show log**

```
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
```



```
-rw-r--r--  1 root  bin      1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin        19656 Oct  1 19:37 wttmp
```

show log filename

user@host> show log rpd

```
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr 192.0.2.21
nhop type local nhop 192.0.2.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr 192.0.2.22
nhop type unicast nhop 192.0.2.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

user@host:LSYS1> show log flow_lsys1.log

```
Nov  7 07:34:09 07:34:09.491800:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491809:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
table lock

Nov  7 07:34:09 07:34:09.491840:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491841:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
table lock

Nov  7 07:34:09 07:34:09.491854:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh
0x0

Nov  7 07:34:09 07:34:09.491868:CID-0:THREAD_ID-00:LSYS_ID-01:RT:got route table
lock

Nov  7 07:34:09 07:34:09.491869:CID-0:THREAD_ID-00:LSYS_ID-01:RT:released route
```

```
table lock
```

```
Nov  7 07:34:09 07:34:09.491881:CID-0:THREAD_ID-00:LSYS_ID-01:RT:cache final sw_nh
0x0
```

```
user@host:TSYS1> show log flow_tsys1.log
```

```
Nov  7 13:21:47
13:21:47.217744:CID-0:THREAD_ID-05:LSYS_ID-32:RT:<192.0.2.0/0->198.51.100.0/9011;1,0x0>
:

Nov  7 13:21:47 13:21:47.217747:CID-0:THREAD_ID-05:LSYS_ID-32:RT:packet [84] ipid
= 39281, @0x7f490ae56d52

Nov  7 13:21:47 13:21:47.217749:CID-0:THREAD_ID-05:LSYS_ID-32:RT:----
flow_process_pkt: (thd 5): flow_ctxt type 0, common flag 0x0, mbuf 0x4882b600,
rtbl7

Nov  7 13:21:47 13:21:47.217752:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow process pak
fast ifl 88 in_ifp lt-0/0/0.101

Nov  7 13:21:47 13:21:47.217753:CID-0:THREAD_ID-05:LSYS_ID-32:RT:
lt-0/0/0.101:192.0.2.0->198.51.100.0, icmp, (0/0)

Nov  7 13:21:47 13:21:47.217756:CID-0:THREAD_ID-05:LSYS_ID-32:RT: find flow: table
0x11d0a2680, hash 20069(0xffff), sa 192.0.2.0, da 198.51.100.0, sp 0, d0

Nov  7 13:21:47 13:21:47.217760:CID-0:THREAD_ID-05:LSYS_ID-32:RT:Found: session
id 0x12. sess tok 28685

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow got session.

Nov  7 13:21:47 13:21:47.217761:CID-0:THREAD_ID-05:LSYS_ID-32:RT: flow session
id 18

Nov  7 13:21:47 13:21:47.217763:CID-0:THREAD_ID-05:LSYS_ID-32:RT: vector bits 0x200
vector 0x84ae85f0

Nov  7 13:21:47 13:21:47.217764:CID-0:THREAD_ID-05:LSYS_ID-32:RT:set nat
0x11e463550(18) timeout const to 2

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT: set_nat_timeout
2 on session 18
```

```

Nov  7 13:21:47 13:21:47.217765:CID-0:THREAD_ID-05:LSYS_ID-32:RT:refresh nat
0x11e463550(18) timeout to 2

Nov  7 13:21:47 13:21:47.217767:CID-0:THREAD_ID-05:LSYS_ID-32:RT:insert usp tag
for apps

Nov  7 13:21:47 13:21:47.217768:CID-0:THREAD_ID-05:LSYS_ID-32:RT:mbuf 0x4882b600,
exit nh 0xffffb0006

```

show log filename (QFabric System)

user@qfabric> show log messages

```

Mar 28 18:00:06 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:06 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 2159)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1486
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 2191)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 242726)
Mar 28 18:00:07 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:07 ED1492
chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 2, jnxFruL3Index 0,
jnxFruName PIC: @ 0/1/*, jnxFruType 11, jnxFruSlot 0, jnxFruOfflineReason 2,
jnxFruLastPowerOff 0, jnxFruLastPowerOn 242757)
Mar 28 18:00:16 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:16 ED1486 file:
UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:27 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:27 ED1486 file:
UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)
Mar 28 18:00:50 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:50
_DCF_default__NW-INE-0_RE0_ file: UI_COMMIT: User 'root' requested 'commit'
operation (comment: none)

```

```

Mar 28 18:00:55 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:00:55 ED1492 file:
  UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:01:10 qfabric file: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:01:10 ED1492 file:
  UI_COMMIT: User 'root' requested 'commit' operation (comment: none)
Mar 28 18:02:37 qfabric chassisd: QFABRIC_INTERNAL_SYSLOG: Mar 28 18:02:37 ED1491
  chassisd: CHASSISD_SNMP_TRAP10: SNMP trap generated: FRU power on
(jnxFruContentsIndex 8, jnxFruL1Index 1, jnxFruL2Index 1, jnxFruL3Index 0,
jnxFruName PIC: 48x 10G-SFP+ @ 0/0/*, jnxFruType 11, jnxFruSlot 0,
jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 33809)

```

show log user

```
user@host> show log user
```

usera	mg2546		Thu Oct 1 19:37	still logged in
usera	mg2529		Thu Oct 1 19:08 - 19:36	(00:28)
usera	mg2518		Thu Oct 1 18:53 - 18:58	(00:04)
root	mg1575		Wed Sep 30 18:39 - 18:41	(00:02)
root	ttyp2	aaa.bbbb.com	Wed Sep 30 18:39 - 18:41	(00:02)
userb	ttyp1	192.0.2.0	Wed Sep 30 01:03 - 01:22	(00:19)

show log accepted-traffic (SRX4600, SRX5400, SRX5600, and SRX5800)

```
user@host> show log accepted-traffic
```

```

Jul 17 20:26:04 sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.5/2->4.4.4.2/63 0x0 None 3.3.3.5/2->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2617282058 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04 sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.4/4->4.4.4.2/63 0x0 None 3.3.3.4/4->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162754 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04 sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.4/1->4.4.4.2/63 0x0 None 3.3.3.4/1->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162755 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04 sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.3/0->4.4.4.2/63 0x0 None 3.3.3.3/0->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162752 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04 sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created

```

```

3.3.3.5/5->4.4.4.2/63 0x0 None 3.3.3.5/5->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162751 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A
Jul 17 20:26:04 sourpunch RT_FLOW: RT_FLOW_SESSION_CREATE: session created
3.3.3.3/3->4.4.4.2/63 0x0 None 3.3.3.3/3->4.4.4.2/63 0x0 N/A N/A N/A N/A 17 p2
TRUST UNTRUST 2550162753 N/A(N/A) xe-7/0/0.0 UNKNOWN UNKNOWN UNKNOWN N/A N/A -1
N/A N/A N/A

```

show services alg conversations

Syntax

```
show services alg conversations
<brief >
<application-protocol protocol>
<extensive>
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 10.4.

h323 option introduced in Junos OS Release 17.1.

ike-esp-nat option introduced in Junos OS Release 17.1.

Description

Display ALG information for Junos OS extension-provider packages.

NOTE: In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.

Options

none—Display standard information about all Junos OS extension-provider packages ALG sessions.

brief —(Optional) Display the specified level of output.

application-protocol—(Optional) Display information about one of the following application protocols:

dce-rpc—Distributed Computing Environment-Remote Procedure Call protocols

dce-rpc-portmap—Distributed Computing Environment-Remote Procedure Call protocols portmap service

dns—Domain Name System protocol

ftp—File Transfer Protocol

h323—H323 protocol

ike-esp-nat—IKE ALG

pptp—Point-to-Point Tunneling Protocol

rpc—Remote Procedure Call protocol

rpc-portmap—Remote Procedure Call protocol portmap service

rtsp—Real-Time Streaming Protocol

rsh—Remote Shell

sip—Session Initiation Protocol

sql—SQLNet

talk—Talk Program

extensive—Display extensive information

interface *interface-name*—(Optional) Display information about a particular interface.

Required Privilege Level

view

List of Sample Output

[show services alg conversations on page 806](#)

[show services alg conversations brief on page 806](#)

[show services alg conversations extensive on page 807](#)

[show services alg conversations application-protocol on page 807](#)

[show services alg conversations interface on page 811](#)

Output Fields

[Table 63 on page 805](#) lists the output fields for the **show services alg conversations** command. Output fields are listed in the approximate order in which they appear.

Table 63: show services alg conversations Output Fields

Field Name	Field Description
Interface	Name of the interface.
ALG	Name of the ALG in use.
Number of conversations	Number of ALG conversations open. A conversation is a group of parent and child sessions.
Group ID	Numeric identifier for the session.
Parent session status	Status of the parent session: <ul style="list-style-type: none"> • Active • Closed

Table 63: show services alg conversations Output Fields (*continued*)

Field Name	Field Description
Parent session ID	Numeric identifier for the parent session.
Protocol	Protocol used for the parent session.
Forward Flow	The source and destination prefixes for forward flow.
Reverse Flow	The source and destination prefixes for reverse flow.
Child session status	Status of the child session: <ul style="list-style-type: none"> • Active • Closed
Child session ID	Numeric identifier for the child session.
Number of Resources	Total number of active child sessions associated with the parent session.
Resource ID	Numeric identifier for the resources associated with the parent session.
Protocol	Protocol used for the child session.

Sample Output

show services alg conversations

```
user@host> show services alg conversations
```

```
Interface name: ms-2/1/0
ALG : SQLV2 ALG, State : active
Number of conversations: 1
Parent session status: closed
Child session : 1, protocol: TCP
Forward Flow : {10.50.50.2:37244 -> 10.40.40.10:4334}
Reverse Flow : {10.40.40.10:4334 -> 10.11.11.10:37244}
```

show services alg conversations brief

The output for the **show services alg conversations brief** command is identical to that for the **show services alg conversations** command. For sample output, see [show services alg conversations on page 806](#).

show services alg conversations extensive

user@host> **show services alg conversations extensive**

```

Interface name: ms-1/0/0
ALG : H323 ALG, State : active
Number of conversations: 1
Group ID : 3499913712, State : active
Parent session state: active
Parent session ID: 33554433, protocol : TCP
Forward Flow : {198.51.100.2:30000 -> 192.0.2.2:1720}
Reverse Flow : {192.0.2.2:1720 -> 203.0.113.1:57730}
Number of resources: 4
Resource ID: 3499927656, State: active
Number of sessions: 1
Child session ID: 33554436, protocol : UDP
Forward Flow : {198.51.100.2:5086 -> 192.0.2.2:5090}
Reverse Flow : {192.0.2.2:5090 -> 203.0.113.3:55916}
Resource ID: 3499927376, State: active
Number of sessions: 1
Child session ID: 67108867, protocol : UDP
Forward Flow : {192.0.2.2:5091 -> 203.0.113.3:55917}
Reverse Flow : {198.51.100.2:5087 -> 192.0.2.2:5091}
Resource ID: 3499926816, State: active
Number of sessions: 1
Child session ID: 33554438, protocol : UDP
Forward Flow : {198.51.100.2:5089 -> 192.0.2.2:5093}
Reverse Flow : {192.0.2.2:5093 -> 203.0.113.2:63435}
Resource ID: 3499926536, State: active
Number of sessions: 1
Child session ID: 33554437, protocol : UDP
Forward Flow : {198.51.100.2:5088 -> 192.0.2.2:5092}
Reverse Flow : {192.0.2.2:5092 -> 203.0.113.2:63434}
ALG : RAS ALG, State : active
Number of conversations: 1
Group ID : 799037592, State : active
Parent session state: closed
Number of resources: 0

```

show services alg conversations application-protocol

This command has the same output for the rpc, dce-rpc, rpc-portmap and dce-rpc-portmap ALGs.

user@router> **show services alg conversations application-protocol rpc**

```

Interface name: ms-1/1/0
ALG : SUNRPC ALG, State : active
Number of conversations: 2
Parent session status: closed
Child session : 1, protocol: UDP
Forward Flow : {192.168.203.198:1019 -> 192.168.203.194:2049}
Reverse Flow : {192.168.203.194:2049 -> 192.168.203.198:1019}
Child session : 2, protocol: UDP
Forward Flow : {192.168.203.198:36595 -> 192.168.203.194:2049}
Reverse Flow : {192.168.203.194:2049 -> 192.168.203.198:36595}
Parent session status: closed
Child session : 1, protocol: UDP
Forward Flow : {192.168.203.198:954 -> 192.168.203.194:613}
Reverse Flow : {192.168.203.194:613 -> 192.168.203.198:954}
Child session : 2, protocol: UDP
Forward Flow : {192.168.203.198:53836 -> 192.168.203.194:613}
Reverse Flow : {192.168.203.194:613 -> 192.168.203.198:53836}

```

user@router> **show services alg conversations application-protocol dns**

```

Interface name: ms-1/1/0
ALG : DNS ALG, State : active
Number of conversations: 1
Parent session status: closed
Child session : 1, protocol: UDP
Forward Flow : {192.168.203.198:1019 -> 192.168.203.194:2049}
Reverse Flow : {192.168.203.194:2049 -> 192.168.203.198:1019}

```

user@router> **show services alg conversations application-protocol ftp**

```

Interface name: ms-1/1/0
ALG : DNS ALG, State : active
Number of conversations: 1
Parent session status: closed
Child session : 1, protocol: UDP
Forward Flow : {192.168.203.198:53836 -> 192.168.203.194:613}
Reverse Flow : {192.168.203.194:613 -> 192.168.203.198:53836}

```

user@router> **show services alg conversations application-protocol ike-esp-nat**

```

Interface name: ms-2/2/0
ALG : IKE ALG, State : active
  Number of conversations: 1
    Parent session status: closed
    Child session : 1, protocol: ESP
      Forward Flow : {198.51.100.101:2623 -> 203.0.113.1:46838}
      Reverse Flow : {192.0.2.101:46838 -> 198.51.10.101:2623}
    Child session : 2, protocol: ESP
      Forward Flow : {192.0.2.101:2666 -> 198.51.10.101:57882}
      Reverse Flow : {198.51.10.101:57882 -> 203.0.113.1:2666}

```

user@router> **show services alg conversations application-protocol pptp**

```

Interface name: ms-2/0/0
ALG : PPTP ALG, State : active
  Number of conversations: 1
    Parent session status: active
    Parent session : 1, protocol : TCP
      Forward Flow : {192.0.2.10:1511 -> 198.51.100.10:1723}
      Reverse Flow : {198.51.100.10:1723 -> 192.0.2.10:1511}
    Child session : 1, protocol: GRE
      Forward Flow : {192.0.2.10:0 -> 198.51.100.10:49913}
      Reverse Flow : {198.51.100.10:49913 -> 192.0.2.10:65001}
    Child session : 2, protocol: GRE
      Forward Flow : {198.51.100.10:0 -> 192.0.2.10:0}
      Reverse Flow : {192.0.2.10:0 -> 198.51.100.10:65000}

```

user@router> **show services alg conversations application-protocol rtsp**

```

Interface name: ms-0/1/0
ALG : RTSP ALG, State : active
  Number of conversations: 1
    Parent session : 1, protocol : TCP
      Forward Flow : {198.51.100.2:3985 -> 192.0.2.1:554}
      Reverse Flow : {203.0.113.2:554 -> 198.51.100.2:3985}
    Child session : 1, protocol: UDP
      Forward Flow : {203.0.113.2:35859 -> 198.51.100.2:38159}
      Reverse Flow : {198.51.100.2:38159 -> 192.0.2.1:35859}
    Child session : 2, protocol: UDP
      Forward Flow : {203.0.113.2:35859 -> 198.51.100.2:37391}
      Reverse Flow : {198.51.100.2:37391 -> 192.0.2.1:35859}

```

user@router> **show services alg conversations application-protocol rsh**

```

Interface name: ms-0/1/0
ALG : RSH ALG, State : active
Number of conversations: 1
  Parent session : 1, protocol : TCP
    Forward Flow : {198.51.100.2:3985 -> 192.0.2.1:554}
    Reverse Flow : {203.0.113.2:554 -> 198.51.100.2:3985}
  Child session : 1, protocol: UDP
    Forward Flow : {203.0.113.2:35859 -> 198.51.100.2:38159}
    Reverse Flow : {198.51.100.2:38159 -> 192.0.2.1:35859}

```

user@router> show services alg conversations application-protocol sip

```

Interface name: ms-1/1/0
ALG : SIP ALG, State : active
Number of conversations: 1
  Parent session status: active
  Parent session : 1, protocol : UDP
    Forward Flow : {192.0.2.2:5060 -> 198.51.100.2:5060}
    Reverse Flow : {198.51.100.2:5060 -> 203.0.113.2:5060}
  Child session : 1, protocol: UDP
    Forward Flow : {192.0.2.2:6000 -> 198.51.100.2:12442}
    Reverse Flow : {198.51.100.2:12442 -> 203.0.113.2:6000}

```

user@router> show services alg conversations application-protocol sql

```

Interface name: ms-2/0/0
ALG : SQLV2 ALG, State : active
Number of conversations: 1
  Parent session : 1, protocol : 0
    Forward Flow : {0.0.0.0:0 -> 0.0.0.0:0}
    Reverse Flow : {0.0.0.0:0 -> 0.0.0.0:0}
  Child session : 1, protocol: TCP
    Forward Flow : {203.0.113.2:19099 -> 198.51.100.10:32773}
    Reverse Flow : {198.51.100.10:32773 -> 192.0.2.1:19099}

```

user@router> show services alg conversations application-protocol talk

```

Interface name: ms-0/1/0
ALG : TALK ALG, State : active
Number of conversations: 1
  Parent session : 1, protocol : TCP
    Forward Flow : {198.51.100.2:3985 -> 192.0.2.1:554}

```

```
Reverse Flow : {203.0.113.2:554 -> 198.51.2:3985}  
Child session : 1, protocol: UDP  
Forward Flow : {203.0.113.2:35859 -> 198.51.2:38159}  
Reverse Flow : {198.51.2:38159 -> 192.0.2.1:35859}
```

show services alg conversations interface

user@router> show services alg conversations interface ms-1/1/0

```
ALG : FTP ALG, State : active  
Number of conversations: 1  
Parent session status: active  
Parent session : 1, protocol : TCP  
Forward Flow : {10.20.20.10:47164 -> 10.30.30.30:21}
```

show services alg statistics

Syntax

```
show services alg statistics
<application-protocol protocol>
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 10.4.

h323 option introduced in Junos OS Release 17.1.

ike-esp-nat option introduced in Junos OS Release 17.1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display ALG statistics for Junos OS extension-provider packages.

NOTE: In Junos OS releases earlier than 12.3, the extension-provider packages were variously referred to as Junos Services Framework (JSF), MP-SDK, and eJunos.

Options

application-protocol—(Optional) Display statistics for one of the following application protocols:

dce-rpc—Distributed Computing Environment-Remote Procedure Call protocols

dce-rpc-portmap—Distributed Computing Environment-Remote Procedure Call protocols portmap service

dns—Domain Name System protocol

ftp—File Transfer Protocol

h323—H323 protocol

ike-esp-nat—IKE ALG

pptp—Point-to-Point Tunneling Protocol

rpc—Remote Procedure Call protocol

rpc-portmap—Remote Procedure Call protocol portmap service

rtsp—Real-Time Streaming Protocol

rsh—Remote Shell

sip—Session Initiation Protocol

sql—SQLNet

talk—Talk Program

tftp—Trivial File Transfer Protocol

interface *interface-name*—(Optional) Display information about a particular interface.

Required Privilege Level

view

List of Sample Output

[show services alg statistics application-protocol on page 822](#)

[show services alg statistics interface on page 827](#)

Output Fields

[Table 64 on page 813](#) lists the output fields for the **show services alg statistics** command. Output fields are listed in the approximate order in which they appear.

Table 64: show services alg statistics Output Fields

Field Name	Field Description
Interface	Name of the interface.
ALG statistics	Name of the ALG for which the statistics are displayed.
Packets with wrong header	Number of packets with wrong header.
Non epm 3.0 packets	Number of non epm 3.0 packets.
Packets with type mismatch	Number of packets with type mismatch.
Packets with id mismatch	Number of packets with id mismatch.
Packets with call mismatch	Number of packets with call mismatch.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
Packets fragmented	Number of packets fragmented.
Packets queued	Number of packets queued.
Packets dropped	Number of packets dropped.
Packets released	Number of packets released.
Invalid packets received	Number of invalid packets received.
Reply packets received	Number of reply packets received.
Oversized packets received	Number of oversized packets received.
ALG parser errors	Number of parsing failed errors.
Packets translated	Number of packets translated.
H323 total calls	Total number of audio/video calls that have been established.
H323 active calls	Current number of active H.323 calls.
H323 gate install failed	Number of gate installation failures for child sessions.
H323 pinhole opened too late	Number of H323 parent sessions that released the resources before pinhole creation.
H323 pinhole hit dropped	Number of H323 gate hits that have been dropped.
H323 gate timeout failed	Number of gate timeout failures due to an error.
H323 packets dropped	Number of packets dropped.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
H323 get virtual ctx failed	Number of failures to get the session virtualization ctx information.
H323 obj alloc failed	Number of memory allocation failures for H323 session cookie.
H323 group alloc failed	Number of H323 session resource/group memory allocation failures.
H323 ce alloc failed	Number of H323 session call entity object memory allocation failures.
H323 Q931 decode error	Number of errors in decoding Q931 packets.
H323 H245 decode error	Number of errors in decoding H245 packets.
H323 Q931 process error	Number of errors in processing Q931 packets.
H323 H245 process error	Number of errors in processing H245 packets.
H323 do nat failed	Number of NAT translation failures after packet decode.
H323 do rm failed	Number of H323 vsip table creation failures.
H323 dscp marked	Number of Differentiated Services code point (DSCP) packets marked.
H323 dscp marked error	Number of Differentiated Services code point (DSCP) packets marked as errors.
RAS obj alloc failed	Number of RAS session object memory allocation failures.
RAS group alloc failed	Number of RAS session group memory allocation failures.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
RAS packets dropped	Number of RAS packets dropped.
RAS packet exists in cookie error	Number of times that some packets exist in existing RAS sessions cookie.
RAS decode error	Number of errors in decoding RAS packets.
RAS flood error	Number of gatekeeper requests that were dropped because of too many RAS request messages.
RAS do nat failed	Number of RAS session payload IP translation errors.
PPTP Objects Active	Number of PPTP objects active.
PPTP Objects Total	Number of PPTP objects in total.
PPTP Objects Error	Number of PPTP objects having errors.
PPTP ASL Group Active	Number of PPTP groups active.
PPTP ASL Group Total	Number of PPTP groups in total.
PPTP ASL Group Error	Number of PPTP groups having errors.
PPTP Packets received	Number of PPTP packets received.
PPTP Packets Discarded	Number of PPTP packets discarded.
PPTP Packets Free	Number of PPTP packets freed.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
PPTP OCRQ Received	Number of Outgoing Call Requests received.
PPTP OCRQ Discarded	Number of Outgoing Call Requests discarded.
PPTP OCRP Received	Number of Outgoing Call Packets received.
PPTP OCRP Discarded	Number of Outgoing Call Packets discarded.
PPTP WEN(SLI) Received	Number of WEN (SLI) packets received.
PPTP WEN(SLI) Discarded	Number of WEN (SLI) packets discarded.
PPTP CCRQ-CDSN Received	Number of Call Clear Requests received.
PPTP CDSN Received	Number of Call Disconnection Notifications received.
PPTP CCRQ-CDSN Discarded	Number of Call Clear Requests discarded.
PPTP Session Create	Number of PPTP sessions created.
PPTP Session Destroy	Number of PPTP sessions destroyed.
PPTP Gate Create	Number of PPTP gates created.
PPTP Gate Hit	Number of PPTP gates hit.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
PPTP Gate Timeout	Number of PPTP gates timed out.
PPTP NAT Events	Number of NAT events.
PPTP DO-NAT Total	Number of DO NATs in total.
PPTP DO-NAT Ok	Number of DO NATs okay.
PPTP DO-NAT Pending	Number of DO NATs pending.
PPTP DO-NAT Fail	Number of DO NATs failed.
PPTP DO-RM Total	Number of DO RMs in total.
PPTP DO-RM Ok	Number of DO RMs okay.
PPTP DO-RM Pending	Number of DO RMs pending.
PPTP DO-RM Fail	Number of DO RMs failed.
PPTP NAT-ASYNC Total	Number of NAT-ASYNCs in total.
PPTP NAT-ASYNC Invalid	Number of NAT-ASYNCs invalid.
PPTP NAT-ASYNC Error1	Number of NAT-ASYNCs error1.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
PPTP NAT-ASYNC Error2	Number of NAT-ASYNCs error2.
PPTP ASL Hole Ok	Number of ASYNC holes okay.
PPTP ASL Hole Error	Number of ASYNC hole errors.
PPTP ASL First Hit	Number of ASYNC holes first hit.
PPTP ASL Hole Timeout	Number of ASYNC holes timed out.
PPTP ASL Invalid	Number of ASYNC holes invalid.
PPTP NAT Ctx Free	Number of NAT Ctxs free.
PPTP Create Resource Error	Number of create resource errors.
PPTP set S2C hole error	Number of server-to-client hole errors.
PPTP set C2S hole error	Number of client-to-server hole errors.
PPTP Inbrk error	Number of PPTP Inbrk errors.
PPTP Mpool Create Error	Number of Mpool create errors.
PPTP RM register client Error	Number of client registration errors.
Call packet with rpcbind2	Number of call packets with rpcbind2.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
Call packet with rpcbind3	Number of call packets with rpcbind3.
Call packet with rpcbind4	Number of call packets with rpcbind4.
Invalid rpcbind call	Number of invalid rpcbind calls.
Reply packet with rpcbind2	Number of reply packets with rpcbind2.
Reply packet with rpcbind3	Number of reply packets with rpcbind3.
Reply packet with rpcbind4	Number of reply packets with rpcbind4.
Invalid rpcbind reply	Number of invalid rpcbind replies.
Packets exceeded maximum length	Number of packets exceeding maximum length.
Packets dropped by ALG	Number of packets dropped by the ALG.
Number of describe messages received	Number of describe messages received.
Number of setup messages received	Number of setup messages received.
Number of teardown messages received	Number of teardown messages received.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
Total packets dropped	Total number of SIP packets dropped.
Unexpected requests dropped	Number of unexpected requests dropped.
Unexpected responses dropped	Number of unexpected responses dropped.
Packets DSCP marked	Number of Differentiated Services code point (DSCP) packets marked.
Packets DSCP marked error	Number of Differentiated Services code point (DSCP) packets marked as error.
NAT errors	Number of Network Address Translation errors.
RR headers exceeded maximum limits	Number of RR headers exceeded maximum limits.
Contact headers exceeded maximum limits	Number of contact headers exceeded maximum limits.
Invite dropped due to call limit	Number of invites dropped due to call limit.
Messages not processed by sip stack	Number of messages not processed by sip stack.
Unknown packets dropped	Number of unknown packets dropped.
Decoding Errors	Number of decoding errors.
Packets received in out of state	Number of packets received in out of state.

Table 64: show services alg statistics Output Fields (*continued*)

Field Name	Field Description
Packets received	Number of packets received.
Packets freed by ALG	Number of packets freed by ALG.
Gate fail errors	Number of gate fail errors.
Lookup packets	Number of lookup packets.
Announce packets	Number of announce packets.
Delete packets	Number of delete packets.
Number of packets received	Number of packets received.
Number of Invalid packets	Number of invalid packets.
Total number of sessions	Total number of sessions.
Number of actives sessions	Number of active sessions.

Sample Output

show services alg statistics application-protocol

While the statistics are the same for dce-rpc and dce-rpc-portmap, both rpc and rpc-portmap have the same output too.

```
user@router> show services alg statistics application-protocol dce-rpc
```

```
Interface name: ms-1/1/0
DCE-RPC ALG statistics:
  Packets with wrong header : 0
  Non epm 3.0 packets       : 0
```



```

Packets with type mismatch: 0
Packets with id mismatch   : 0
Packets with call mismatch: 0
Packets fragmented         : 0
Packets queued             : 0
Packets dropped            : 0
Packets released           : 0

```

user@router> **show services alg statistics application-protocol dns**

```

Interface name: ms-2/0/0
DNS ALG statistics:
  Invalid packets received   : 0
  Reply packets received     : 3509
  Oversized packets received : 0

```

user@router> **show services alg statistics application-protocol ftp**

```

Interface name: ms-1/1/0
FTP ALG statistics:
  Packets dropped           : 0
  ALG parser errors         : 0
  Packets translated        : 0

```

user@router> **show services alg conversations application-protocol h323**

```

Interface name: ms-1/0/0
H323 ALG statistics:
  H323 total calls: 1
  H323 active calls: 1
  H323 gate install failed: 0
  H323 pinhole opened too late: 0
  H323 pinhole hit dropped: 0
  H323 gate timeout failed: 0
  H323 packets dropped: 0
  H323 get virtual ctx failed: 0
  H323 obj alloc failed: 0
  H323 group alloc failed: 0
  H323 ce alloc failed: 0
  H323 Q931 decode error: 0
  H323 H245 decode error: 0
  H323 Q931 process error: 0

```

```

H323 H245 process error: 0
H323 do nat failed: 0
H323 do rm failed: 0
H323 dscp marked: 0
H323 dscp marked error: 0
RAS obj alloc failed: 0
RAS group alloc failed: 0
RAS packets dropped: 0
RAS packet exists in cookie error: 0
RAS decode error: 0
RAS flood error: 0
RAS do nat failed: 0

```

user@router> **show services alg statistics application-protocol ike-esp-nat**

```

Interface name: ms-4/1/0
IKE ESP ALG statistics:
  Session interests processed: 2
  Sessions created: 2
  Sessions destroyed: 1
  Control sessions created: 2
  Control sessions destroyed: 1
  Data sessions created: 0
  Data sessions destroyed: 0
  Gates created: 4
  Gate hits: 0
  Gates timedout: 4

```

user@router> **show services alg statistics application-protocol pptp**

```

Interface name: ms-2/0/0
PPTP ALG statistics:
  PPTP Objects Active   : 1
  PPTP Objects Total    : 1
  PPTP Objects Error    : 0
  PPTP ASL Group Active : 1
  PPTP ASL Group Total  : 1
  PPTP ASL Group Error  : 0
  PPTP Packets received : 11
  PPTP Packets Discarded : 0
  PPTP Packets Free     : 0
  PPTP OCRQ Received    : 1
  PPTP OCRQ Discarded   : 0

```

```

PPTP OCRP Received      : 1
PPTP OCRP Discarded     : 0
PPTP WEN(SLI) Received  : 3
PPTP WEN(SLI) Discarded : 0
PPTP CCRQ-CDSN Received : 0
PPTP CDSN Received      : 0
PPTP CCRQ-CDSN Discarded : 0
PPTP Session Create     : 3
PPTP Session Destroy    : 0
PPTP Gate Create        : 0
PPTP Gate Hit           : 2
PPTP Gate Timeout       : 0
PPTP NAT Events         : 0
PPTP DO-NAT Total       : 1
PPTP DO-NAT Ok          : 1
PPTP DO-NAT Pending     : 0
PPTP DO-NAT Fail        : 0
PPTP DO-RM Total        : 1
PPTP DO-RM Ok           : 2
PPTP DO-RM Pending      : 0
PPTP DO-RM Fail         : 0
PPTP NAT-ASYNC Total    : 0
PPTP NAT-ASYNC Invalid  : 0
PPTP NAT-ASYNC Error1   : 0
PPTP NAT-ASYNC Error2   : 0
PPTP ASL Hole Ok        : 2
PPTP ASL Hole Error     : 0
PPTP ASL First Hit      : 2
PPTP ASL Hole Timeout   : 0
PPTP ASL Invalid        : 0
PPTP NAT Ctx Free       : 0
PPTP Create Resource Error : 0
PPTP set S2C hole error : 0
PPTP set C2S hole error : 0
PPTP lnbrk error        : 0
PPTP Mpool Create Error : 0
PPTP RM register client Error : 0

```

user@router> **show services alg statistics application-protocol rpc**

```

Interface name: ms-1/1/0
RPC ALG statistics:
  Call packet with rpcbind2 : 2
  Call packet with rpcbind3 : 0

```

```

Call packet with rpcbind4 : 0
Invalid rpcbind call      : 0
Reply packet with rpcbind2: 2
Reply packet with rpcbind3: 0
Reply packet with rpcbind4: 0
Invalid rpcbind reply     : 0
Packets fragmented       : 0
Packets dropped           : 0
Packets released         : 0

```

user@router> **show services alg statistics application-protocol rtsp**

```

Interface name: ms-0/1/0
RTSP ALG statistics:
  Packets exceeded maximum length : 0
  Packets dropped by ALG : 0
  Number of describe messages received : 8
  Number of setup messages received : 30
  Number of teardown messages received : 7

```

user@router> **show services alg statistics application-protocol rsh**

```

Interface name: ms-2/0/0
RSH ALG statistics:
  Invalid packets received : 0
  Packets dropped by ALG : 0
  ALG parser errors : 0
  Packets freed by ALG : 0

```

user@router> **show services alg statistics application-protocol sip**

```

Interface name: ms-2/0/0
SIP ALG statistics:
  Total packets dropped : 0
  Unexpected requests dropped : 0
  Unexpected responses dropped : 0
  Packets DSCP marked : 0
  Packets DSCP marked error : 0
  NAT errors : 0
  RR headers exceeded maximum limits : 0
  Contact headers exceeded maximum limits : 0

```

```

Invite dropped due to call limit : 0
Messages not processed by sip stack : 0
Unknown packets dropped : 0
Decoding Errors : 0
Packets received in out of state : 0

```

user@router> **show services alg statistics application-protocol sql**

```

Interface name: ms-2/0/0
SQLNET ALG statistics:
  Packets received : 5
  ALG parser errors : 0
  Packets freed by ALG : 0
  Gate fail errors : 0

```

user@router> **show services alg statistics application-protocol talk**

```

Interface name: ms-2/0/0
TALK ALG statistics:
  Lookup packets : 5
  Announce packets : 0
  Delete packets : 0

```

user@router> **show services alg statistics application-protocol tftp**

```

Interface name: ms-0/0/0
TFTP ALG statistics:
  Number of packets received : 0
  Number of Invalid packets : 0
  Total number of sessions : 0
  Number of active sessions: 0

```

show services alg statistics interface

user@router> **show services alg statistics interface ms-1/1/0**

```

Interface name: ms-1/1/0
FTP ALG statistics:
Packets dropped : 0

```

```
ALG parser errors      : 0
Packets translated    : 0
```

show services cos statistics (Next Gen Services)

Syntax

```
show services cos statistics
<brief | detail | extensive>
<diffserv | forwarding-class>
<interface interface-name>
<service-set service-set-name>
<summary>
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns and the mapping of forwarding class names to queue numbers as configured in CoS services for Next Gen Services services PICs.

Options

none—Display all services CoS statistics.

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

diffserv | forwarding-class—(Optional) Display only the selected information, either DiffServ codepoints or forwarding classes.

interface *interface-name*—(Optional) Display statistics for the specified interface only.

service-set *service-set-name*—(Optional) Display statistics for the specified service set only.

summary—(Optional) Display summary of statistics on a per-interface basis.

Required Privilege Level

view

List of Sample Output

[show services cos statistics on page 830](#)

[show services cos statistics brief on page 832](#)

[show services cos statistics detail on page 832](#)

[show services cos statistics extensive on page 832](#)

Output Fields

Table 65 on page 830 describes the output fields for the **show services cos statistics** command. Output fields are listed in the approximate order in which they appear.

Table 65: show services cos statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Name of interface.	All levels
Service set	Name of service set.	All levels
DSCP	DiffServ code point bit pattern.	All levels
Packets in	Number of packets received.	All levels
Packets out	Number of packets transmitted.	All levels
Forwarding class	Forwarding class queue number.	All levels

Sample Output

show services cos statistics

user@host> **show services cos statistics details**

```

Interface: vms-0/2/0, Service set: ssl
DSCP          Packets in      Packets out
DSCP          Packets in      Packets out
000000          0              0
000001          0              0
000010          0              0
000011          0              0
000100          0              0
000101          0              0
000110          0              0
000111          0              0
001000          0              0
001001          0              0
001010          0              0
001011          0              0
001100          0              0
001101          0              0

```


001110	0	0
001111	0	0
010000	0	0
010001	0	0
010010	0	0
010011	0	0
010100	0	0
010101	0	0
010110	0	0
010111	0	0
011000	0	0
011001	0	0
011010	0	0
011011	0	0
011100	0	0
011101	0	0
011110	0	0
011111	0	0
100000	0	0
100001	0	0
100010	0	0
100011	0	0
100100	0	0
100101	0	0
100110	0	0
100111	0	0
101000	0	0
101001	0	0
101010	0	0
101011	0	0
101100	0	0
101101	0	0
101110	0	0
101111	0	0
110000	0	0
110001	0	0
110010	0	0
110011	0	0
110100	0	0
110101	0	0
110110	0	0
110111	0	0
111000	0	0
111001	0	0

111010	0	0
111011	0	0
111100	0	0
111101	0	0
111110	0	0
111111	0	0
Forwarding class	Packets in	Packets out
0	0	0
1	0	0
2	0	0
3	0	0
4	0	0
5	0	0
6	0	0
7	0	0
8	0	0
9	0	0
10	0	0
11	0	0
12	0	0
13	0	0
14	0	0
15	0	0

show services cos statistics brief

The output for the **show services cos statistics brief** command is identical to that for the **show services cos statistics** command.

show services cos statistics detail

The output for the **show services cos statistics detail** command is identical to that for the **show services cos statistics** command.

show services cos statistics extensive

The output for the **show services cos statistics extensive** command is identical to that for the **show services cos statistics** command.

show services inline software statistics

Syntax

```
show services inline software statistics
<interface interface-name>
<mape name>
<v6rd>
```

Release Information

Command introduced in Junos OS Release 13.3R3.

map-e option introduced in Junos OS Release 18.2R1 for MX Series Routers with MPC and MIC interfaces.

map-e option introduced in Junos OS Release 20.2R1 for Next Gen Services on MX240, MX480 and MX960 routers.

Description

Display information about inline software activity.

Options

- interface interface-name**—(Optional) Display information about the specified services-inline interface only.
When a specific interface is not specified, statistics for all services-inline interfaces are shown.
- mape name**—(Optional) Display information on per physical service interface basis.
- v6rd**—(Optional) Display information for 6rd.

Required Privilege Level

view

List of Sample Output

- [show services inline software statistics on page 835](#)
- [show services inline software statistics mape \(Adaptive Services si- interfaces\) on page 836](#)
- [show services inline software statistics mape \(Next Gen Services si- interfaces\) on page 837](#)

Output Fields

[Table 66 on page 833](#) lists the output fields for the **show services inline software statistics** command. Output fields are listed in the order in which they appear.

Table 66: show services inline software statistics Output Fields

Field Name	Field Description
Service PIC Name	Name of the service PIC for which statistics are displayed.

Table 66: show services inline softwire statistics Output Fields (*continued*)

Field Name	Field Description
Control Plane Statistics	Statistics on the control plane.
ICMPv4 echo requests to softwire concentrator	Number of ICMPv4 echo received by the softwire concentrator. IPv6 ICMP type = 128, code =0. destined to BR IPv6 address
ICMPv4 echo responses from softwire concentrator	Number of ICMPv4 echo responses sent from the softwire concentrator or BR. IPv6 ICMP type = 129
Dropped ICMPv4 packets to softwire concentrator	Number of ICMP packets (except ICMP request) received by the softwire concentrator or BR. All these packets are dropped in by the packet forwarding engine Ukernel.
Trace route UDP packets to softwire concentrator	Number of UDP trace route packets (port numbers 33434 through 33534) received by the softwire concentrator.
ICMPv4 Port unreachable errors sent from softwire concentrator	Number of ICMP port unreachable errors sent by the softwire concentrator after receiving the UDP trace route packets.
Other dropped IPv4 packets to softwire concentrator	Number of non-ICMP packets that were received and dropped because of fragmentation during encapsulation or decapsulation.
Data Plane Statistics	Statistics of the data plane.
6rd decaps	Number of 6rd decapsulated packets and bytes in the data plane. Decapsulation includes removing the outer IPv4 header and routing the inner IPv6 packet.
6rd encaps	Number of 6rd encapsulated (IPv4) packets and bytes in the data plane.

Table 66: show services inline software statistics Output Fields (continued)

Field Name	Field Description
6rd decap errors	Number of all the packets and bytes that are not IPv4-IPv6, IPv4-UDP, or IPV4-ICMP packets.
6rd decap fragment errors	Number of IPv4 fragmented packets and bytes.
6rd decap spoof attacks	Number of spoof attack packets and bytes, which includes packets for which the 6rd derived IPv4 address does not match with the source IPv4 address and packets for which the source IPv6 prefix does not match the 6rd IPv6 prefix.
6rd encap v4 mtu errors	Count of packets and bytes with IPv4 encapsulation MTU errors. For downlink packets after encapsulating with an IPv4 header, if the packet length is more than Tunnel MTU then it is dropped as v4 MTU errors. For these packet drops, an ICMPv6 packet too big error is sent back to the sender.
Data Plane Statistics (MAP-E upstream)	
MAPE decaps	IPv6 packets successfully decapsulated by BR (includes reassembled IPv6)
MAPE ICMP decap errors	IPv6 packets dropped due to unsupported type/code of inner ICMPv4
MAPE decap spoof errors	IPv6 Packets that failed MAPE spoof check

Sample Output

show services inline software statistics

user@host> **show services inline software statistics**

```
Border Router v6rd statistics:
```

```
Service PIC Name
```

```
si-0/0/0
```

Control Plane Statistics

ICMPv4 echo requests to software concentrator	0
ICMPv4 echo responses from software concentrator	0
Dropped ICMPv4 packets to software concentrator	0
Trace route UDP packets to software concentrator	0
ICMPv4 Port unreachable errors sent from software concentrator	0
Other dropped IPv4 packets to software concentrator	0

Data Plane Statistics	Packets	Bytes
6rd decaps	32222173891	3061106519645
6rd encaps	415480622	28252710148
6rd decap errors	0	0
6rd decap fragment errors	0	0
6rd decap spoof attacks	0	0

Service PIC Name si-0/2/0

Control Plane Statistics

ICMPv4 echo requests to software concentrator	0
ICMPv4 echo responses from software concentrator	0
Dropped ICMPv4 packets to software concentrator	0
Trace route UDP packets to software concentrator	0
ICMPv4 Port unreachable errors sent from software concentrator	0
Other dropped IPv4 packets to software concentrator	0

Data Plane Statistics	Packets	Bytes
6rd decaps	0	0
6rd encaps	0	0
6rd decap errors	0	0
6rd decap fragment errors	0	0
6rd decap spoof attacks	0	0
6rd encap v4 mtu errors	0	0

show services inline software statistics mape (Adaptive Services si- interfaces)

user@host> **show services inline software statistics mape**

Service PIC Name		si-0/0/0
Statistics	Packets	Bytes
MAP-E decaps	0	

MAP-E encaps	0	0
MAP-E decap errors	0	0
MAP-E encap errors	0	0
MAP-E decap spoof attacks	0	0
MAP-E decap v4 fragmented	0	0
MAP-E decap v4 reassembled	0	0
MAP-E encap v4 mtu errors	0	0

show services inline software statistics mape (Next Gen Services si- interfaces)

user@host> show services inline software statistics mape

Service PIC Name	si-2/0/0	
Control Plane Statistics		
MAPE ICMPv6 echo requests to softwire concentrator		0
MAPE ICMPv6 echo responses from softwire concentrator		0
MAPE Dropped ICMPv6 packets to softwire concentrator		0
Data Plane Statistics (v6-to-v4)	Packets	Bytes
MAPE decaps	0	0
MAPE ICMP decap errors	0	0
MAPE decap spoof errors	0	0
MAPE v6 reassembled	0	0
MAPE dropped v6 fragments	0	0
MAPE v6 unsupp protocol drops	0	0
Data Plane Statistics (v4-to-v6)	Packets	Bytes
MAPE encaps	0	0
MAPE ICMP encap errors	0	0
MAPE v6 mtu errors	0	0
MAPE v4 reassembled	0	0
MAPE dropped v4 fragments	0	0

show services inline ip-reassembly statistics

Syntax

```
show services inline ip-reassembly statistics
<fpc fpc-slot>
<pfe pfe-slot>
```

Release Information

Statement introduced in Junos OS Release 12.2X49.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display the inline IP reassembly statistics for the Packet Forwarding Engines on one or more MPCs or Next Gen Services MX-SPC3 services card. Inline IP reassembly statistics are collected at the Packet Forwarding Engine level.

NOTE: For more information on MPCs that support inline IP reassembly, refer to *Protocols and Applications Supported on the MPC1E for MX Series Routers*.

Options

none—Displays standard inline IP reassembly statistics for all MPCs or MX-SPC3 services card.

fpc fpc—(Optional) Displays inline IP reassembly statistics for the specified MPC or MX-SPC3 services card.

NOTE: Starting with Junos OS Release 14.2, the FPC option is not displayed for MX Series routers that do not contain switch fabrics, such as MX80 and MX104 routers.

pfe pfe—(Optional) Displays inline IP reassembly for the specified Packet Forwarding Engine slot. You must specify an FPC slot number before specifying a Packet Forwarding Engine slot.

Required Privilege Level

view

RELATED DOCUMENTATION

List of Sample Output

[show services inline ip-reassembly statistics fpc on page 842](#)

Output Fields

[Table 67 on page 839](#) lists the output fields for the **show services inline ip-reassembly statistics** command. Output fields are listed in the approximate order in which they appear.

Table 67: show services inline ip-reassembly statistics Output Fields

Field Name	Field Description
FPC	MPC or MX-SPC3 services card slot number for which the statistics are displayed.
PFE	Packet Forwarding Engine on the MPC or MX-SPC3 services card for which the statistics are displayed.

NOTE: The output fields displayed (per Packet Forwarding Engine) are arranged in a logical sequence from top to bottom to enable users to understand how the inline IP reassembly statistics are gathered.

The information about total number of fragments received is displayed first, and then the information about the reassembled packets and those pending reassembly are displayed. Then, the reasons why the fragments were dropped or not reassembled are displayed. Finally, the information about the fragments reassembled, fragments dropped, and fragments sent to the backup user plane PIC (services PIC) are displayed.

Total Fragments Received	<p>Total number of fragments received and the current rate of fragments received for inline IP reassembly. The following information is also displayed:</p> <ul style="list-style-type: none"> • First Fragments—Number of first fragments received and current rate of first fragments processed. • Intermediate Fragments—Number of intermediate fragments received and current rate of intermediate fragments processed. • Last Fragments—Number and rate of last fragments received. <p>NOTE: Current rate refers to the current number of fragments processed per second in the instant preceding the command's execution.</p>
Total Packets Reassembled	Total number of packets reassembled and current rate, in the instant preceding the command's execution, at which the packets are reassembled.

Table 67: show services inline ip-reassembly statistics Output Fields (*continued*)

Field Name	Field Description
Approximate Packets Pending Reassembly	Approximate number of packets pending reassembly.
Fragments Dropped Reasons	<p>Total number of fragments dropped reasons and the current rate of total fragment dropped reasons. The number of dropped reasons and rate corresponding to each of the following reasons are also displayed:</p> <ul style="list-style-type: none"> • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error • Record in use error • Duplicate first fragments • Duplicate last fragments • Missing first fragment <p>NOTE:</p> <ul style="list-style-type: none"> • These fields indicate why a fragment was dropped. When a fragment is dropped, the corresponding reason field is incremented by 1. For example, when a fragment is dropped because the memory runs out, the Buffers not available field increases by 1. • The maximum number of fragments allowed for reassembly is 16. If the interface encounters a 17th fragment, it drops the entire packet and increments the Fragment per packet exceeded field by 17. • Current rate refers to the current number of fragment dropped reasons per second in the instant preceding the command's execution.
Reassembly Errors Reasons	<p>Number of errors during reassembly and the current rate of reassembly errors. The number of errors and the rate for each of the following types of errors are also displayed:</p> <ul style="list-style-type: none"> • Fragment not found • Fragment not in sequence • ASIC errors <p>NOTE: Current rate refers to the current number of reassembly errors processed per second in the instant preceding the command's execution.</p>

Table 67: show services inline ip-reassembly statistics Output Fields (*continued*)

Field Name	Field Description
Aged out packets	<p>Number of aged out packets and the current number of packets aged out per second in the instant preceding the command's execution.</p> <p>NOTE: In some cases, aged out packets can refer to aged out fragments. If previous fragments of the packet have already been discarded then linking of the dropped fragments to the aged out fragments cannot occur.</p>
Total Fragments Successfully Reassembled	Number of fragments successfully reassembled and the current number of fragments reassembled per second in the instant preceding the command's execution.
Total Fragments Dropped	<p>Total number of fragments dropped and the current rate of total number of fragments dropped. The number of fragments dropped and rate corresponding to each of the following reasons are also displayed:</p> <ul style="list-style-type: none"> • Buffers not available • Fragments per packet exceeded • Packet length exceeded • Record insert error • Record in use error • Duplicate first fragments • Duplicate last fragments • Missing first fragment • Fragment not found • Fragment not in sequence • ASIC errors • Aged out fragments
Total fragments punted to UPIC	Number of fragments sent to the backup user plane PIC (services PIC) and current rate of fragments sent per second in the instant preceding the command's execution

The following information applies to the **Total Fragments Dropped** field.

- These fields indicate *how many* of the packet fragments received were then dropped due to a particular reason.

For example, consider a packet that has 10 fragments, 9 of which have been received and stored in memory. When the tenth fragment arrives, if the memory runs out (Buffers not available), then this fragment is dropped. Because the tenth fragment has been dropped, the other 9 fragments must also be dropped. In this case, the **Buffers not available** field (under the **Fragments Dropped Reasons** field) is incremented by 1 and the **Buffers not available** field (under the **Total Fragments Dropped** field) is incremented by 10.

For the next packet arriving, which also has 10 fragments, the first four fragments are stored but the memory runs out for the fifth fragment. Then the first 5 fragments (fifth and the first four) are dropped. In this case, the **Buffers not available** field (under the **Fragments Dropped Reasons** field) is incremented by 1 and the **Buffers not available** field (under the **Total Fragments Dropped** field) is incremented by 5.

For fragments of the packet, if memory becomes available, the next 5 fragments (6 through 10) that arrive are stored in memory. The fragments are stored until the timeout period elapses, and are eventually dropped. In this case, the **Aged out packets** field is incremented by 1 and the **Aged out fragments** field (under the **Total Fragments Dropped** field) is incremented by 5.

The fragment counters (after both packets have been processed) are as follows:

- **Fragments Dropped Reasons**
 - **Buffers not available** 2
 - **Aged out packets** 1
- **Total Fragment Dropped**
 - **Buffers not available** 15
 - **Aged out packets** 5
- Current rate refers to the current total number fragments dropped per second in the instant preceding the command's execution.

Sample Output

show services inline ip-reassembly statistics fpc

user@host> **show services inline ip-reassembly statistics fpc 0**

```
FPC: 0 PFE: 0
=====
```

	Total	Current Rate
Total Fragments Received	728177644	83529
First Fragments	260759430	29924
Intermediate Fragments	206658784	23681
Last Fragments	260759430	29924

Total Packets Successfully Reassembled	260746982	29924
Approximate Packets Pending Reassembly	4	
Fragments Dropped Reasons	34558	3
Buffers not available	0	0
Fragments per packet exceeded	0	0
Packet length exceeded	0	0
Record insert error	0	0
Record in use error	34558	3
Duplicate first fragments	0	0
Duplicate last fragments	0	0
Missing first fragment	0	0
Reassembly Errors Reasons	0	0
Fragment not found	0	0
Fragment not in sequence	0	0
ASIC errors	0	0
Aged out packets	63	0
Total Fragments Successfully Reassembled	728142977	83528
Total Fragments Dropped	34673	3
Buffers not available	0	0
Fragments per packet exceeded	0	0
Packet length exceeded	0	0
Record insert error	0	0
Record in use error	34558	3
Duplicate first fragments	0	0
Duplicate last fragments	0	0
Missing first fragment	0	0
Fragment not found	0	0
Fragment not in sequence	0	0
ASIC errors	0	0
Aged out fragments	115	0
Total fragments punted to UPIC	0	0

show services nat destination pool

Syntax

```
show services nat destination pool
<interface interface-name>
<service-set service set>
<all>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display destination NAT address pool information.

Options

interface *interface-name*—Optional. Display destination NAT information specific to the interface.

service-set *service-set*—Optional. Display destination NAT information specific to the service set.

all—Optional. Display all destination NAT address pool information.

Required Privilege Level

view

List of Sample Output

[show services nat destination pool on page 845](#)

Output Fields

[Table 68 on page 844](#) lists the output fields for the **show services nat destination pool** command. Output fields are listed in the approximate order in which they appear.

Table 68: show services nat destination pool Output Fields

Field Name	Description
Interface	Interface name.
Service set	Service set name.
Pool name	Pool name.
Pool id	Pool identification.
Total address	Number of IP addresses that are in use.

Table 68: show services nat destination pool Output Fields (*continued*)

Field Name	Description
Translation hits	Number of times a translation in the translation table is used for a source NAT rule.
Address range	IP address range in the source pool.
Port	Port number used to access the pool.

Sample Output

show services nat destination pool

```
user@host> show services nat destination pool service-set ss1_interface_style1 interface vms-0/2/0
all
```

```
ss1_interface_style1 interface vms-0/2/0 all | no-more
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Pool name      : dest_pool
Pool id        : 1
Total address  : 253
Translation hits: 11
Address range              Port
      30.1.1.2 - 30.1.1.254      0
```

show services nat destination rule

Syntax

```
show services nat destination rule
rule-name
<service-set service-set>
<interface interface-name>
<all>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display destination NAT rule-set information.

Options

rule-name—Display information about the specified destination NAT rule.

service-set service-set— Display information specific to the service-set.

interface interface-name — Display information specific to the interface.

all—Display all NAT rule-set information.

Required Privilege Level

view

List of Sample Output

[show services nat destination rule service-set ss1_interface_style1 interface vms-0/2/0 all | no-more on page 847](#)

Output Fields

[Table 69 on page 846](#) lists the output fields for the **show services nat destination rule** command. Output fields are listed in the approximate order in which they appear.

Table 69: show services nat destination rule Output Fields

Field Name	Description
Interface	Interface name.
Service set	Service set name.
Destination NAT rule	Name of the destination NAT rule.

Table 69: show services nat destination rule Output Fields (*continued*)

Field Name	Description
Rule-Id	Rule identification number.
Rule-position	Position of the destination NAT rule.
Match-direction	Three options: <ul style="list-style-type: none"> • input—Apply the rule match on the input side of the interface. • input-output—Apply the rule match bidirectionally. • output—Apply the rule match on the output side of the interface.
Destination addresses	Name of the destination addresses that match the rule. The default value is any.
Action	The action taken when a packet matches the rule's tuples. Actions include the following: <ul style="list-style-type: none"> • destination NAT pool—Use user-defined destination NAT pool to perform destination NAT. • off—Do not perform destination NAT.
Translation hits	Number of times a translation in the translation table is used for a source NAT rule.
Successful sessions	Number of successful session installations after the NAT rule is matched.
Failed sessions	Number of unsuccessful session installations after the NAT rule is matched.
Number of sessions	Number of sessions that reference the specified rule.

Sample Output

```
show services nat destination rule service-set ss1_interface_style1 interface vms-0/2/0 all | no-more
user@host> show services nat destination rule service-set ss1_interface_style1 interface vms-0/2/0
all | no-more
```

```
ss1_interface_style1 interface vms-0/2/0 all | no-more

Interface: vms-0/2/0 , Service set: ss1_interface_style1
Destination NAT rule: r1                               Rule-set: rs2
Rule-Id                               : 2
```

```
Rule position      : 1
Match-direction    : input
  Destination addresses : 50.1.1.2      - 50.1.1.2
Action             : dest_pool
Translation hits    : 34
  Successful sessions  : 34
  Failed sessions     : 0
Number of sessions  : 0
```

show services nat destination summary

Syntax

```
show services nat destination summary
<interface interface-name>
<service-set service-set>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display summary destination NAT information.

Options

interface interface-name—Display summary destination NAT information for the specified service interface.

service-set service-set—Display summary destination NAT information for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services nat destination summary service-set ss1_interface_style1 interface vms-0/2/0 on page 850](#)

Output Fields

[Table 70 on page 849](#) lists the output fields for the **show services nat destination summary** command. Output fields are listed in the approximate order in which they appear.

Table 70: show services nat destination summary Output Fields

Field Name	Description
Interface	Interface name.
Service set	Service set name.
Pool name	Name of the destination address pool.
Address Range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
Port	Port number.

Table 70: show services nat destination summary Output Fields (*continued*)

Field Name	Description
Total Address	Number of IP addresses that are in use.
Rule name	Rule name.
Rule set	The set of rules for destination NAT.
Match-direction	Three options: <ul style="list-style-type: none"> • input—Apply the rule match on the input side of the interface. • input-output—Apply the rule match bidirectionally. • output—Apply the rule match on the output side of the interface.
Action	The action taken when a packet matches the rule's tuples. Actions include the following: <ul style="list-style-type: none"> • destination NAT pool—Use user-defined destination NAT pool to perform destination NAT. • off—Do not perform destination NAT.

Sample Output

show services nat destination summary service-set ss1_interface_style1 interface vms-0/2/0

user@host> **show services nat destination summary service-set ss1_interface_style1 interface vms-0/2/0**

```

Interface: vms-0/2/0 , Service set: ss1_interface_style1
Pool name      Address                               Routing      Port  Total
                                     Range          Instance      Address
dest_pool      30.1.1.2      - 30.1.1.254                                0      253
Interface: vms-0/2/0 , Service set: ss1_interface_style1
Rule name      Rule set      Match-direction  Action
r1             rs2           input            dest_pool

```

show services nat ipv6-multicast-interfaces

Syntax

```
show services nat ipv6-multicast-interfaces
```

Release Information

Command introduced in Junos OS Release 8.5.

Description

Displays a list of interfaces enabled for IPv6 multicast.

Required Privilege Level

view

List of Sample Output

[show services nat ipv6-multicast-interfaces on page 851](#)

Output Fields

[Table 71 on page 851](#) lists the output fields for the **show services nat ipv6-multicast-interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 71: show services nat ipv6-multicast-interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a service interface.	All levels
Admin State	Configured IPv6 multicast capability of an interface ,	All levels
Operational State	Operation IPv6 multicast status of an interface.	All levels

Sample Output

show services nat ipv6-multicast-interfaces

user@host> **show services nat ipv6-multicast-interfaces**

Interface	Admin State	Operational State
ge-5/1/9	Enabled	Enabled
ge-5/1/8	Enabled	Enabled

ge-5/1/7	Enabled	Enabled
ge-5/1/6	Enabled	Enabled
ge-5/1/5	Enabled	Enabled
ge-5/1/4	Enabled	Enabled
ge-5/1/3	Enabled	Enabled
ge-5/1/2	Enabled	Enabled
ge-5/1/1	Enabled	Enabled
ge-5/1/0	Enabled	Enabled
ge-5/0/9	Enabled	Enabled
ge-5/0/8	Enabled	Enabled
ge-5/0/7	Enabled	Enabled
ge-5/0/6	Enabled	Enabled
ge-5/0/5	Enabled	Enabled
ge-5/0/4	Enabled	Enabled
ge-5/0/3	Enabled	Enabled
ge-5/0/2	Enabled	Enabled
ge-5/0/1	Enabled	Enabled
ge-5/0/0	Enabled	Enabled
ge-1/3/9	Enabled	Enabled
ge-1/3/8	Enabled	Enabled
ge-1/3/7	Enabled	Enabled
ge-1/3/6	Enabled	Enabled
ge-1/3/5	Enabled	Enabled
ge-1/3/4	Enabled	Enabled
ge-1/3/3	Enabled	Enabled
ge-1/3/2	Enabled	Enabled
ge-1/3/1	Enabled	Enabled
ge-1/3/0	Enabled	Enabled
ge-1/2/9	Enabled	Enabled
ge-1/2/8	Enabled	Enabled
ge-1/2/7	Enabled	Enabled
ge-1/2/6	Enabled	Enabled
ge-1/2/5	Enabled	Enabled
ge-1/2/4	Enabled	Enabled
ge-1/2/3	Enabled	Enabled
ge-1/2/2	Enabled	Enabled
ge-1/2/1	Enabled	Enabled
ge-1/2/0	Enabled	Enabled
ge-1/1/9	Enabled	Enabled
ge-1/1/8	Enabled	Enabled
ge-1/1/7	Enabled	Enabled
ge-1/1/6	Enabled	Enabled
ge-1/1/5	Enabled	Enabled
ge-1/1/4	Enabled	Enabled

ge-1/1/3	Enabled	Enabled
ge-1/1/2	Enabled	Enabled
ge-1/1/1	Enabled	Enabled
ge-1/1/0	Enabled	Enabled
ge-1/0/9	Enabled	Enabled
ge-1/0/8	Enabled	Enabled
ge-1/0/7	Enabled	Enabled
ge-1/0/6	Enabled	Enabled
ge-1/0/5	Enabled	Enabled
ge-1/0/4	Enabled	Enabled
ge-1/0/3	Enabled	Enabled
ge-1/0/2	Enabled	Enabled
ge-1/0/1	Enabled	Enabled
ge-1/0/0	Enabled	Enabled
xe-0/3/0	Enabled	Enabled
xe-0/2/0	Enabled	Enabled
xe-0/1/0	Enabled	Enabled
xe-0/0/0	Enabled	Enabled

show services nat resource-usage source-pool

Syntax

```
show services nat resource-usage source-pool
<all>
pool-name
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display NAT resource usage.

Options

<all>—Display all NAT resource usage statistics.

pool-name—Display NAT resource usage statistics for the specified pool.

Required Privilege Level

view

List of Sample Output

[show services nat resource-usage source-pool all on page 855](#)

[show services nat resource-usage source-pool src-nat-pool-2 on page 855](#)

Output Fields

[Table 72 on page 854](#) lists the output fields for the **show services nat resource-usage** command. Output fields are listed in the approximate order in which they appear.

Table 72: show services nat resource-usage Output Fields

Field Name	Description
Pool	Name of the pool.
Address	Address of the pool.
Used	Number of used resources in the pool.
Available	Number of available resources in the pool.
Total	Total number of addresses in the pool.
Usage	Percent of resources used.


```
user@host> show services nat resource-usage source-pool all
```

```
show services nat resource-usage source-pool src-nat-pool-2
```

```
show services nat resource-usage source-pool src-nat-pool-2
```

```

Pool name: src-nat-pool-2
Total address: 4
Port-overloading-factor: 1
Total ports: 258048 Used: 0 Avail: 258048
Current usage: 0% Peak usage: 0% at 1970-01-01 00:00:00 UTC

```

Address	Factor-index	Port-range	Used	Avail	Total	Usage
1.1.1.20	0	Single Ports	0	64512	64512	0%
1.1.1.21	0	Single Ports	0	64512	64512	0%
1.1.1.22	0	Single Ports	0	64512	64512	0%
1.1.1.23	0	Single Ports	0	64512	64512	0%

show services nat source deterministic

Syntax

```
show services source nat deterministic
host-address-range
host-ip ip-address
pool pool-name
xlated-ip translated-ip-address
xlated-port translated-port-number
```

Release Information

This command was introduced in Junos OS 19.3R2.

Description

Display deterministic port block allocation information.

Options

host-address-range—Display the deterministic host address range without overlap.

host-ip ip-address—Display the internal host IP address.

pool pool-name—Display the source NAT pool.

xlated-ip translated-ip-address—Display translated IP address.

xlated-port translated-port-number—Display the translated port number.

Required Privilege Level

view

List of Sample Output

[show services nat source deterministic on page 857](#)

[show services nat source deterministic host-address-range on page 858](#)

Output Fields

[Table 73 on page 856](#) lists the output fields for the `show services nat source deterministic` command. Output fields are listed in the approximate order in which they appear.

Table 73: show services nat source deterministic Output Fields

Field Name	Field Description
Pool name	Name of the NAT source pool.

Table 73: show services nat source deterministic Output Fields (*continued*)

Field Name	Field Description
Port overloading factor	Factor of port overloading for the source pool.
Used/total port blocks	Port block used number and port block total number for this source NAT pool.
Host IP	Host IP address.
External IP	IP address of external router.
Port Block Range	The range of ports in a block, ranging from lowest to highest.
Ports Used/Ports Total	Number of ports used and total ports.
Total host ranges number	Host ranges in total.
Min Host Address	Minimum host address.
Max Host Address	Maximum host address.

Sample Output

show services nat source deterministic

user@host> **show services nat source deterministic**

```
Pool name: src-nat-pool-1
Port-overloading-factor: 1 Port block size: 256
Used/total port blocks: 0/12
Host_IP External_IP Port_Block Ports_Used/
                        Range           Ports_Total
10.1.1.1 202.0.0.1    1280-1535      0/256*1
10.1.1.2 202.0.0.1    1536-1791      0/256*1
```

show services nat source deterministic host-address-range

user@host> **show services nat source deterministic host-address-range**

```
Pool name: src-nat-pool-1
Total host ranges number: 1
Min Host Address Max Host Address
10.1.1.1 10.1.1.2
```

show services nat source mappings address-pooling-paired

Syntax

```
show services nat source mappings address-pooling-paired
```

Description

Options

address-pooling-paired—(Optional) Display only information about address-pooling paired mappings.

endpoint-independent—(Optional) Display only information about endpoint-independent mappings.

pcp—(Optional) Display only information about port control protocol mappings.

NOTE: PCP requests with the prefer-failure option request a particular external IP address and port. When the request cannot be fulfilled, the mapping is not created. In this case, the subscriber does not have a mapped IP address. Such a subscriber is counted in the summary of the number or address mappings, but is not displayed in the list of address mappings, as shown in the following examples:

```
user@host# show services nat mappings summary
```

```
Service Interface:                sp-2/0/0
Total number of address mappings: 1
Total number of endpoint independent port mappings: 0
Total number of endpoint independent filters: 0
```

```
user@host# show services nat mappings address-pooling-paired
```

```
[edit]
```

This is expected behavior because unfulfilled address mappings (IP of 0.0.0.0) are not displayed in the output of the second CLI command. These address mappings will time out based on configured or default values.

Required Privilege Level

view

List of Sample Output

[show services nat source mappings address-pooling-paired on page 860](#)

[show services nat source mappings address-pooling-paired private 1.1.1.100 on page 860](#)
[show services nat source mappings address-pooling-paired public 30.30.30.2 on page 860](#)
[show services nat source mappings address-pooling-paired pool-name sp1 on page 861](#)
[show services nat mappings address-pooling-paired on page 861](#)
[show services nat mappings address-pooling-paired \(mapping of active B4 for a subscriber\) on page 861](#)
[show services nat mappings endpoint-independent on page 861](#)
[show services nat mappings pcg on page 862](#)
[show services nat mappings nptv6 internal on page 862](#)
[show services nat mappings nptv6 external on page 862](#)

Sample Output

show services nat source mappings address-pooling-paired

user@host> **show services nat source mappings address-pooling-paired**

```

Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address      External address      Session Count      Mapping State
1.1.1.100             30.30.30.1           1                  Active
1.1.1.101             30.30.30.2           1                  Active

```

show services nat source mappings address-pooling-paired private 1.1.1.100

user@host> **show services nat source mappings address-pooling-paired private 1.1.1.100**

```

Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address      External address      Session Count      Mapping State
1.1.1.100             30.30.30.1           1                  Active

```

show services nat source mappings address-pooling-paired public 30.30.30.2

user@host> **show services nat source mappings address-pooling-paired public 30.30.30.2**

```

Interface: ms-2/0/0, Service set: ss1
Pool name: sp1
Internal address      External address      Session Count      Mapping State
1.1.1.101             30.30.30.2           1                  Active

```

show services nat source mappings address-pooling-paired pool-name sp1

```
user@host> show services nat source mappings address-pooling-paired pool-name sp1
```

```
Interface: ms-2/0/0, Service set: ssl
Pool name: sp1
Internal address      External address      Session Count      Mapping State
1.1.1.100             30.30.30.1           1                  Active
1.1.1.101             30.30.30.2           1                  Active
```

show services nat mappings address-pooling-paired

```
user@host> show services nat mappings address-pooling-paired
```

```
Interface: sp-3/0/0, Service set: NAPT44-SS1
NAT pool: napt44-SS1-pl
Mapping          : 29.32.38.255    --> 192.168.75.23
Ports In Use     :          9
Session Count    :          1
Mapping State     : Active
```

show services nat mappings address-pooling-paired (mapping of active B4 for a subscriber)

```
user@host> show services nat mappings address-pooling-paired
```

```
Interface: sp-0/0/0, Service set: sset_1

NAT pool: nat_pool1

Mapping          : 2001::          --> 33.33.33.2
Ports In Use     :          1
Session Count    :          9
Mapping State     : Timeout
```

show services nat mappings endpoint-independent

```
user@host> show services nat mappings endpoint-independent
```

```
Interface: sp-3/0/0, Service set: NAPT44-SS1
NAT pool: napt44-SS1-pl
Mapping          : 29.32.38.255:10000    --> 192.168.75.23:1024
Session Count    : 1
Mapping State     : Active
```

show services nat mappings pcip

```
user@host> show services nat mappings pcip
```

```
PCP Client      : 172.16.0.1      PCP Lifetime : 45
Mapping         : 29.32.38.255:10000 --> 192.168.75.23:1024
Session Count   : 1
Mapping State   : Active
```

show services nat mappings nptv6 internal

```
user@host> show services nat mappings nptv6 internal 1111:2222:3333:aaaa:bbbb::1
```

```
Interface      Service-set  NAT-Pool      Address Mapping
vms-0/1/0      ss_nptv6     ss_nptv6_pool 1111:2222:3333:aaaa:bbbb::1 ->
aaaa:bbbb:cccc:dddd:bbbb::1
```

show services nat mappings nptv6 external

```
user@host> show services nat mappings nptv6 external aaaa:bbbb:cccc:dddd:bbbb::1
```

```
Interface      Service-set  NAT-Pool      Address Mapping
vms-0/1/0      ss_nptv6     ss_nptv6_pool 1111:2222:3333:aaaa:bbbb::1
-> aaaa:bbbb:cccc:dddd:bbbb::1
```


show services nat source mappings endpoint-independent

Syntax

```
show services nat source mappings endpoint-independent
<pool-name>
<private | public>
```

Release Information

Command introduced in Junos OS 19.3R2.
Support for Next Gen Services with the MX-SPC3 security services card added in Junos OS Release 20.2.

Description

Displays NAT endpoint independent mapping.

Options

- <pool-name>—Name of address pool.
- <private>—Private IPv4/IPv6 prefix to use as a filter.
- <public>—Public IP prefix to use as a filter.

Required Privilege Level

view

List of Sample Output

- [show services nat source mappings endpoint-independent \(ms- interfaces\) on page 864](#)
- [show services nat source mappings endpoint-independent private 15.4.4.2 public 20.20.20.1 \(ms-interfaces\) on page 864](#)
- [show services nat source mappings endpoint-independent pool-name p1 \(ms-interfaces\) on page 864](#)
- [show services nat source mappings address-pooling-paired pool-name sp1 \(sp- interfaces\) on page 865](#)
- [show services nat source mappings endpoint-independent \(vms- interfaces\) on page 865](#)

Output Fields

[Table 74 on page 863](#) lists the output fields from the **show services nat source mappings endpoint-independent** command. Output fields are listed in the approximate order in which they appear.

Table 74: show services nat source mappings endpoint-independent Output Fields

Field Name	Description
Interface	Name of the interface.
Service set	Name of the service set.

Table 74: show services nat source mappings endpoint-independent Output Fields (*continued*)

Field Name	Description
NAT pool	Name of the NAT pool.
Mapping	Shows the mapping of IP addresses.
Session Count	Number of sessions currently using the mapping.
Mapping State	NAT mapping state. The following states are possible: <ul style="list-style-type: none"> • ACTIVE—Indicates that the entry is active and in use. • TIMEOUT—Indicates that the mapping is not in use. After the mapping-timeout, configured at the [edit services nat pool pool-name] hierarchy level, lapses, the mapping is deleted. This field also displays the number of seconds after which the timeout occurs.

Sample Output

show services nat source mappings endpoint-independent (ms- interfaces)

user@host> **show services nat source mappings endpoint independent**

```
Interface: ms-2/0/0, Service set: ss1
NAT pool: test-pool
Mapping      : 2.1.1.1          : 1026  --> 123.0.0.5          :10926
Session Count :      1
Mapping State  : Active
```

show services nat source mappings endpoint-independent private 15.4.4.2 public 20.20.20.1 (ms-interfaces)

user@host> **show services nat source mappings endpoint-independent private 15.4.4.2 public 20.20.20.1**

```
Interface: ms-2/0/0, Service set: ss1
NAT pool: p1
Mapping      : 15.4.4.2          :12841  --> 20.20.20.1          :11205
Session Count :      1
Mapping State  : Active
```

show services nat source mappings endpoint-independent pool-name p1 (ms-interfaces)

user@host> **show services nat source mappings endpoint-independent pool-name p1**

```

Interface: ms-2/0/0, Service set: ss1
NAT pool: pl
Mapping      : 15.4.4.2          :12841  --> 20.20.20.1      :11205
Session Count :      1
Mapping State  : Active

```

show services nat source mappings address-pooling-paired pool-name sp1 (sp- interfaces)

```
user@host> show services nat source mappings address-pooling-paired pool-name sp1
```

```

Interface: ms-2/0/0, Service set: ss1
Pool name: sp1

```

Internal address	External address	Session Count	Mapping State
1.1.1.100	30.30.30.1	1	Active
1.1.1.101	30.30.30.2	1	Active

Sample Output

show services nat source mappings endpoint-independent (vms- interfaces)

```
user@host> show services nat source mappings endpoint-independent
```

```

Interface: vms-2/0/0, Service set: vms-sset10
Pool name: napt44-pool12

```

Mapping	: 20.1.0.101	: 1024	--> 50.0.12.1	: 1024
Session Count	: 1			
Mapping State	: Active			
B4 Address	: 2002:2010::1401:4	>>>>>>> B4 Address in mapping		

show services nat source mappings pcp

Syntax

```
show services nat source mappings pcp
<interface interface-name>
<service-set service-set.>
```

Release Information

Command introduced in Junos OS 20.1R1.

Description

Display NAT source mapping for PCP.

Options

interface *interface-name*—Display PCP source NAT mapping for the specified interface.

service-set *service-set*—Display PCP source NAT mapping for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services nat source mappings pcp on page 866](#)

Sample Output

show services nat source mappings pcp

user@host> show services nat source mappings pcp Interface: vms-0/0/0, Service set: in

```
NAT pool: p
PCP Client      : 10.1.1.2          PCP lifetime : 995
Mapping         : 10.1.1.2          : 9000 --> 8.8.8.8          : 1025
Session Count   : 1
Mapping State    : Active

DS-LITE output:
=====
PCP Client      : 2222::1          PCP lifetime : 106
Mapping         : 88.1.0.47        : 47 --> 70.70.70.1        : 41972
Session Count   : 1
```

```
Mapping State      : Active
B4 Address         : 2222::1
```

show services nat source mappings summary

Syntax

```
show services nat source mappings summary
<interface interface-name>
<service-set service-set.>
```

Release Information

Command introduced in Junos OS 19.3R2.

Description

Display NAT mapping summary information.

Options

interface *interface-name*—Display source NAT mapping information for the specified interface.

service-set *service-set*—Display source NAT mapping information for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services nat source mappings summary on page 869](#)

Output Fields

[Table 75 on page 868](#) lists the output fields for the **show services nat source mappings summary** command. Output fields are listed in the approximate order in which they appear.

Table 75: show services nat source mappings summary Output Fields

Field Name	Field Description
Service Interface	Name of the service interface.
Total number of address mappings	Displays total number of address mappings.
Total number of endpoint independent port mappings	Displays total number of endpoint independent port mappings.
Total number of endpoint independent filters	Displays total number of endpoint independent filters.

Sample Output

show services nat source mappings summary

user@host> **show services nat source mappings summary**

```
Service Interface:                ms-2/0/0
Total number of address mappings: 2
Total number of endpoint independent port mappings: 1
Total number of endpoint independent filters: 1
```

show services nat source pool

Syntax

```
show services nat source pool pool-name
<all>
<interface interface-name>
<service-set service-set>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display source NAT information for a pool.

Options

pool-name—Display information about the specified pool.

all—Display all source NAT pool information.

interface interface-name—Display information specific to the adaptive services interface.

service-set service-set—Display information specific to the service set.

Required Privilege Level

view

List of Sample Output

- [show services nat source pool JNPR-CGNAT-PUB-POOL \(NAT Pool\) on page 872](#)
- [show services nat source pool JNPR-CGNAT-PUB-POOL \(PBA Pool\) on page 873](#)
- [show services nat source pool JNPR-CGNAT-PUB-POOL \(Deterministic\) on page 874](#)
- [show services nat source pool service-set ss1_interface_style1 interface vms-0/2/0 all on page 874](#)

Output Fields

[Table 76 on page 870](#) lists the output fields for the **show services nat source pool** command. Output fields are listed in the approximate order in which they appear.

Table 76: show services nat source pool Output Fields

Field Name	Description
Pool name	Name of the source pool.
Pool id	Pool identification number.
Routing instance	Name of the routing instance.

Table 76: show services nat source pool Output Fields (*continued*)

Field Name	Description
Host address base	Base address of the original source IP address range.
Port	Port numbers used for the source pool.
Port overloading	Number of port overloading for the source pool.
Address assignment	Type of address assignment.
Total addresses	Number of IP addresses that are in use.
Translation hits	Number of times there is traffic that matches the source rule.
Limit ports per host	
Include-boundary-addresses	Include the lowest and highest addresses in the source address range of the NAT rule to be translated when the NAT pool is used.
Ei-mapping-timeout	Duration for endpoint independent translations that use the specified NAT pool.
Mapping-timeout	Duration for mappings that use the specified NAT pool.
EIF Inbound session count	Number of EIF inbound sessions.
EIF Inbound session limit exceeded drops	Number of EIF inbound sessions that exceed the drop limit.
Address range	IP address range for the source pool.
Ports	
Total used ports	

Table 76: show services nat source pool Output Fields (continued)

Field Name	Description
Error Counters	The following bullets describe the fields:
• Out of port errors	• No ports available.
• Out of address errors	• No room in the pool for another address.
• Parity port errors	•
• Preserve Range errors	•
• APP port allocation errors	•
• App port limit allocation errors	•
• Port block allocation errors	•
• Port blocks limit exceeded errors	•

Sample Output

show services nat source pool JNPR-CGNAT-PUB-POOL (NAT Pool)

user@host> show services nat source pool JNPR-CGNAT-PUB-POOL

```

Interface: vms-0/2/0 , Service set: JNPR-IF-SSET
Pool name      : JNPR-CGNAT-PUB-POOL
Pool id       : 4
Routing instance : default
Host address base : 0.0.0.0
Port          : [1024, 65535]
Port overloading : 1
Address assignment : no-paired
Total addresses  : 254
Translation hits : 0
+Limit ports per host : 10
Include-boundary-addresses: Disable
Ei-mapping-timeout : 300
Mapping-timeout    : 300
EIF Inbound session count: 0
EIF Inbound session limit exceeded drops: 0
Address range      Ports
20.20.20.1 - 20.20.20.254 0
Total used ports    : 0
+Error Counters:
+ Out of port errors : 0

```

```

+   Out of address errors           : 0
+   Parity port errors              : 0
+   Preserve Range errors           : 0
+   APP port allocation errors       : 0
+   APP port limit allocation errors : 0
+   Port block allocation errors     : 0
+   Port blocks limit exceeded errors: 0

```

show services nat source pool JNPR-CGNAT-PUB-POOL (PBA Pool)

user@host> show services nat source pool JNPR-CGNAT-PUB-POOL

```

Interface: vms-0/2/0 , Service set: JNPR-IF-SSET
Pool name      : JNPR-CGNAT-PUB-POOL
Pool id       : 4
Routing instance : default
Port          : [1024, 65535]
Port overloading : 1
Address assignment : no-paired
Total addresses  : 510
Translation hits : 0
Port block size  : 256
Max blocks per host : 8
Active block timeout : 0
Interim logging interval : 0
PBA block log    : Enable
Used/total port blocks: 0/128520
+Max number of port blocks used: 0
Include-boundary-addresses: Disable
Ei-mapping-timeout : 300
Mapping-timeout    : 300
EIF Inbound session count: 0
EIF Inbound session limit exceeded drops: 0
Address range      Ports
      100.0.0.1 - 100.0.1.254      0
Total used ports   :                0
Error Counters:
    Out of port errors           : 0
    Out of address errors        : 0
    Parity port errors           : 0
    Preserve Range errors        : 0
    APP port allocation errors    : 0
    APP port limit allocation errors : 0

```

```

Port block allocation errors      : 0
Port blocks limit exceeded errors : 0

```

show services nat source pool JNPR-CGNAT-PUB-POOL (Deterministic)

user@host> **show services nat source pool JNPR-CGNAT-PUB-POOL**

```

Interface: vms-0/2/0 , Service set: JNPR-IF-SSET
Pool name      : JNPR-CGNAT-PUB-POOL
Pool id       : 4
Routing instance : default
Port          : [1024, 65535]
Port overloading : 1
Address assignment : no-paired
Total addresses  : 510
Translation hits  : 0
Port block size   : 256
Determ host range num: 1
+Unique pool users: 0
Include-boundary-addresses: Disable
Ei-mapping-timeout : 300
Mapping-timeout    : 300
EIF Inbound session count: 0
EIF Inbound session limit exceeded drops: 0

```

Address range	Single Ports	Twin Ports
100.0.0.1 - 100.0.1.254	0	0
Total used ports :	0	0

```

Error Counters:
  Out of port errors      : 0
  Out of address errors   : 0
  Parity port errors      : 0
  Preserve Range errors   : 0
  APP port allocation errors : 0
  APP port limit allocation errors : 0
  Port block allocation errors : 0
  Port blocks limit exceeded errors : 0

```

show services nat source pool service-set ss1_interface_style1 interface vms-0/2/0 all

user@router> **show services nat source pool service-set ss1_interface_style1 interface vms-0/2/0 all**

```

Interface: vms-0/2/0 , Service set: ss1_interface_style1
Pool name      : src_pool1

```

Pool id	: 4		
Routing instance	: default		
Host address base	: 0.0.0.0		
Port	: [1024, 63487]		
Twin port	: [63488, 65535]		
Port overloading	: 1		
Address assignment	: no-paired		
Total addresses	: 254		
Translation hits	: 3		
Address range		Single Ports	Twin Ports
44.0.0.1 - 44.0.0.254		1	0
Total used ports	:	1	0

show services nat source port-block

Syntax

```
show services nat source port-block
<host-ip ip-address>
<pool pool-name>
<xlated-ip translated-ip-address>
<xlated-port translated-port-number>
```

Release Information

Command introduced in Junos OS 19.3R2.

Description

Display port block allocation information.

Options

host-ip ip-address—Display port block allocation information for the specified host.

pool pool-name—Display port block allocation information for the specified pool.

xlated-ip translated-ip-address—Display port block allocation information for the specified translated IP address.

xlated-port translated-port-number —Display port block allocation information for the specified translated port number.

Required Privilege Level

view

List of Sample Output

[show services nat source port-block on page 878](#)

Output Fields

[Table 77 on page 876](#) lists the output fields for the **show services nat source port block** command. Output fields are listed in the approximate order in which they appear.

Table 77: show services nat source port block Output Fields

Field Name	Field Description
Pool name	Name of the pool.
Port-overloading-factor	Factor of port overloading for the source pool.
Port block size	Number of ports that a port block contains.

Table 77: show services nat source port block Output Fields (*continued*)

Field Name	Field Description
Max port blocks per host	Maximum number of blocks that one host can use for translation.
Port block active timeout	Longest duration that a block remains active for port allocation.
Used/total port blocks	Current number of used ports and total number of ports in this source pool.
Host IP	Host IP address.
External IP	External IP address.
Port Block Range	Port range of one PBA port block entry from the lowest to the highest port number that can be allowed to allocate ports for this block.
Ports Used/Ports Total	Current number of used ports and total number of ports in this source pool.
Block State/Left Time (s)	<p>PBA port block entry state for NAT port allocation, including Active, Inactive, Query, and the time left for a port block that is in the Active or Query state.</p> <ul style="list-style-type: none"> • Active—When an internal subscriber initiates a NAT request, a port block is allocated from the pool, and the status is set to Active. When there is a subsequent request from the same subscriber, a port is allocated from the existing Active block. • Inactive—When there is a request from an internal subscriber who had previously had a port allocated from this port block, but the time on the Active port block has expired or the ports are used up, the port block status changes from Active to Inactive. • InactiveB—When a chassis cluster is in active/passive mode, and a port block is created on the active node, the status for the synced port block on the backup node is InactiveB. • Query—When no ports are used in an Active port block, the status changes from Active to Query.
Failed sessions	Number of failed sessions.
Number of sessions	Total number of sessions.

Sample Output

show services nat source port-block

user@host> **show services nat source port-block**

```
Pool name: spl
Port-overloading-factor:      1      Port block size:  512
Max port blocks per host:    8      Port block active timeout:  100
Used/total port blocks: 1/64260

Host_IP      External_IP      Port_Block      Ports_Used/
      Block_State/
      Left_Time(s)      Range      Ports_Total
1.1.1.100    30.30.30.1      13824-14335      1/512*1
      Active/71

Failed sessions      : 0
Number of sessions   : 0
```


show services nat source rule

Syntax

```
show services nat source rule
rule-name
<all>
<interface interface-name>
<service-set service-set>
```

Release Information

Command introduced in Junos OS 19.3R2.

Description

Display source NAT rule-set information.

Options

rule-name—Display source NAT rule-set information for the specified rule.

all—Display all source NAT rule-set information.

interface interface-name—Display rule-set information about the adaptive services interface.

service-set service-set—Display rule-set information about the service set.

Required Privilege Level

view

List of Sample Output

[show services nat source rule on page 880](#)

[show services nat source rule \(Mapping and EIF Configuration\) on page 881](#)

Output Fields

[Table 78 on page 879](#) lists the output fields for the **show services nat source rule** command. Output fields are described in the approximate order in which they appear.

Table 78: show services nat source rule Output Fields

Field Name	Description
Interface	Interface name.
Service set	Service set name.
Rule Id	Rule identification number.


```

Application          : configured
Action               : src_pool1
Persistent NAT type  : N/A
Persistent NAT mapping type : address-port-mapping
Inactivity timeout   : 0
Max session number   : 0
Translation hits     : 3
Successful sessions  : 3
Failed sessions      : 0
Number of sessions   : 1

```

show services nat source rule (Mapping and EIF Configuration)

show services nat source rule all

```

Total rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/0
source NAT rule: r1                      Rule-set: rs1
Rule-Id                               : 1
Rule position                          : 1
From zone                             : nh-JNPR-NH-SSET-ZoneIn
To zone                               : nh-JNPR-NH-SSET-ZoneOut
Match
Source addresses                      : 30.30.30.0      - 30.30.30.255
Action                               : p2
+Mapping-type : endpoint-independent;
+Mapping-refresh : inbound
+Filtering-type: endpoint-independent
+Prefix-list :
1.2.2.0 --- 2.2.2.3
3.3.3.0 --- 3.3.3.3 except
Translation hits                      : 0
Successful sessions                   : 0
Failed sessions                       : 0
Number of sessions                    : 0

```

show services nat source rule-application

Syntax

```
show services nat source rule-application
<all>
<interface interface-name>
<service-set service-set>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display source NAT rule application information.

Options

all—Display all source NAT rule application information.

interface *interface-name*—Display source NAT rule application information for the specified interface.

service-set *service-set*—Display source NAT rule application information for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services nat source rule-application on page 883](#)

Output Fields

[Table 79 on page 882](#) lists the output fields for the **show services nat source rule-application** command. Output fields are described in the approximate order in which they appear.

Table 79: show services nat source rule-application Output Fields

Field Name	Description
Interface	Displays rule application for the specified interface.
Service set	Displays rule application for the specified service set.

Table 79: show services nat source rule-application Output Fields (*continued*)

Field Name	Description
Source NAT rule	The name of the source NAT rule.
<ul style="list-style-type: none"> • Rule-set • Rule-Id • Match-direction • Application • IP Protocol • Source port range • Destination port range 	<ul style="list-style-type: none"> • Set of rules for matching traffic. • Rule identification number. • Specifies the direction in which to match traffic that meets the rule conditions. • Name of the application or application set. • IP protocol identifier. • Source port range identifier. • Destination port range identifier.

Sample Output

show services nat source rule-application

```
user@host> show services nat source rule-application service-set ss1_interface_style1 interface
vms-0/2/0 all
```

```
Interface: vms-0/2/0 , Service set: ss1_interface_style1
source NAT rule: r1           Rule-set: rs1
  Rule-Id                  : 1
  Match-direction          : input
  Application: any
  IP protocol: 0
  Source port range: [0-0]
  Destination port range: [0-0]
```

show services nat source summary

Syntax

```
show services nat source summary
<interface interface-name>
<service-set service-set>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Displays source NAT summary information.

Options

interface *interface-name*—Display source NAT summary information for the specified interface.

service-set *service-set*—Display source NAT summary information for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services nat source summary on page 885](#)

Output Fields

[Table 80 on page 884](#) lists the output fields for the **show services nat source summary** command. Output fields are listed in the approximate order in which they appear.

Table 80: show services nat source summary Output Fields

Field Name	Description
Interface	Interface name.
Service set	Service set name.
Pool Name	Name of the source address pool.
Address Range	IP address or IP address range for the pool.
Routing Instance	Name of the routing instance.
PAT	Whether Port Address Translation (PAT) is enabled (yes or no).

Table 80: show services nat source summary Output Fields (*continued*)

Field Name	Description
Total Address	Number of IP addresses that are in use.
Rule name	Name of the rule.
Rule set	Set of rules.
Match-direction	Specifies the direction in which to match traffic that meets the rule conditions.
Action	Action taken for a packet that matches a rule.

Sample Output

show services nat source summary

user@host> show services nat source summary service-set ssl_interface_style11 interface vms-0/2/0

```

Interface: vms-0/2/0 , Service set: ssl_interface_style1
Pool          Address          Routing          PAT   Total
Name          Range            Instance         Address
src_pool1     44.0.0.1-44.0.0.254  default         yes   254
Interface: vms-0/2/0 , Service set: ssl_interface_style1
Rule name Rule set      Match-direction  Action
r1        rs1          input           src_pool1

```

show services pcsp statistics

Syntax

```
show services pcsp statistics
```

Release Information

Command introduced in Junos OS Release 13.2

Command introduced in Junos OS Release 20.1R1 for Next Gen Services

Description

Display information PCP mappings.

Required Privilege Level

view

List of Sample Output

[show services pcsp statistics pcsp on page 888](#)

Output Fields

[Table 81 on page 886](#) lists the output fields for the **show services pcsp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 81: show services pcsp statistics Output Fields

Field Name	Field Description
Services PIC Name	Name of a service interface.
Protocol Statistics	Overall PCP statistics, consisting of: operational, option, and results statistics.
Operational Statistics	Operational statistics group.
Map request received	Total PCP MAP requests received from PCP clients.
Peer request received	Number of peer requests received.
Option Statistics	Number of requests using available options.
Unprocessed requests received	Number of requests received with no option specified.
Third party requests received	Number of third-party requests received.
Prefer fail option received	Number of prefer fail requests received.

Table 81: show services pcsp statistics Output Fields (*continued*)

Field Name	Field Description
Filter option received	Number of filter option requests received.
Other options counters	Number of packets received with options other than prefer-fail and third-party .
Other optional received	
Results Statistics	Information about the results of PCP requests.
PCP success	Number of PCP MAP requests successfully processed by the server.
PCP unsupported version	Number of PCP packets received with version other than 1.
Not authorized	Number of unauthorized MAP delete requests.
Bad requests	Number of requests with invalid PCP packets.
Unsupported opcode	Number of packets that have an unsupported opcode.
Unsupported option	Number of packets that have an unsupported option.
Bad option	Number of packet that have a malformed option.
Network failure	Number of times a mapping could not be provided due to a network failure.
Out of resources	Number of times a mapping could not be provided because the PCP server ran out of pool resources.
Unsupported protocol	Number of requests for which the protocol was neither TCP nor UDP.
User exceeded quota	Number of requests for which the PCP client requested more than the configured number of ports.
Cannot provide external	Number of requests for which the PCP server cannot provide the external address or port requested by the client.
Address mismatch	Number of requests for which the PCP client IP address and the layer-3 source IP do not match.
Excessive number of remote peers	This counter is not currently used.

Table 81: show services pcip statistics Output Fields (*continued*)

Field Name	Field Description
Processing error	Number of requests with malformed PCP packets information, such as an invalid IP address in a third-party request .
Other result counters	Not currently used.

Sample Output

show services pcip statistics pcip

user@host> show services pcip statistics pcip

```

Services PIC Name:      sp-2/1/0

Protocol Statistics:

Operational Statistics

Map request received           : 0
Peer request received          : 0
Other operational counters     : 0

Option Statistics

Unprocessed requests received  : 0
Third party requests received  : 0
Prefer fail option received    : 0
Filter option received         : 0
Other options counters        : 0
Option optional received       : 0

Result Statistics

PCP success                    : 0
PCP unsupported version        : 0
Not authorized                 : 0
Bad requests                   : 0
Unsupported opcode             : 0
Unsupported option             : 0
Bad option                     : 0

```

Network failure	: 0
Out of resources	: 0
Unsupported protocol	: 0
User exceeded quota	: 0
Cannot provide external	: 0
Address mismatch	: 0
Excessive number of remote peers	: 0
Processing error	: 0
Other result counters	: 0

show services policies

Syntax

```
show services policies
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display services policy information.

Required Privilege Level

view

List of Sample Output

[show services policies on page 891](#)

Output Fields

[Table 82 on page 890](#) lists the output fields for the **show services policies** command. Fields are listed in the approximate order in which they appear.

Table 82: show services policies Output Fields

Field Name	Description
Default policy	
Policy	Name of the applicable policy.
State	Status of the policy: <ul style="list-style-type: none"> • enabled: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • disabled: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Scope policy	

Table 82: show services policies Output Fields (*continued*)

Field Name	Description
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1,2,3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1,2,3,4.
Stateful firewall rule	
Service set	Name of the service set.
Interface	Name of the interface.
Match direction	
Source addresses	Names of the source addresses for a policy. Address sets are resolved to their individual Names of the source addresses for a policy. Address sets are resolved to their individual
Destination addresses	Name of the destination address (or address set as it was entered on the destination zone's address book.
Application	

Sample Output

show services policies

user@host> **show services policies**

```

Default policy: deny-all
  Policy: p1, State: enabled, Index: 1007, Scope Policy: 0, Sequence number: 1
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
  Source addresses: any-ipv4
  Destination addresses: any
  Applications: junos-ftp
  Policy: p2, State: enabled, Index: 1008, Scope Policy: 0, Sequence number: 2
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
  Source addresses: any

```

Destination addresses: any

Applications: any

show services policies detail

Syntax

```
show services policies detail
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display detailed information about configured services policies.

Required Privilege Level

view

List of Sample Output

[show services policies detail on page 894](#)

Output Fields

[Table 83 on page 893](#) lists the output fields for the **show services policies detail** command. Output fields are listed in the approximate order in which they appear.

Table 83: show services policies detail

Field Name	Description
Default policy	
Policy	
Action type	
State	<p>Status of the policy:</p> <ul style="list-style-type: none"> • enabled: The policy can be used in the policy lookup process, which determines access rights for a packet and the action taken in regard to it. • disabled: The policy cannot be used in the policy lookup process, and therefore it is not available for access control.
Index	Internal number associated with the policy.
Scope policy	
Policy type	

Table 83: show services policies detail (continued)

Field Name	Description
Sequence number	Number of the policy within a given context. For example, three policies that are applicable in a from-zoneA-to-zoneB context might be ordered with sequence numbers 1,2,3. Also, in a from-zoneC-to-zoneD context, four policies might have sequence numbers 1,2,3,4.
Stateful firewall rule	
Service set	Service set name.
Interface	Interface name.
Source addresses	The names and corresponding IP addresses for the policy. Address sets are resolved to their individual address name-IP address pairs.
Destination addresses	Name of the destination address (or address set) as it was entered in the destination zone's address book. A packet's destination address must match this value for the policy to apply to it.
Application	
IP protocol	
Inactivity timeout	
Source port range	
Destination port range	
Per policy TCP Options	

Sample Output

show services policies detail

user@host> show services policies detail

```
Default policy: deny-all
Policy: p1, action-type: permit, State: enabled, Index: 1007, Scope Policy: 0
Policy Type: Configured
```



```

Sequence number: 1
Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
Source addresses:
  any-ipv4(global): 0.0.0.0/0
Destination addresses:
  any-ipv4(global): 0.0.0.0/0
  any-ipv6(global): ::/0
Application: junos-ftp
  IP protocol: tcp, ALG: ftp, Inactivity timeout: 1800
  Source port range: [0-0]
  Destination port range: [21-21]
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No
Policy: p2, action-type: permit, State: enabled, Index: 1008, Scope Policy: 0
Policy Type: Configured
Sequence number: 2
Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
Source addresses:
  any-ipv4(global): 0.0.0.0/0
  any-ipv6(global): ::/0
Destination addresses:
  any-ipv4(global): 0.0.0.0/0
  any-ipv6(global): ::/0
Application: any
  IP protocol: 0, ALG: 0, Inactivity timeout: 0
  Source port range: [0-0]
  Destination port range: [0-0]
Per policy TCP Options: SYN check: No, SEQ check: No, Window scale: No

```

show services policies hit-count

Syntax

```
show services policies hit-count
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display the hit count of policies.

Required Privilege Level

view

List of Sample Output

[show services policies hit-count on page 896](#)

Output Fields

Sample Output

show services policies hit-count

user@host> show services policies hit-count

Index	Service Set	Interface	Name	Sfw rule	Direction
	Policy count				
1	JNPR-NH-SSET	vms-0/2/0	p1	sfw1	
input	0				
2	JNPR-NH-SSET	vms-0/2/0	p2	sfw1	
input	0				
Number of policy: 2					

show services policies interface

Syntax

```
show services policies interface interface-name
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display services policies for the specified interface.

Required Privilege Level

view

List of Sample Output

[show services policies interface vms-0/2/0 on page 897](#)

Output Fields

Sample Output

```
show services policies interface vms-0/2/0
```

```
user@host> show services policies interface vms-0/2/0
```

```
Default policy: deny-all
  Policy: p1, State: enabled, Index: 1007, Scope Policy: 0, Sequence number: 1
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
  Source addresses: any-ipv4
  Destination addresses: any
  Applications: junos-ftp
  Policy: p2, State: enabled, Index: 1008, Scope Policy: 0, Sequence number: 2
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
  Source addresses: any
  Destination addresses: any
  Applications: any
```

show services policies service-set

Syntax

```
show services policies service-set service-set
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display policy information for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services policies service-set on page 898](#)

Output Fields

Sample Output

show services policies service-set

```
user@host> show services policies service-set JNPR-NH-SSET
```

```
Default policy: deny-all
  Policy: p1, State: enabled, Index: 1007, Scope Policy: 0, Sequence number: 1
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
  Source addresses: any-ipv4
  Destination addresses: any
  Applications: junos-ftp
  Policy: p2, State: enabled, Index: 1008, Scope Policy: 0, Sequence number: 2
    Stateful firewall rule: sfw1, Service set: JNPR-NH-SSET, Interface: vms-0/2/0,
Match Direction: input
  Source addresses: any
  Destination addresses: any
  Applications: any
```

show services redundancy-group

Syntax

```
show services redundancy-group
<rg-id>
<brief | extensive | terse>
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Description

Display redundancy group status information for all redundancy groups or a specified redundancy group.

Options

rg-id—(Optional) Name of a specific redundancy group.

brief | extensive | terse—(Optional) Display the specified level of output. When no level is specified, display terse level output.

Default: terse

Required Privilege Level

view

List of Sample Output

[show services redundancy-group terse on page 905](#)

[show services redundancy-group brief \(Health Status Passed\) on page 905](#)

[show services redundancy-group brief \(Health Status Failed\) on page 906](#)

[show services redundancy-group extensive on page 907](#)

Output Fields

[Table 84 on page 899](#) lists the output fields for the **show services redundancy-group** command. Output fields are listed in the approximate order in which they appear.

Table 84: show services redundancy-group Output Fields

Field Name	Field Description	Level of Output
ICCP process connection	Status of the connection between the srd and iccpd. <ul style="list-style-type: none"> Connected Not connected 	all levels
Redundancy Group ID	Identifier of the redundancy group.	all levels

Table 84: show services redundancy-group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Number of peer RG connections	Total number of peers in the redundancy group.	brief, extensive
Local RG IP	IP address of the local redundancy group.	all levels
RS ID		terse
Local RS state	State of the local redundancy set. <ul style="list-style-type: none"> • MASTER • STANDBY • INITIALIZING • STANDBY (WARNED) 	terse
Peer RS state	State of the peer redundancy set. <ul style="list-style-type: none"> • MASTER • STANDBY • INITIALIZING • STANDBY (WARNED) 	terse
Peer RG IP	Peer redundancy group IP address.	all
Status	Status of redundancy group connection with this peer. <ul style="list-style-type: none"> • Connected • Not Connected 	terse
Number of peer RG connections	Total number of peers in the redundancy group.	brief
Redundancy Set ID	Identifier of the redundancy set.	brief, extensive
Connection status	Status of the connection between the srd and iccpd. <ul style="list-style-type: none"> • Connected • Not Connected 	brief, extensive

Table 84: show services redundancy-group Output Fields (continued)

Field Name	Field Description	Level of Output
Redundancy Set state	State of the local redundancy set state. <ul style="list-style-type: none">• INITIALIZING• MASTER• STANDBY• STANDBY (WARNED)	brief, extensive
Redundancy Set peer state	State of the peer redundancy set state. <ul style="list-style-type: none">• INITIALIZING• MASTER• STANDBY• STANDBY (WARNED)	brief, extensive
Redundancy Set health status	<ul style="list-style-type: none">• Passed• Failed	brief, extensive
Number of Monitored interface down	Number of monitored interfaces that are d	brief, extensive
Failed Interfaces	List of all monitored interfaces that are down.	brief, extensive
Service Set	Service set used for stateful sync.	brief, extensive
Service Interface	Service set used for	brief, extensive
Type	Type of redundancy and stateful sync for the listed service interface. <ul style="list-style-type: none">• Inter-chassis• Intra-chassis	brief, extensive
Role	Role of the listed service interface. <ul style="list-style-type: none">• active• backup	brief, extensive
Connection	Status of connection with peer service PIC. <ul style="list-style-type: none">• Up• Down	brief, extensive

Table 84: show services redundancy-group Output Fields (continued)

Field Name	Field Description	Level of Output
Synchronization	<p>Type of synchronization. When all eligible sessions are still synchronizing, it is cold synchronization. When all current existing sessions are synchronized, it is a HOT synchronization, When long lived sessions are eligible, they are synchronized.</p> <ul style="list-style-type: none"> • Hot—All current existing sessions are synced. When long-lived sessions are eligible, they are synchronized. • Cold—Eligible sessions are in the processing of synchronizing. 	brief, extensive
ICCP process connection open complete count	Number of completed opens of ICCP process connections.	extensive
ICCP process connection close complete count	Number of completed closes of ICCP process connections.	
ICCP packet sent count	Number of ICCP packets sent.	extensive
ICCP packet receive count	Number of ICCP packets received.	extensive
ICCP process keepalive receive count	Number of ICCP process keepalive messages received.	extensive
ICCP process keepalive sent count	Number of ICCP process keepalive messages sent.	extensive
ICCP redundancy group add count	Number of redundancy group add messages received by srd from ICCP.	extensive
ICCP redundancy group delete count	Number of redundancy group delete messages received by srd from ICCP.	extensive
RG connection up count	Number of redundancy group connection up messages received by srd from ICCP.	extensive
RG connection down count	Number of redundancy group connection down messages received by srd from ICCP.	extensive

Table 84: show services redundancy-group Output Fields (*continued*)

Field Name	Field Description	Level of Output
RG join count	Number of redundancy group join messages sent from srd to ICCP.	extensive
RG data receive count	Number of packets of messages received by srd from a peer.	extensive
RG data sent count	Number of packets of messages sent from srd to a peer.	extensive
RG connect message sent count	Number of connect messages sent from srd to ICCP.	extensive
RG connect message receive count	Number of connect messages received by srd from ICCP.	extensive
RG disconnect message sent count	Number of disconnect messages sent from srd to ICCP.	extensive
RG disconnect message receive count	Number of disconnect messages received by srd from ICCP.	extensive
RG ack sent count	Number of RG ack messages sent.	extensive
RG nack sent count	Number of RG nack messages sent.	extensive
RG nack receive count	Number of RG nack messages received.	extensive
Transition Events Received	Number of transition events received in each of the following categories: <ul style="list-style-type: none"> • Acquire mastership auto • Acquire mastership manual • Release mastership auto • Release mastership manual 	extensive

Table 84: show services redundancy-group Output Fields (*continued*)

Field Name	Field Description	Level of Output
Transition Events Ignored	<p>Number of transition events ignored in each of the following categories:</p> <ul style="list-style-type: none"> • Acquire mastership auto • Acquire mastership manual • Release mastership auto • Release mastership manual <p>In a high-availability or redundancy pair of SDGs, in which one SDG is the master and the other is the standby, when perform a double failover of the SDGs, the second failover event is not ignored, which is the expected behavior. The event is not disregarded because it arrives as a critical redundancy-event based on the redundancy-policy. However, because the SDG is already be in Standby state, the finite state machine transitions to the Standby-Warned state until it recovers. Therefore, the event is honored and not ignored. Although there was no mastership transition, it is because of a valid reason that the SDG is already in Standby state. The redundancy-event is associated with to a mastership release policy based on the configuration and the Release mastership field under the Transition Events Ignored column displays a number that corresponds to the redundancy event.</p> <p>The services redundancy daemon (SRD) finite state machine quickly recovers (transitions from Standby-Warned to Standby) during restart-routing because the rpd restart-handling and recovery are fast and the following critical event is not ignored. However, disabling or deactivating the interface results in the FSM remaining in Standby-Warned until the interface is up. Any critical events during the time when the interface is down are ignored because the state is already Standby-Warned and does not transition to a different state. In summary, the following is the manner in which critical events are analyzed during state transitions:</p> <ul style="list-style-type: none"> • Standby -> Standby Warned = Critical Event Not ignored [valid state transition] • Standby Warned -> Standby Warned = Critical Event Ignored [no state transition] 	extensive

Table 84: show services redundancy-group Output Fields (continued)

Field Name	Field Description	Level of Output
Monitored Events Received	Number of monitored events received in each of the following categories: <ul style="list-style-type: none"> • Link-down • Routing restart/abort • Route update error • Peer mastership-acquire • Peer mastership-release 	extensive
Monitored Events Ignored	Number of monitored events ignored in each of the following categories: <ul style="list-style-type: none"> • Link-down • Routing restart/abort • Route update error • Peer mastership-acquire • Peer mastership-release 	extensive

Sample Output

show services redundancy-group terse

user@host> **show services redundancy-group terse**

```

ICCP process connection           : Connected

Redundancy Group ID              : 1
Number of peer RG connections    : 1
Local RG IP                      : 172.19.39.70
RS ID      Local RS state      Peer RS state    Peer RG IP      Status
1          MASTER              STANDBY          172.19.39.69    Connected

```

show services redundancy-group brief (Health Status Passed)

user@host> **show services redundancy-group brief**

```

ICCP process connection           : Connected
Redundancy Group ID              : 1

```

```

Number of peer RG connections      : 1
Local RG IP                       : 172.19.39.70
Redundancy Set ID                 : 1
  Connection status                : Connected
  Redundancy Set state             : MASTER
  Redundancy Set peer state       : STANDBY
Peer RG IP                       : 172.19.39.69
Redundancy Set health status      : Passed
  Service Set : IPv6-SFW
    Service interface  Type      Role      Connection
Synchronization
    ms-1/3/0          Inter-chassis active    Up          Hot
    ms-1/2/0          Inter-chassis active    Up          Hot
    ms-1/1/0          Inter-chassis active    Up          Hot
    ms-1/0/0          Inter-chassis active    Up          Hot
  Service Set : NAPT44-SS1-SS4
    Service interface  Type      Role      Connection
Synchronization
    ms-1/3/0          Inter-chassis active    Up          Hot
    ms-1/2/0          Inter-chassis active    Up          Hot
    ms-1/1/0          Inter-chassis active    Up          Hot
    ms-1/0/0          Inter-chassis active    Up          Hot

```

show services redundancy-group brief (Health Status Failed)

user@host> show services redundancy-group brief

```

ICCP Process Connection          : Connected
Redundancy Group ID             : 1
  Number of Members              : 2
Redundancy Set ID               : 1
  Remote IP address              : 203.0.113.2
  Connection Status              : Connected
  Redundancy Set State           : STANDBY (WAIT)
  Redundancy Set Peer State      : MASTER
  Redundancy Set Health Status   : Failed
    Number of Monitored interface down : 1          <<<<<<< Failure Reasons

```

```

Failed Interfaces
<<<<<< Name of the monitored interfaces which have gone down
ms-2/3/0
Service Set : ss2
Service Interface      Type                Role                Connection
Synchronization
ms-2/2/0                Inter-chassis      backup              Up
Hot
ms-2/1/0                Inter-chassis      backup              Down
Off
ms-2/0/0                Inter-chassis      backup              Down
Off
Service Set : ss_new
Service Interface      Type                Role                Connection
Synchronization
ms-2/3/0

```

show services redundancy-group extensive

```
user@host> show services redundancy-group extensive
```

```

ICCP process connection           : Connected
ICCP process connection close count : 0
ICCP process connection open complete count : 1
ICCP packet sent count           : 7303
ICCP packet receive count         : 7321
ICCP process keepalive receive count : 7253
ICCP process keepalive sent count  : 7253
ICCP redundancy group add count    : 0
ICCP redundancy group delete count : 0
Redundancy Group ID               : 1
Number of peer RG connections     : 1
Local RG IP                       : 172.19.39.70
RG connection up count            : 4
RG connection down count          : 2
RG join count                     : 4
RG data receive count             : 37
RG data sent count                : 0
RG connect message sent count     : 4
RG connect message receive count  : 4
RG disconnect message sent count  : 0
RG disconnect message receive count : 4
RG ack sent count                 : 4

```

```

RG nack sent count           : 0
RG nack receive count        : 4
Redundancy Set ID            : 1
  Connection status           : Connected
  Redundancy Set state         : MASTER
  Redundancy Set peer state    : STANDBY
  Peer RG IP                   : 172.19.39.69
  Redundancy Set health status : Passed

```

Service Set : IPv6-SFW

Service interface	Type	Role	Connection	
Synchronization				
ms-1/3/0	Inter-chassis	active	Up	Hot
ms-1/2/0	Inter-chassis	active	Up	Hot
ms-1/1/0	Inter-chassis	active	Up	Hot
ms-1/0/0	Inter-chassis	active	Up	Hot

Service Set : NAPT44-SS1-SS4

Service interface	Type	Role	Connection	
Synchronization				
ms-1/3/0	Inter-chassis	active	Up	Hot
ms-1/2/0	Inter-chassis	active	Up	Hot
ms-1/1/0	Inter-chassis	active	Up	Hot
ms-1/0/0	Inter-chassis	active	Up	Hot

Transition events	Received	Ignored
Acquire mastership auto	3	0
Acquire mastership manual	0	0
Release mastership auto	3	0
Release mastership manual	0	0

Monitored events	Received	Ignored
Link-down	145	31
Routing restart/abort	1	0
Route update error	0	0
Peer mastership-acquire	3	0
Peer mastership-release	3	0

show services screen ids-option (Next Gen Services)

Syntax

```
show services screen <ids-option>
screen-name
logical-system
root-logical-system
tenant
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display the configuration information about the specified services screen. You can configure a **ids-option** to enable screen protection on the MX Series devices.

Options

- **screen-name** —Name of the screen.
- **logical-system**—Name of the logical system.
- **root-logical-system**—Displays root logical system as default.
- **tenant | all**—Name of the tenant system or all tenants.

Required Privilege Level

view

RELATED DOCUMENTATION

| *ids-option*

List of Sample Output

[show services screen ids-option on page 909](#)

Sample Output

```
show services screen ids-option
```

```
user@host> show services screen ids-option <option1>
```

Screen object status:

Name	Value
ICMP flood threshold	0
UDP flood threshold	0
TCP winnuke	enabled
TCP port scan threshold	0
ICMP address sweep threshold	0
TCP sweep threshold	0
UDP sweep threshold	0
IP tear drop	enabled
TCP SYN flood attack threshold	0
TCP SYN flood alarm threshold	0
TCP SYN flood source threshold	0
TCP SYN flood destination threshold	0
TCP SYN flood timeout	0
ICMP ping of death	enabled
IP source route option	enabled
TCP land attack	enabled
TCP SYN fragment	enabled
TCP no flag	enabled
IP unknown protocol	enabled
IP bad options	enabled
IP record route option	enabled
IP timestamp option	enabled
IP security option	enabled
IP lose source route option	enabled
IP stream option	enabled
ICMP fragmentation	enabled
ICMP large packet	enabled
TCP SYN FIN	enabled
TCP FIN no ACK	enabled
Session source limit threshold	0
Session destination limit threshold	0
Alarm without drop	enabled

show services screen-statistics service-set (Next Gen Services)

Syntax

```
show services screen statistics service-set service-set
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display intrusion detection service (IDS) screen statistics.

Options

- **screen-name**—Name of the screen.
- **logical-system**—Name of the logical system.
- **root-logical-system**—Displays root logical system as default.
- **tenant**—Name of the tenant system.

Required Privilege Level

view

RELATED DOCUMENTATION

ids-option

Example: Configuring Multiple Screening Options

List of Sample Output

[show services screen statistics service-set on page 914](#)

Output Fields

[Table 85 on page 911](#) lists the output fields for the **show services screen statistics service-set** command. Output fields are listed in the approximate order in which they appear.

Table 85: show services screen statistics service-set Output Fields

Field Name	Field Description
ICMP flood	Internet Control Message Protocol (ICMP) flood counter. An ICMP flood typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.

Table 85: show services screen statistics service-set Output Fields *(continued)*

Field Name	Field Description
UDP flood	User Datagram Protocol (UDP) flood counter. UDP flooding occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.
TCP winnuke	Number of Transport Control Protocol (TCP) WinNuke attacks. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP port scan	Number of TCP port scans. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
ICMP address sweep	Number of ICMP address sweeps. An IP address sweep can occur with the intent of triggering responses from active hosts.
IP tear drop	Number of teardrop attacks. Teardrop attacks exploit the reassembly of fragmented IP packets.
TCP SYN flood	Number of TCP SYN attacks.
IP spoofing	Number of IP spoofs. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
ICMP ping of death	ICMP ping of death counter. Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).
IP source route option	Number of IP source route attacks.
TCP address sweep	Number of TCP address sweeps.
TCP land attack	Number of land attacks. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and source IP address.
TCP SYN fragment	Number of TCP SYN fragments.
TCP no flag	Number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
IP unknown protocol	Number of IPs.

Table 85: show services screen statistics service-set Output Fields (continued)

Field Name	Field Description
IP bad options	Number of invalid options.
IP record route option	Number of packets with the IP record route option enabled. This option records the IP addresses of the network devices along the path that the IP packet travels.
IP timestamp option	Number of IP timestamp option attacks. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination.
IP security option	Number of IP security option attacks.
IP loose source route option	Number of IP loose source route option attacks. This option specifies a partial route list for a packet to take on its journey from source to destination.
IP strict source route option	Number of IP strict source route option attacks. This option specifies the complete route list for a packet to take on its journey from source to destination.
IP stream option	Number of stream option attacks. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams.
ICMP fragment	Number of ICMP fragments. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.
ICMP large packet	Number of large ICMP packets.
TCP SYN FIN	Number of TCP SYN FIN packets.
TCP FIN no ACK	Number of TCP FIN flags without the acknowledge (ACK) flag.
Source session limit	Number of concurrent sessions that can be initiated from a source IP address.
TCP SYN-ACK-ACK proxy	Number of TCP flags enabled with SYN-ACK-ACK. To prevent flooding with SYN-ACK-ACK sessions, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold and SRX Series devices running Junos OS reject further connection requests from that IP address.
IP block fragment	Number of IP block fragments.

Table 85: show services screen statistics service-set Output Fields (continued)

Field Name	Field Description
Destination session limit	Number of concurrent sessions that can be directed to a single destination IP address.

Sample Output

show services screen statistics service-set

```
user@host> show services screen statistics service-set USF-Service-Set-X
```

```

Screen statistics:

IDS attack type                                Statistics
ICMP flood                                     0
UDP flood                                     0
TCP winnuke                                   0
TCP port scan                                 0
ICMP address sweep                           0
TCP sweep                                     0
UDP sweep                                     0
IP tear drop                                 0
TCP SYN flood                                0
ICMP ping of death                           0
IP source route option                       0
TCP land attack                              0
TCP SYN fragment                             0
TCP no flag                                  0
IP unknown protocol                          0
IP bad options                               0
IP record route option                       0
IP timestamp option                          0
IP security option                           0
IP loose source route option                 0
IP strict source route option                0
IP stream option                             0
ICMP fragment                                0
ICMP large packet                            0
TCP SYN FIN                                  0
TCP FIN no ACK                              0

```

Source session limit	0
TCP SYN-ACK-ACK proxy	0
IP block fragment	0
Destination session limit	0

show services security-intelligence category summary

Syntax

```
show services security-intelligence category summary category-name
```

Release Information

Statement introduced before Junos OS Release 18.4.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card.
Support for threat feed status (enabled, disabled, or user disabled) is added in Junos OS Release 20.1R1.

Description

Display summary for the specified Security Intelligence category.

Options

category-name—Name of the category.

Required Privilege Level

View

RELATED DOCUMENTATION

[security-intelligence](#) | 665

List of Sample Output

[show services security-intelligence category summary on page 917](#)

Output Fields

[Table 86 on page 916](#) lists the output fields for the **show services security-intelligence category summary** command. Output fields are listed in the approximate order in which they appear.

Table 86: show services security-intelligence category summary Output Fields

Field Name	Field Description
Category name	Name of the Security Intelligence category.
Status	Status of the Security Intelligence category.
Description	Description of the Security Intelligence category
Update interval	Amount of time after which Policy Enforcer sends an update for the feed.

Table 86: show services security-intelligence category summary Output Fields (*continued*)

Field Name	Field Description
TTL	Length of time (in minutes) the file remains open, receiving statistics before it is closed, transferred, and rotated. When either the time or the file size is exceeded, the file is closed and a new one is opened, whether or not a transfer site is specified.
Feed name	Information about the feed, including: <ul style="list-style-type: none"> • Version • Object umber • Create time • Update time • Update status • Expired • Options • Status

Sample Output

show services security-intelligence category summary

user@host> show services security-intelligence category summary

```

node1:
-----

Category name      :CC
Status             :Enable
Description        :Command and Control data schema
Update interval    :1800s
TTL                :3456000s
Feed name          :cc_ip_data
  Version          :N/A
  Objects number   :0
  Create time      :2018-03-16 05:57:39 PDT
  Update time      :2018-03-19 12:30:32 PDT
  Update status    :N/A
  Expired          :No
  Options          :N/A
  Status           :Enabled

```

```
Feed name      :cc_ipv6_data
Version        :20180228.1
Objects number:1
Create time    :2018-03-16 05:57:39 PDT
Update time    :2018-03-16 06:19:47 PDT
Update status  :Store succeeded
Expired        :No
Options        :N/A
Status         :Disabled
```


show services security-intelligence update status

Syntax

```
show services security-intelligence update status
```

Release Information

Statement introduced before Junos OS Release 18.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card.

Description

Display the status of the connection with Policy Enforcer.

Required Privilege Level

View

RELATED DOCUMENTATION

[security-intelligence](#) | [665](#)

List of Sample Output

[show services security-intelligence update status on page 919](#)

Sample Output

```
show services security-intelligence update status
```

```
user@host> show services security-intelligence update status
```

```
node1:
-----
Current action      :Start downloading the latest manifest.
Last update status  :Download manifest failed.
Last connection status:succeeded
Last update time    :2018-03-21 16:59:59 PDT
```

show services service-sets cpu-usage

Syntax

```
show services service-sets cpu-usage
<interface interface-name>
<service-set service-set-name>
```

Release Information

Command introduced before Junos OS Release 7.4.
 Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display service set CPU usage as a percentage. The command is supported only on Adaptive Services PICs (SP PICs).

Options

- none**—Display CPU usage for all adaptive services interfaces and service sets.
- interface interface-name**—(Optional) Display CPU usage for a particular interface. On M Series and T Series routers, the *interface-name* parameter can have the value *sp-fpc/pic/port* or *rspnumber*.
- service-set service-set-name**—(Optional) Display CPU usage for a particular service set. For the Layer 2 Tunneling Protocol (L2TP), you can use a tunnel group to represent a service set.

Required Privilege Level

view

List of Sample Output

[show services service-sets cpu-usage on page 921](#)

Output Fields

[Table 87 on page 920](#) lists the output fields for the **show services service-sets cpu-usage** command. Output fields are listed in the approximate order in which they appear.

Table 87: show services service-sets cpu-usage Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface

Table 87: show services service-sets cpu-usage Output Fields (*continued*)

Field Name	Field Description
Service set (system category)	<p>Name of the CPU usage category:</p> <ul style="list-style-type: none"> • idp_recommended—Name of the service sets (displays all the service sets attached to the service PICs) • Idle • System • Receive • Transmit
CPU utilization %	Percentage of the CPU resources being used

Sample Output

show services service-sets cpu-usage

user@host> **show services service-sets cpu-usage**

Interface	Service set (system category)	CPU utilization %
sp-4/1/0	idp_recommended	18.20 %
sp-4/1/0	Idle	44.69 %
sp-4/1/0	System	7.01 %
sp-4/1/0	Receive	15.10 %
sp-4/1/0	Transmit	15.00 %

show services service-sets memory-usage

Syntax

```
show services service-sets memory-usage
<interface interface-name>
<service-set service-set-name>
<zone>
```

Release Information

Command introduced before Junos OS Release 7.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display service set memory usage.

Options

none—Display service set memory usage.

interface *interface-name*—(Optional) Display memory usage for a particular interface. On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port*, or *rspnumber*.

NOTE: This command is not supported on Multilink Protocol-based services PICs.

The interface option is not supported on Multiservice PICs.

service-set *service-set-name*—(Optional) Display memory usage for a particular service set. For Layer 2 Tunneling Protocol (L2TP), you can use a tunnel group to represent a service set.

zone—(Optional) Display the memory usage zone of the adaptive services interface or an individual service set.

Required Privilege Level

view

List of Sample Output

[show services service-sets memory-usage on page 923](#)

[show services service-sets memory-usage zone on page 923](#)

[show services service-sets memory-usage interface on page 923](#)

Output Fields

Table 88 on page 923 lists the output fields for the **show services service-sets memory-usage** command. Output fields are listed in the approximate order in which they appear.

Table 88: show services service-sets memory-usage Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface
Service set	Name of a service set
Bytes Used	Number of bytes of memory being used
Memory zone	Memory zone in which the adaptive services interface is currently operating: <ul style="list-style-type: none"> • Green—All new flows are allowed. • Yellow—Unused memory is reclaimed. All new flows are allowed. • Orange—New flows are allowed only for service sets that are using less than their equal share of memory. • Red—No new flows are allowed.

Sample Output

show services service-sets memory-usage

```
user@host> show services service-sets memory-usage
```

Interface	Service set	Bytes Used
ms-4/0/0	N/A	14817036
ms-4/1/0	N/A	14691700

show services service-sets memory-usage zone

```
user@host> show services service-sets memory-usage zone
```

Interface	Memory zone
-----------	-------------

show services service-sets memory-usage interface

```
user@host> show services service-sets memory-usage interface ms-4/1/0
```

Interface	Service Set	Bytes Used
ms-4/1/0	N/A	14691700

show services service-sets plug-ins

Syntax

```
show services service-sets plug-ins <interface interface-name>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display service set plug-ins summary.

Options

interface *interface-name*—Display service set plug-ins information for the specified interface.

Required Privilege Level

view

List of Sample Output

[show services service-sets plug-ins on page 925](#)

Sample Output

show services service-sets plug-ins

user@host> **show services service-sets plug-ins**

```
Interface: vms-0/2/0
  Service-set: ssl_interface_style1, State: Ready
  Plugins configured: 1
  Plugin: junos-alg, ID: 25
```

show services service-sets statistic screen-drops (Next Gen Services)

Syntax

```
show services service-sets statistic screen-drops [service-set] interface interface-name
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display statistics for packet drops resulting from header-integrity, suspicious packet pattern, and session-limit checks performed by an MS-MPC or MS-MIC.

Options

none—Display statistics for all configured service interfaces and service sets.

<interface *interface-name*>—(Optional) Display statistics for the specified services interface.

<service-set *service-set-name* >—(Optional) Display statistics for the specified service set.

Required Privilege Level

view

RELATED DOCUMENTATION

| [Configuring Protection Against Network Attacks on an MS-MPC](#)

List of Sample Output

[show services service-sets statistic screen-drops on page 933](#)

Output Fields

[Table 89 on page 926](#) lists the output fields for the **show services service-set integrity-drops** command. Output fields are listed in the approximate order in which they appear.

Table 89: show services service-set statistics screen drops Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.

Table 89: show services service-set statistics screen drops Output Fields (*continued*)

Field Name	Field Description
Errors	<p>Total errors, categorized by protocol:</p> <ul style="list-style-type: none">• IP—Total IP version 4 errors.• TCP—Total Transmission Control Protocol (TCP) errors.• UDP—Total User Datagram Protocol (UDP) errors.• ICMP—Total Internet Control Message Protocol (ICMP) errors.

Table 89: show services service-set statistics screen drops Output Fields (*continued*)

Field Name	Field Description
IP Errors	

Table 89: show services service-set statistics screen drops Output Fields (*continued*)

Field Name	Field Description
	<p>Number of IPv4 errors for the following categories:</p> <ul style="list-style-type: none"> • IP packet length inconsistencies—IP packet length did not match the Layer 2 reported length. • Minimum IP header length check failures—Minimum IP header length is 20 bytes. The received packet contained less than 20 bytes. • Reassembled packet exceeds maximum IP length—After fragment reassembly, the reassembled IP packet length exceeded 65,535. • Illegal source address 0—Source address is not a valid address. Invalid addresses are loopback, broadcast, multicast, and reserved addresses. Source address 0, however, is allowed to support BOOTP and the destination address 0xffffffff. • Illegal destination address —Destination address was not a valid address. The address is reserved. • TTL zero errors—Received packet had a time-to-live (TTL) value of 0. • Illegal IP protocol number 0 or 255—IP protocol is 0 or 255. • Land attack—IP source address is the same as the destination address. • Non-IP packets—Packet did not conform to the IP standard. • IP option—Packet had a non-allowed IP option. • Non-IPv4 packets—Packet was not of the IPv4 type. • Non-IPv6 packets—Packet was not of the IPv6 type. • Bad checksum—Packet had an invalid IP checksum. • Illegal IP fragment length—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • IP fragment overlap—Fragments had overlapping fragment offsets. • IP fragment limit exceeded —Configured number of allowed fragments for a packet was exceeded. • IP fragment reassembly timeout—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. Whenever a fragment is received, it is maintained in a chain until all other fragments are received. If other fragments do not arrive

Table 89: show services service-set statistics screen drops Output Fields (*continued*)

Field Name	Field Description
	<p>within the configured value of reassembly-timeout, this packet is dropped and the value of the counter shown in this field is incremented. If other fragments arrive in time but the total number of fragments is more than the configured value of fragment-limit, all the fragments (of this packet) are dropped and the value of the counter shown in this field is incremented.</p> <ul style="list-style-type: none"> ● IPv4 bad options—Packet IP header contained IPv4 option that is not allowed. ● IPv6 bad extension headers—Packet contained IPv6 extension header type that is not allowed. ● session-limit exceeded for source—Number of concurrent sessions from an individual source address or subnet exceeded limit. ● session-limit exceeded for destination—Number of concurrent sessions to an individual destination address or subnet exceeded limit. ● connections/second limit exceeded for source—Number of connections per second for an individual source address or subnet exceeded limit. ● connections/second limit exceeded for destination—Number of connections per second for an individual destination address or subnet exceeded limit. ● packets/second limit exceeded for source—Number of packets per second for an individual source address or subnet exceeded limit. ● packet/second limit exceeded for destination—Number of packets per second for an individual destination address or subnet exceeded limit. ● Unknown —Unknown fragments.

Table 89: show services service-set statistics screen drops Output Fields (*continued*)

Field Name	Field Description
TCP Errors	<p>Number of TCP protocol errors for the following categories:</p> <ul style="list-style-type: none"> • TCP header length inconsistencies—Minimum TCP header length is 20 bytes, and the IP packet received did not contain at least 20 bytes. • Source or destination port number is zero—TCP source or destination port was zero. • Illegal sequence number, flags combination—Packet had any type of TCP header anomaly. • TCP winnuke—TCP segments destined for port 139 with the urgent (URG) flag set. • TCP SYN Fragment—TCP SYN packet was a fragment. • TCP connection closed due to SYN defense—Unestablished TCP connection closed because open-timeout value expired. • TCP session-limit exceeded for source—Number of concurrent TCP sessions from an individual source address or subnet exceeded limit. • TCP session-limit exceeded for destination—Number of concurrent TCP sessions to an individual destination address or subnet exceeded limit. • TCP connections/second limit exceeded for source—Number of TCP connections per second for an individual source address or subnet exceeded limit. • TCP connections/second limit exceeded for destination—Number of TCP connections per second for an individual destination address or subnet exceeded limit. • TCP packets/second limit exceeded for source—Number of TCP packets per second for an individual source address or subnet exceeded limit. • TCP packet/second limit exceeded for destination—Number of TCP packets per second for an individual destination address or subnet exceeded limit.

Table 89: show services service-set statistics screen drops Output Fields (*continued*)

Field Name	Field Description
UDP Errors	<p>Number of UDP protocol errors for the following categories:</p> <ul style="list-style-type: none"> • IP data length less than minimum UDP header length (8 bytes)—Minimum UDP header length is 8 bytes. The received IP packets contained less than 8 bytes. • Source or destination port is zero—UDP source or destination port was 0. • UDP session-limit exceeded for source—Number of concurrent UDP sessions from an individual source address or subnet exceeded limit. • UDP session-limit exceeded for destination—Number of concurrent UDP sessions to an individual destination address or subnet exceeded limit. • UDP connections/second limit exceeded for source—Number of UDP connections per second for an individual source address or subnet exceeded limit. • UDP connections/second limit exceeded for destination—Number of UDP connections per second for an individual destination address or subnet exceeded limit. • UDP packets/second limit exceeded for source—Number of UDP packets per second for an individual source address or subnet exceeded limit. • UDP packet/second limit exceeded for destination—Number of UDP packets per second for an individual destination address or subnet exceeded limit.

Table 89: show services service-set statistics screen drops Output Fields (*continued*)

Field Name	Field Description
ICMP Errors	<p>Number of ICMP protocol errors for the following categories:</p> <ul style="list-style-type: none"> • IP data length less than minimum ICMP header length (8 bytes)—ICMP header length contained less than 8 bytes. • ICMP error length inconsistencies—ICMP error packet length was outside range of 48 bytes through 576 bytes. • ICMP fragments— ICMP packet was an IP fragment. • ICMP session-limit exceeded for source—Number of concurrent ICMP sessions from an individual source address or subnet exceeded limit. • ICMP session-limit exceeded for destination—Number of concurrent ICMP sessions to an individual destination address or subnet exceeded limit. • ICMP connections/second limit exceeded for source—Number of ICMP connections per second for an individual source address or subnet exceeded limit. • ICMP connections/second limit exceeded for destination—Number of ICMP connections per second for an individual destination address or subnet exceeded limit. • ICMP packets/second limit exceeded for source—Number of ICMP packets per second for an individual source address or subnet exceeded limit. • ICMP packet/second limit exceeded for destination—Number of ICMP packets per second for an individual destination address or subnet exceeded limit.

Sample Output

show services service-sets statistic screen-drops

user@host> show services service-sets statistic screen-drops USF-Service-Set-X interface vms-0/2/0

```
Interface: vms-0/2/0
Service set: sset1
Errors:
  IP: 0, TCP: 0
  UDP: 0, ICMP: 0
IP errors:
```

```

IP packet length inconsistencies: 0
Illegal source address: 0
Illegal destination address: 0
TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
Land attack: 0
Non-IPv4 packets: 0
Non-IPv6 packets: 0
Bad checksum: 0
Illegal IP fragment length: 0
IP fragment overlap: 0
IP fragment reassembly timeout: 0
IP fragment limit exceeded: 0
IPv4 bad options: 0
IPv6 bad extension headers: 0
session-limit exceeded for source: 0
session-limit exceeded for destination: 0
connections/second limit exceeded for source: 0
connections/second limit exceeded for destination: 0
packets/second limit exceeded for source: 0
packet/second limit exceeded for destination: 0
Unknown: 0
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combinations: 0
  TCP winnuke: 0
  TCP SYN Fragment: 0
  TCP connection closed due to SYN defense: 0
  TCP session-limit exceeded for source: 0
  TCP session-limit exceeded for destination: 0
  TCP connections/second limit exceeded for source: 0
  TCP connections/second limit exceeded for destination: 0
  TCP packets/second limit exceeded for source: 0
  TCP packet/second limit exceeded for destination: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
  Source or destination port number is zero: 0
  UDP session-limit exceeded for source: 0
  UDP session-limit exceeded for destination: 0
  UDP connections/second limit exceeded for source: 0
  UDP connections/second limit exceeded for destination: 0
  UDP packets/second limit exceeded for source: 0
  UDP packet/second limit exceeded for destination: 0
ICMP errors:

```



```
IP data length less than minimum ICMP header length (8 bytes): 0
ICMP error length inconsistencies: 0
ICMP fragments: 0
ICMP session-limit exceeded for source: 0
ICMP session-limit exceeded for destination: 0
ICMP connections/second limit exceeded for source: 0
ICMP connections/second limit exceeded for destination: 0
ICMP packets/second limit exceeded for source: 0
ICMP packet/second limit exceeded for destination: 0
```

show services service-sets statistic screen-session-limit-counters (Next Gen Services)

Syntax

```
show services service-set statistic screen-session-limit-counters
<interface interface>
<service-set service-set>
```

Release Information

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display counters for session drops and packet drops resulting from session-limit checks performed by an IDS rule on an MS-MPC or MS-MIC.

Options

none—Display statistics for all configured services interfaces.

interface *interface-name*—(Optional) Display statistics for the specified services interface.

service *service-set*—Display statistics for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services service-sets statistic screen-session-limit-counters on page 940](#)

Output Fields

[Table 90 on page 936](#) lists the output fields for the **show services service-set statistics ids session-limits counters** command. Output fields are listed in the approximate order in which they appear.

Table 90: show services service-sets statistics ids session-limits counters Output Fields

Field Name	Field Description

Table 90: show services service-sets statistics ids session-limits counters Output Fields (*continued*)

Field Name	Field Description
TCP Counters	<p>Session-limit TCP counters in the ingress direction for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of TCP sessions allowed by the IDS rule. • Sessions ignored—Number of TCP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of TCP sessions dropped because the number of TCP sessions exceeded the limit. • Sessions dropped due to high rate—Number of TCP sessions dropped because the number of TCP connections per second exceeded the limit. • Packets allowed—Number of TCP packets that the IDS rule allowed. • Packets dropped due to high pps—Number of TCP packets dropped because the number of TCP packets per second exceeded the limit.
UDP Counters	<p>Session-limit UDP counters in the ingress direction for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of UDP sessions allowed by the IDS rule. • Sessions ignored—Number of UDP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of UDP sessions dropped because the number of UDP sessions exceeded the limit. • Sessions dropped due to high rate—Number of UDP sessions dropped because the number of UDP connections per second exceeded the limit. • Packets allowed—Number of UDP packets that the IDS rule allowed. • Packets dropped due to high pps—Number of UDP packets dropped because the number of TCP packets per second exceeded the limit.

Table 90: show services service-sets statistics ids session-limits counters Output Fields (*continued*)

Field Name	Field Description
ICMP Counters	<p>Session-limit ICMP counters in the ingress direction for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of ICMP sessions allowed by the IDS rule. • Sessions ignored—Number of ICMP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of ICMP sessions dropped because the number of ICMP sessions exceeded the limit. • Sessions dropped due to high rate—Number of ICMP sessions dropped because the number of ICMP connections per second exceeded the limit. • Packets allowed—Number of ICMP packets that the IDS rule allowed. • Packets dropped due to high pps—Number of ICMP packets dropped because the number of ICMP packets per second exceeded the limit.
Other-Protocols Counters	<p>Session-limit counters in the ingress direction for protocols other than TCP, UDP, and ICMP for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of sessions allowed by the IDS rule. • Sessions ignored—Number of sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of sessions dropped because the number of sessions exceeded the limit. • Sessions dropped due to high rate—Number of sessions dropped because the number of connections per second exceeded the limit. • Packets allowed—Number of packets that the IDS rule allowed. • Packets dropped due to high pps—Number of packets dropped because the number of packets per second exceeded the limit.
Egress General Info	<p>Information for IDS rules for the service set in the egress direction.</p> <ul style="list-style-type: none"> • Match-direction—Displays output. • Rule name—Name of the IDS rule. • Term name—Name of the term in the IDS rule.

Table 90: show services service-sets statistics ids session-limits counters Output Fields (*continued*)

Field Name	Field Description
Egress TCP Counters	<p>Session-limit TCP counters in the egress direction for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of TCP sessions allowed by the IDS rule. • Sessions ignored—Number of TCP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of TCP sessions dropped because the number of TCP sessions exceeded the limit. • Sessions dropped due to high rate—Number of TCP sessions dropped because the number of TCP connections per second exceeded the limit. • Packets allowed—Number of TCP packets that the IDS rule allowed. • Packets dropped due to high pps—Number of TCP packets dropped because the number of TCP packets per second exceeded the limit.
Egress UDP Counters	<p>Session-limit UDP counters in the egress direction for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of UDP sessions allowed by the IDS rule. • Sessions ignored—Number of UDP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of UDP sessions dropped because the number of UDP sessions exceeded the limit. • Sessions dropped due to high rate—Number of UDP sessions dropped because the number of UDP connections per second exceeded the limit. • Packets allowed—Number of UDP packets that the IDS rule allowed. • Packets dropped due to high pps—Number of UDP packets dropped because the number of TCP packets per second exceeded the limit.

Table 90: show services service-sets statistics ids session-limits counters Output Fields (*continued*)

Field Name	Field Description
Egress ICMP Counters	<p>Session-limit ICMP counters in the egress direction for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of ICMP sessions allowed by the IDS rule. • Sessions ignored—Number of ICMP sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of ICMP sessions dropped because the number of ICMP sessions exceeded the limit. • Sessions dropped due to high rate—Number of ICMP sessions dropped because the number of ICMP connections per second exceeded the limit. • Packets allowed—Number of ICMP packets that the IDS rule allowed. • Packets dropped due to high pps—Number of ICMP packets dropped because the number of ICMP packets per second exceeded the limit.
Egress Other-Protocols Counters	<p>Session-limit counters in the egress direction for protocols other than TCP, UDP, and ICMP for the following:</p> <ul style="list-style-type: none"> • Sessions allowed—Number of sessions allowed by the IDS rule. • Sessions ignored—Number of sessions that did not undergo IDS processing because traffic matched a stateful firewall rule that included accept skip-ids. • Sessions dropped due to maximum reached—Number of sessions dropped because the number of sessions exceeded the limit. • Sessions dropped due to high rate—Number of sessions dropped because the number of connections per second exceeded the limit. • Packets allowed—Number of packets that the IDS rule allowed. • Packets dropped due to high pps—Number of packets dropped because the number of packets per second exceeded the limit.

Sample Output

```
show services service-sets statistic screen-session-limit-counters
```

```
user@host> show services service-sets statistic screen-session-limit-counters
```

```
IDS Option Name: option-1
```

```
-----
```

TCP Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets allowed: 0
Packets dropped due to high pps: 0

UDP Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets allowed: 0
Packets dropped due to high pps: 0

ICMP Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets allowed: 0
Packets dropped due to high pps: 0

Other-Protocols Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets allowed: 0
Packets dropped due to high pps: 0

IDS Option Name: option-2

TCP Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets allowed: 0
Packets dropped due to high pps: 0

UDP Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets dropped due to high pps: 0

ICMP Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets allowed: 0
Packets dropped due to high pps: 0

Other-Protocols Counters:

Sessions allowed: 0
Sessions ignored: 0
Sessions dropped due to maximum reached: 0
Sessions dropped due to high rate: 0
Packets allowed: 0
Packets dropped due to high pps: 0 Destination session limit 0

show services service-sets statistics integrity-drops

Syntax

```
show services service-sets statistics integrity-drops
<interface interface-name>
<service-set service-set-name>
```

Release Information

Command introduced in Junos OS Release 13.1

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display integrity-drops statistics for one adaptive services interface, for all adaptive services interfaces, or for one service-set. You can configure use the output of this command to verify the packet header for anomalies in IP, TCP, UDP, and IGMP information and to examine any anomalies and errors.

Options

none—Display integrity-drops statistics for all configured adaptive service interfaces/ service-set.

service-set *service-set-name* —(Optional) Display integrity-drops statistics for the specified service-set

interface *interface-name*—(Optional) Display integrity-drops statistics for the specified adaptive services interface.

Required Privilege Level

view

RELATED DOCUMENTATION

| *clear services service-sets statistics integrity-drops*

List of Sample Output

[show services service-sets statistics integrity-drops on page 948](#)

Output Fields

[Table 91 on page 944](#) lists the output fields for the **show services service-sets integrity-drops** command. Output fields are listed in the approximate order in which they appear.

Table 91: show services service-sets integrity-drops Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.
Errors	<p>Total errors, categorized by protocol:</p> <ul style="list-style-type: none"> • IP—Total IP version 4 errors. • TCP—Total Transmission Control Protocol (TCP) errors. • UDP—Total User Datagram Protocol (UDP) errors. • ICMP—Total Internet Control Message Protocol (ICMP) errors.

Table 91: show services service-sets integrity-drops Output Fields (continued)

Field Name	Field Description
IP Errors	

Table 91: show services service-sets integrity-drops Output Fields (*continued*)

Field Name	Field Description
	<p>IPv4 errors:</p> <ul style="list-style-type: none"> • IP packet length inconsistencies—IP packet length does not match the Layer 2 reported length. • Minimum IP header length check failures—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes. • Reassembled packet exceeds maximum IP length—After fragment reassembly, the reassembled IP packet length exceeds 65,535. • Illegal source address 0—Source address is not a valid address. Invalid addresses are, loopback, broadcast, multicast, and reserved addresses. Source address 0, however, is allowed to support BOOTP and the destination address 0xffffffff. • Illegal destination address —Destination address is not a valid address. The address is reserved. • TTL zero errors—Received packet had a time-to-live (TTL) value of 0. • Illegal IP protocol number 0 or 255—IP protocol is 0 or 255. • Land attack—IP source address is the same as the destination address. • Non-IP packets—Packet did not conform to the IP standard. • IP option—Packet dropped because of a nonallowed IP option. • Non-IPv4 packets—Packet was not of the IPv4 type. • Non-IPv6 packets—Packet was not of the IPv6 type. • Bad checksum—Packet had an invalid IP checksum. • Illegal IP fragment length—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • IP fragment overlap—Fragments have overlapping fragment offsets. • IP fragment limit exceeded: —Fragments dropped because the configured number of allowed fragments for a packet was exceeded. • IP fragment reassembly timeout—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. Whenever a

Table 91: show services service-sets integrity-drops Output Fields (*continued*)

Field Name	Field Description
	<p>fragment is received, it is maintained in a chain until all other fragments are received. If other fragments do not arrive within the configured value of reassembly-timeout, this packet is dropped and the value of the counter shown in this field is incremented. If other fragments arrive in time but the total number of fragments is more than the configured value of fragment-limit, all the fragments (of this packet) are dropped and the value of the counter shown in this field is incremented.</p> <ul style="list-style-type: none"> ● Unknown: —Unknown fragments.
TCP Errors	<p>TCP protocol errors:</p> <ul style="list-style-type: none"> ● TCP header length inconsistencies—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes. ● Source or destination port number is zero—TCP source or destination port is zero. ● Illegal sequence number, flags combination—Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set.
UDP Errors	<p>UDP protocol errors:</p> <ul style="list-style-type: none"> ● IP data length less than minimum UDP header length (8 bytes)—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes. ● Source or destination port is zero—UDP source or destination port is 0.
ICMP Errors	<p>ICMP protocol errors:</p> <ul style="list-style-type: none"> ● IP data length less than minimum ICMP header length (8 bytes)—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes. ● ICMP error length inconsistencies—Minimum length of an ICMP error packet is 48 bytes, and the maximum length is 576 bytes. This counter is incremented when the received ICMP error falls outside this range.

Sample Output

show services service-sets statistics integrity-drops

user@host> **show services service-sets statistics integrity-drops**

```
Interface: ms-1/0/0
Service set: sset1
Errors:
  IP: 0, TCP: 0
  UDP: 0, ICMP: 0
IP errors:
  IP packet length inconsistencies: 0
  Illegal source address: 0
  Illegal destination address: 0
  TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
  Land attack: 0
  Non-IPv4 packets: 0
  Non-IPv6 packets: 0
  Bad checksum: 0
  Illegal IP fragment length: 0
  IP fragment overlap: 0
  IP fragment limit exceeded: 0
  IP fragment reassembly timeout: 0
  Unknown: 0
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combinations: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
  Source or destination port number is zero: 0
ICMP errors:
  IP data length less than minimum ICMP header length (8 bytes): 0
  ICMP error length inconsistencies: 0
```

show services service-sets statistics packet-drops

Syntax

```
show services service-sets statistics packet-drops
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 7.4.
 Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display the number of dropped packets for service sets exceeding CPU limits or memory limits.

Options

- none**—Display the number of dropped service sets packets for all adaptive services interfaces.
- interface *interface-name***—(Optional) Display the number of dropped service sets packets for a particular interface. On M Series and T Series routers, *interface-name* can be *ms-fpc/pic/port*, *sp-fpc/pic/port*, or *rspnumber*.

Required Privilege Level

view

RELATED DOCUMENTATION

| [clear services flow-collector statistics](#)

List of Sample Output

[show services service-sets statistics packet-drops on page 950](#)

Output Fields

[Table 91 on page 944](#) lists the output fields for the **show services service-sets packet-drops** command. Output fields are listed in the approximate order in which they appear.

Table 92: show services service-sets packet-drops Output Fields

Field Name	Field Description
<i>Interface</i>	Name of an adaptive services interface.
<i>Service set</i>	Name of a service set.

Table 92: show services service-sets packet-drops Output Fields (*continued*)

Field Name	Field Description
<i>CPU limit Drops</i>	Number of packets dropped because the service set exceeded the average CPU limit.
<i>Memory limit Drops</i>	Number of packets dropped because the service set exceeded the memory limit.
<i>Flow limit Drops</i>	Number of packets dropped because the service set exceeded the flow limit.

Sample Output

show services service-sets statistics packet-drops

user@host> **show services service-sets statistics packet-drops**

```
Interface: vms-1/0/0
Service set: ssl
  CPU limit drops: 0
  Memory limit drops: 0
  Flow limit drops: 0
```


show services service-sets statistics syslog

Syntax

```
show services service-sets statistics syslog
<interface interface-name>
<service-set service-set-name>
<brief | detail>
```

Release Information

Command introduced in Junos OS Release 11.1.

Support for this command introduced in Junos OS Release 19.3R2 for Next Gen Services with the MX-SPC3 services card on MX240, MX480 and MX960 routers.

Description

Display the system log statistics with optional filtering by interface and service set name.

Options

none—Display the system log statistics for all services interfaces and all service sets.

brief—(Default) (Optional) Display abbreviated system log statistics.

detail—(Optional) Display detailed system log statistics.

interface *interface-name*—(Optional) Display the system log statistics for a specific adaptive service interface.
On M Series and T Series routers, *interface-name* can be **ms-fpc/pic/port**, **sp-fpc/pic/port**, or **rspnumber**.

service-set *service-set-name*—(Optional) Display the system log statistics for a specific named service-set.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear services service-sets statistics syslog](#)

List of Sample Output

[show services service-sets statistics syslog brief on page 955](#)

[show services service-sets statistics syslog detail on page 956](#)

Output Fields

[Table 93 on page 952](#) lists the output fields for the **show services service-sets statistics syslog** command. Output fields are listed in the approximate order in which they appear.

Table 93: show services service-sets statistics syslog Output Fields

Field Name	Field Description	Level
Interface	Name of a services interface.	all
Rate limit	Maximum number of messages per second written to the interface's system log.	all
Sent	Number of messages sent that are not associated with a service set.	all
Dropped	Number of messages dropped that are not associated with a service set.	all
Service-set		
Service-set	Name of a service set.	all
Sent	Number of sent messages that are associated with the service set.	all
Dropped	Number of dropped messages that are associated with the service set.	all
Session open logs	<p>The following information is displayed for system log messages for session open events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail

Table 93: show services service-sets statistics syslog Output Fields (*continued*)

Field Name	Field Description	Level
Session close logs	<p>The following information is displayed for system log messages for session close events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail
Packet logs	<p>The following information is displayed for system log messages for packet events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail
Stateful firewall logs	<p>The following information is displayed for system log messages for stateful firewall events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail

Table 93: show services service-sets statistics syslog Output Fields (*continued*)

Field Name	Field Description	Level
ALG logs	<p>The following information is displayed for system log messages for ALG events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail
NAT logs	<p>The following information is displayed for system log messages for NAT events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail
IDS logs	<p>The following information is displayed for system log messages for IDS events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail

Table 93: show services service-sets statistics syslog Output Fields (*continued*)

Field Name	Field Description	Level
Other logs	<p>The following information is displayed for system log messages for other types of events that are logged and are associated with the service set:</p> <ul style="list-style-type: none"> • Sent—Number of messages sent. • Dropped—Number of messages dropped. Counts are given for these drop reasons: <ul style="list-style-type: none"> • low priority—Priority of the message was too low for the message to be sent. • no class set—Specific classes of event messages were configured and this class was not selected. • above rate limit—Maximum number of system log messages per second was exceeded. 	detail

Sample Output

show services service-sets statistics syslog brief

user@host> show services service-sets statistics syslog brief

```

Interface: sp-1/1/0
  Rate limit: 200000
  Sent: 0
  Dropped: 0
  Service-set: sset-sfw-sp1
    Sent: 20
    Dropped: 3488
  Service-set: sset-nat-sp1
    Sent: 18
    Dropped: 91
Interface: sp-1/2/0
  Rate limit: 15000
  Sent: 0
  Dropped: 0
  Service-set: sset-sfw-sp2
    Sent: 210
    Dropped: 579

```

Sample Output

show services service-sets statistics syslog detail

user@host> **show services service-sets statistics syslog detail**

```
Interface: ms-2/1/0
  Rate limit: 0
  Sent: 0
  Dropped: 0
  Service-set: sset1
    Sent: 0
    Dropped: 0
    Session open logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    Session close logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    Packet logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    Stateful firewall logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    ALG logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    NAT logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    IDS logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    PCP MAP logs:
      Sent: 0
      Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
```

```

limit: 0)
    PCP protocol logs:
        Sent: 0
        Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    PCP protocol error logs:
Sent: 0
    Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    PCP debug logs:
        Sent: 0
        Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)
    Other logs:
        Sent: 0
        Dropped: 0 (low priority: 0, none severity: 0, no class set: 0, above rate
limit: 0)

```

For Next Gen Services MX-SPC3 Services Card

Following shows the output for the **show services service-sets statistics syslog** on the MX-SPC3 services cards **vms-x/y/z** interfaces.

```
user@host> show services service-sets statistics syslog
```

```

show services service-sets statistics syslog
Log Module Statistics
Interface-Name- vms-2/0/0
Service-set Name- Sset1
Name              Generated      Discarded
-----
UTM                0              0
FW_AUTH            0              0
SCREEN             0              0
ALG                0              0
NAT                0              0
FLOW               0              0
SCTP               0              0
GTP                0              0
IPSEC              0              0

```

IDP	0	0
RTLOG	0	0
PST_DS_LITE	0	0
APPQOS	0	0
SECINTEL	0	0
AAMW	0	0
OTHERS	0	0
Log stream Statistics		
Interface-Name- vms-2/0/0		
Service-set Name- Sset1		
Name	send	Fail

database	0	0

show services service-sets statistics tcp

Syntax

```
show services service-sets statistics tcp
<interface interface-name>
<service-set service-set-name>
```

Release Information

Command introduced in Junos OS Release 17.2.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display TCP-related statistics.

Options

interface *interface-name*—Name of adaptive services interface.

service-set *service-set-name*—Name of service set.

Required Privilege Level

view

RELATED DOCUMENTATION

| [Configuring TFO](#)

List of Sample Output

[show services service-sets statistics tcp on page 959](#)

Output Fields

Sample Output

```
show services service-sets statistics tcp
```

```
user@host> show services service-sets statistics tcp
```

```
Interface:vms-0/2/0
  Service set: ssl_interface_style1
```

```
TCP open/close statistics:
  TCP first packet non-syn: 1
  TCP first packet reset: 0
  TCP first packet FIN: 0
  TCP non syn discard: 0
  TCP extension alloc fail: 0
  TFO SYN with cookie request: 0
  TFO SYN with cookie: 0
  TFO SYN ACK with cookie: 0
  TFO packets forwarded: 0
  TFO packets dropped: 0
  TFO packets stripped: 0
  TCP invalid syn ack: 0
  TCP invalid ack window check: 0
  TCP invalid syn transmit: 0
  TCP invalid reset in listen: 0
  TCP invalid reset in syn received: 0
  TCP invalid reset in syn sent: 0
  TCP invalid flags handshake: 0
TCP MSS statistics:
  TCP SYN MSS Received: 0
  TCP SYN MSS Modified: 0
```

show services service-sets summary

Syntax

```
show services service-sets summary
<interface interface-name>
```

Release Information

Command introduced before Junos OS Release 7.4.
 Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display service set summary information.

Options

- none**—Display service set summary information for all adaptive services interfaces.
- interface *interface-name***—(Optional) Display service set summary information for a particular interface.
 On M Series and T Series routers, *interface-name* can be *ms-fpc/pic/port*, *sp-fpc/pic/port*, or *rspnumber*.

 On MX Series MX240, MX480, and MX960 routers, *interface-name* can be *vms-fpc/pic/port* for the MX-SPC3 services card for Next Gen Services.

Required Privilege Level

view

List of Sample Output

- [show services service-sets summary on page 962](#)
- [show services service-sets summary interface on page 962](#)

Output Fields

[Table 94 on page 961](#) lists the output fields for the **show services service-sets summary** command. Output fields are listed in the approximate order in which they appear.

Table 94: show services service-sets summary Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface
Service type	Type of adaptive service, such as stateful firewall (SFW), Network Address Translation (NAT), intrusion detection service (IDS), Layer 2 Tunneling Protocol (L2TP), Compressed Real-Time Transport Protocol (CRTP), or IP Security (IPsec)

Table 94: show services service-sets summary Output Fields (*continued*)

Field Name	Field Description
Service sets configured	Total number of service sets configured on the PIC that use internal service set IDs and do not consume external service sets, including CRTP and L2TP
Bytes used	Bytes used by a particular service or all services
Policy bytes used	Policy bytes used by a particular service or all services
CPU utilization	Percentage of the CPU resources being used

Sample Output

show services service-sets summary

user@host> **show services service-sets summary**

Service sets				
Interface	CPU configured	Bytes used	Session bytes used	Policy
bytes used	utilization			
vms-3/0/0	1	3453621040 (24.93%)	0 (0.00%)	8161168
(0.90%)	0.14 %			

show services service-sets summary interface

user@host> **show services service-sets summary interface sp-1/3/0**

Interface: sp-1/3/0				
Service sets				CPU
Service type	configured	Bytes used		utilization
SFW/NAT/IDS	1	54 (0.00 %)		N/A
L2TP	1	58 (0.00 %)		N/A
CRTP	1	58 (0.00 %)		N/A
System	0	920831 (0.44 %)		N/A
Idle	0	0 (0.00 %)		N/A
Total	3	921001 (0.44 %)		N/A

show services sessions (Next Gen Services)

Syntax

```
show services sessions
<brief | extensive | terse>
<application-protocol protocol>
<count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
<utilization>
```

Release Information

Command introduced in Junos OS Release 19.3R2 on MX Series for Next Gen Services for CGNAT 6rd softwires running inline on the MPC card and specifying the si-1/0/0 interface naming convention. Support added in Junos OS Release 20.2R1 for Next Gen Services CGNAT DS-Lite softwires on the MX-SPC3 security services card .

Description

Display session information.

NOTE: On MX Series routers (with interchassis redundancy configured), the idle timeout for every flow is displayed in the **show services session extensive** and **show services flows extensive** commands.

Options

none—Display standard information about all sessions.

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

application-protocol *protocol*—(Optional) Display information about one of the following application protocols:

- **bootp**—Bootstrap protocols
- **dce-rpc**—Distributed Computing Environment-Remote Procedure Call protocols

- **dce-rpc-portmap**—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—Domain Name System protocol
- **exec**—Remote Execution Protocol
- **ftp**—File Transfer Protocol
- **h323**—H.323
- **icmp**—ICMP
- **icmpv6**—ICMPv6
- **iiop**—Internet Inter-ORB Protocol
- **ike-esp-nat**—IKE ALG
- **ip**—IP
- **login**—LOGIN
- **netbios**—NETBIOS
- **netshow**—NETSHOW
- **pptp**—Point-to-Point Tunneling Protocol
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **rsh**—Remote Shell
- **sip**—Session Initiation Protocol
- **shell**—Shell
- **snmp**—SNMP
- **sql**—SQLNet
- **talk**—Talk Program
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

NOTE: You can use the **none** option with the **show services sessions count application-protocol** command to display information about sessions other than ALG sessions.

count—(Optional) Display a count of the matching entries.

destination-port *destination-port*—(Optional) Display information for the specified destination port. The range of values is from 0 to 65,535.

destination-prefix *destination-prefix*—(Optional) Display information for the specified destination prefix.

interface *interface-name*—(Optional) Display information about the specified services interface.

limit *number*—(Optional) Maximum number of entries to display.

protocol *protocol*—(Optional) Display information about one of the following IP types:

- ***number***—Numeric protocol value from 0 to 255
- ***ah***—IPsec Authentication Header protocol
- ***egp***—An exterior gateway protocol
- ***esp***—IPsec Encapsulating Security Payload protocol
- ***gre***—A generic routing encapsulation protocol
- ***icmp***—Internet Control Message Protocol
- ***icmp6***—Internet Control Message Protocol version 6
- ***igmp***—Internet Group Management Protocol
- ***ipip***—IP-within-IP Encapsulation Protocol
- ***ospf***—Open Shortest Path First protocol
- ***pim***—Protocol Independent Multicast protocol
- ***rsvp***—Resource Reservation Protocol
- ***sctp***—Stream Control Transmission Protocol
- ***tcp***—Transmission Control Protocol
- ***udp***—User Datagram Protocol

service-set *service-set*—(Optional) Display information for the specified service set.

source-port *source-port*—(Optional) Display information for the specified source port. The range of values is from 0 to 65,535.

source-prefix *source-prefix*—(Optional) Display information for the specified source prefix.

utilization—(Optional) Display statistical details about session utilization.

Required Privilege Level

view

List of Sample Output

[show services sessions on page 967](#)

[show services sessions brief on page 968](#)

[show services sessions extensive on page 968](#)

[show services sessions terse on page 968](#)

[show services sessions analysis on page 969](#)

[show services sessions application-protocol on page 971](#)

[show services sessions count on page 974](#)

[show services sessions destination-port on page 974](#)

[show services sessions destination-prefix on page 975](#)

[show services sessions interface on page 975](#)

[show services sessions protocol on page 975](#)

[show services sessions service-set on page 975](#)

[show services sessions source-port on page 976](#)

[show services sessions source-prefix on page 976](#)

Output Fields

[Table 95 on page 966](#) lists the output fields for the **show services sessions** command. Output fields are listed in the approximate order in which they appear.

Table 95: show services sessions Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the services interface.	application-protocol
Session	Session ID that uniquely identifies the session.	All levels
ALG	Name of the application.	terse
Flags	Session flag for the ALG: <ul style="list-style-type: none"> • 0x1—Found an existing session. • 0x2—Reached session or flow limit. • 0x3—No memory available for new sessions. • 0x4—No free session ID available. • 0x0000—No session ID found. 	All levels

Table 95: show services sessions Output Fields (*continued*)

Field Name	Field Description	Level of Output
IP Action	Flag indicating whether IP action has been set for the session.	All levels
Offload	Flag indicating whether the session has been offloaded to the Packet Forwarding Engine.	All levels
Asymmetric	Flag indicating whether the session is uni-directional.	terse application-protocol
Service set	Name of a service set. Individual empty service sets are not displayed.	count
Sessions Count	Number of sessions.	count

Sample Output

show services sessions

user@host> show services sessions

```

Session ID: 536870913, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 26, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1:ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1:ipip, Conn Tag: 0x0, If:
vms-2/0/0.0, Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 26, Valid
Logical system: root-logical-system
  Software          2002:2010::1401:4          -> 2002:2010::1401:1
  In: 30.1.0.101/1024 --> 30.2.0.101/1024:udp, Conn Tag: 0x0, If: vms-2/0/0.16391,
Pkts: 1, Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024:udp, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,
Total sessions: 2

```

show services sessions brief

The output for the **show services flows brief** command is identical to that for the **show services sessions** command. For sample output, see [show services sessions on page 967](#).

show services sessions extensive

user@host> **show services sessions extensive**

```

Session ID: 536870917, Service-set: vms-sset10, Status: Normal
Flags: 0x40/0x0/0x4000/0x2000103
Policy name: default-service-set-policy/32779
Source NAT pool: Null, Destination NAT pool: Null
Dynamic application: junos:UNKNOWN,
Encryption: Unknown
Application traffic control rule-set: INVALID, Rule: INVALID
Maximum timeout: 30, Current timeout: 28
Session State: Valid
Logical system: root-logical-system
Start time: 1878, Duration: 2
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip,
  Conn Tag: 0x0, Interface: vms-2/0/0.16391,
    Session token: 0xfcc, Flag: 0x400023
    Route: 0x0, Gateway: 2002:2010::1401:4, Tunnel ID: 0, Tunnel type: None
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 1, Bytes: 110
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip,
  Conn Tag: 0x0, Interface: vms-2/0/0.0,
    Session token: 0x4fcc, Flag: 0x400022
    Route: 0x0, Gateway: 2002:2010::1401:1, Tunnel ID: 0, Tunnel type: None
    Port sequence: 0, FIN sequence: 0,
    FIN state: 0,
    Pkts: 0, Bytes: 0
Total sessions: 1

```

show services sessions terse

user@router> **show services sessions terse**

```

vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP          10.2.2.2:52138 ->      10.1.1.2:21      Forward I          33
TCP          10.1.1.2:21      ->      10.2.2.2:52138 Forward O          31

```

show services sessions analysis

user@router>**show services sessions analysis**

```
vms-1/0/0
  Interface:    vms-1/0/0

Session Analysis Statistics:

Total sessions Active           :0
Total TCP Sessions Active       :0
  Tcp sessions from gate       :0
  Tunneled TCP sessions        :0
  Regular TCP sessions         :0
  IPv4 active Session          :0
  IPv6 active Session          :0
Total UDP sessions Active       :0
  UDP sessions from gate       :0
  Tunneled UDP sessions        :0
  Regular UDP sessions         :0
  IPv4 active Session          :0
  IPv6 active Session          :0
Total Other sessions Active     :0
  IPv4 active Session          :0
  IPv6 active Session          :0
Created sessions per Second     :0
Deleted sessions per Second     :0
Peak Total sessions Active      :0
Peak Total TCP sessions Active  :0
Peak Total UDP sessions Active  :0
Peak Total Other sessions Active :0
Peak Created Sessions per Second :0
Peak Deleted Sessions per Second :0
Packets received                :0
Packets transmitted             :0
Slow path forward               :0
Slow path discard               :0

Session Rate Data:
  Number of Samples: 638051

Session Rate Distribution(sec)

Session Operation :Creation
```

```

400000+          :0
350001 - 400000 :0
300001 - 350000 :0
250001 - 300000 :0
200001 - 250000 :0
150001 - 200000 :0
50001  - 150000 :0
40001  - 50000  :0
30001  - 40000  :0
20001  - 30000  :0
10001  - 20000  :0
1001   - 10000  :0
1      - 1000   :0
          0      :638051

```

Session Operation :Deletion

```

400000+          :0
350001 - 400000 :0
300001 - 350000 :0
250001 - 300000 :0
200001 - 250000 :0
150001 - 200000 :0
50001  - 150000 :0
40001  - 50000  :0
30001  - 40000  :0
20001  - 30000  :0
10001  - 20000  :0
1001   - 10000  :0
1      - 1000   :0
          0      :638051

```

Session Lifetime Distribution(sec):

	TCP	UDP	HTTP
240+	:0	0	0
120 - 240	:0	0	0
60 - 120	:0	0	0
30 - 60	:0	0	0
15 - 30	:0	0	0
5 - 15	:0	0	0

1 - 5	:0	0	0
0 - 1	:0	0	0

show services sessions application-protocol

This command has the same output for the rpc, dce-rpc, rpc-portmap and dce-rpc-portmap ALGs.

user@router> **show services sessions application-protocol dce-rpc**

```
Interface name: vms-1/1/0
Session: 8, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:1019 ->192.168.203.194:2049 Forward I      4
UDP    192.168.203.194:2049 ->192.168.203.198:1019 Forward O      4
Session: 7, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:954  ->192.168.203.194:613 Forward I      1
UDP    192.168.203.194:613  ->192.168.203.198:954 Forward O      1
Session: 6, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:53836 ->192.168.203.194:613 Forward I      1
UDP    192.168.203.194:613  ->192.168.203.198:53836 Forward O      1
Session: 5, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:59813 ->192.168.203.194:111 Forward I      1
UDP    192.168.203.194:111  ->192.168.203.198:59813 Forward O      1
Session: 4, ALG: portmapper, Flags: 0x1800, IP Action: no, Offload: no
UDP    192.168.203.198:36595 ->192.168.203.194:2049 Forward I      1
UDP    192.168.203.194:2049 ->192.168.203.198:36595 Forward O      1
Session: 3, ALG: portmapper, Flags: 0x1000, IP Action: no, Offload: no
UDP    192.168.203.198:56050 ->192.168.203.194:111 Forward I      1
UDP    192.168.203.194:111  ->192.168.203.198:56050 Forward O      1
```

user@router> **show services sessions application-protocol dns**

```
Interface name: vms-2/0/0
Session: 293, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    198.51.100.2:43677 -> 203.0.113.10:53 Forward I      1
UDP    203.0.113.10:53    -> 192.0.2.1:43677 Forward O      1
Session: 53, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    198.51.100.2:37494 -> 203.0.113.10:53 Forward I      1
UDP    203.0.113.10:53    -> 192.0.2.1:37494 Forward O      1
Session: 66, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    198.51.100.2:48161 -> 203.0.113.10:53 Forward I      1
UDP    203.0.113.10:53    -> 192.0.2.1:48161 Forward O      1
Session: 17, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    198.51.100.2:38908 -> 203.0.113.10:53 Forward I      1
```

```

UDP    203.0.113.10:53    ->      192.0.2.1:38908 Forward  O      1
Session: 42, ALG: 16, Flags: 0x0040, IP Action: no, Offload: no
UDP    198.51.100.2:58189 ->      203.0.113.10:53    Forward  I      1
UDP    203.0.113.10:53    ->      192.0.2.1:58189 Forward  O      1

```

user@router> **show services sessions application-protocol ftp**

```

Interface name: vms-4/1/0
Session: 1, ALG: 1, Flags: 0x0040, IP Action: no, Offload: no
TCP    192.0.2.129:32843 ->      198.51.100.129:21    Forward  I      26
TCP    198.51.100.129:21  ->      192.0.2.0:32843 Forward  O      30

```

user@router> **show services sessions application-protocol ike-esp-nat**

```

Service Set: ss_ipv4, Session: 33554435, ALG: ike-esp-nat, Flags: 0x0800, IP Action:
no, Offload: no, Asymmetric: no
ESP 198.51.100.2:4689 ->      203.0.113.1:62108 Forward  O 2199
ESP 192.0.2.2:62108 ->      198.51.100.2:4689 Forward  I 0
Service Set: ss_ipv4, Session: 33554434, ALG: ike-esp-nat, Flags: 0x0800, IP Action:
no, Offload: no, Asymmetric: no
ESP 192.0.2.2:44179 ->      198.51.100.2:43809 Forward  I 2199
ESP 198.51.100.2:43809 ->      203.0.113.1:44179 Forward  O 0
Service Set: ss_ipv4, Session: 33554433, ALG: ike-esp-nat, Flags: 0x0000, IP Action:
no, Offload: no, Asymmetric: no
UDP 192.0.2.2:500 ->      198.51.100.2:500 Forward  I 8
UDP 198.51.100.2:500 ->      203.0.113.1:57730 Forward  O

```

user@router> **show services sessions application-protocol pptp**

```

Interface name: vms-2/0/0
Session: 3, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE    203.0.113.138:0    ->      203.0.113.138:0      Forward  O
21
GRE    192.0.2.794:0      ->      203.0.113.138:0:65000 Forward  I
0
Session: 2, ALG: pptp, Flags: 0x2800, IP Action: no, Offload: no, Asymmetric: no
GRE    192.0.2.794:0      ->      203.0.113.138:0:49913 Forward  I
88
GRE    203.0.113.138:0:49913 ->      192.0.2.794:65001 Forward  O
0
Session: 1, ALG: pptp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP    192.0.2.794:1511 ->      203.0.113.138:0:1723 Forward  I

```

```

13
TCP      203.0.113.138:0:1723  ->    192.0.2.794:1511  Forward  O
12

```

user@router> **show services sessions application-protocol rtsp**

```

Interface name: vms-0/1/0
Session: 13, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
UDP      203.0.113.66:5004  ->    198.51.100.66:3989  Forward  O      152
UDP      198.51.100.66:3989  ->    192.0.2.161:5004  Forward  I      0
Session: 9, ALG: rtsp, Flags: 0x0800, IP Action: no, Offload: no
UDP      203.0.113.66:5004  ->    198.51.100.66:3986  Forward  O      3
UDP      198.51.100.66:3986  ->    192.0.2.161:5004  Forward  I      0

```

user@router> **show services sessions application-protocol rsh**

```

Interface name: vms-2/0/0
Session: 3, ALG: 2, Flags: 0x0840, IP Action: no, Offload: no
TCP      203.0.113.10:1023  ->    198.51.100.2:1020  Forward  O      4
TCP      198.51.100.2:1020  ->    203.0.113.10:1023  Forward  I      3
Session: 1, ALG: 2, Flags: 0x0040, IP Action: no, Offload: no
TCP      198.51.100.2:1021  ->    203.0.113.10:514   Forward  I     1331
TCP      203.0.113.10:514   ->    198.51.100.2:1021  Forward  O     2485

```

user@router> **show services sessions application-protocol sip**

```

Interface name: vms-2/0/0
Session: 4, ALG: sip, Flags: 0x0800, IP Action: no, Offload: no
UDP      198.51.100.130:6000  ->    192.0.2.129:12682  Forward  I
      246
UDP      192.0.2.129:12682  ->    198.51.100.162:6000  Forward  O
      0
Session: 1, ALG: sip, Flags: 0x0000, IP Action: no, Offload: no
UDP      198.51.100.130:5060  ->    192.0.2.130:5060   Forward  I
      10
UDP      192.0.2.130:5060   ->    198.51.100.162:5060  Forward  O
      9

```

user@router> **show services sessions application-protocol sql**

```
Interface name: vms-2/0/0
Session: 3934, ALG: sqlnet, Flags: 0x0800, IP Action: no, Offload: no
TCP      198.51.100.2:39754 ->      203.0.113.138:0:1408 Forward I      26
TCP      203.0.113.138:0:1408 ->      192.0.2.1:39754 Forward O      23
```

user@router> **show services sessions application-protocol talk**

```
Interface name: vms-0/2/0
Session: 4, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
TCP      203.0.113.162:36888 ->      192.0.2.2:33294 Forward O
      4
TCP      192.0.2.1:33294 ->      203.0.113.162:36888 Forward I
      3
Session: 7, ALG: 65, Flags: 0x0800, IP Action: no, Offload: no
UDP      203.0.113.162:1165 ->      192.0.2.2:518 Forward O
      1
UDP      192.0.2.2:518 ->      203.0.113.162:1165 Forward I
      1
Session: 8, ALG: 65, Flags: 0x0000, IP Action: no, Offload: no
UDP      192.0.2.2:1509 ->      203.0.113.162:518 Forward I
      3
UDP      203.0.113.162:518 ->      192.0.2.2:1509 Forward O
      3
Session: 6, ALG: 0, Flags: 0x0000, IP Action: no, Offload: no
UDP      192.0.2.1:123 ->      192.0.2.2:123 Forward O
      4
```

show services sessions count

user@host> **show services sessions count**

Interface	Service set	Valid	Invalid	Pending	Other
vms-0/2/0	ssl_interface_style1		1	0	0
0					

show services sessions destination-port

user@router> **show services sessions destination-port 21**

```
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
```



```
TCP          10.2.2.2:52138 ->      10.1.1.2:21    Forward  I          25
TCP          10.1.1.2:21    ->      10.2.2.2:52138 Forward  O          24
```

show services sessions destination-prefix

user@router> show services sessions destination-prefix 10.1.1.2

```
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP          10.2.2.2:52138 ->      10.1.1.2:21    Forward  I          25
TCP          10.1.1.2:21    ->      10.2.2.2:52138 Forward  O          24
```

show services sessions interface

user@router> show services sessions interface vms-1/1/0

```
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP          10.2.2.2:52138 ->      10.1.1.2:21    Forward  I          30
TCP          10.1.1.2:21    ->      10.2.2.2:52138 Forward  O          29
```

show services sessions protocol

user@router> show services sessions protocol tcp

```
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP          10.2.2.2:52138 ->      10.1.1.2:21    Forward  I          30
TCP          10.1.1.2:21    ->      10.2.2.2:52138 Forward  O          29
```

show services sessions service-set

user@router> show services sessions service-set ss1_interface_style1

```
Session ID: 3, Service-set: ss1_interface_style1, Policy name: R11/7, Timeout: 30,
Valid
  In: 20.1.1.2/48102 --> 30.1.1.2/22;tcp, Conn Tag: 0x0, If: vms-0/2/0.16387, Pkts:
70, Bytes: 6257,
  Out: 30.1.1.2/22 --> 44.0.0.3/29071;tcp, Conn Tag: 0x0, If: vms-0/2/0.0, Pkts:
59, Bytes: 8193,
Total sessions: 1
```

show services sessions source-port

```
user@router> show services sessions source-port 21
```

```
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP          10.2.2.2:52138 ->      10.1.1.2:21    Forward  I          33
TCP          10.1.1.2:21    ->      10.2.2.2:52138 Forward  O          31
```

show services sessions source-prefix

```
user@router> show services sessions source-prefix 10.2.2.2
```

```
vms-1/1/0
Session: 1, ALG: ftp, Flags: 0x2000, IP Action: no, Offload: no, Asymmetric: no
TCP          10.2.2.2:52138 ->      10.1.1.2:21    Forward  I          33
TCP          10.1.1.2:21    ->      10.2.2.2:52138 Forward  O          31
```

show services sessions (Aggregated Multiservices)

Syntax

```
show services sessions
<brief | extensive | terse>
<application-protocol protocol>
<count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display the session information for each service set in each member interface of the AMS interface.

Options

none—Display standard information about all sessions.

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

application-protocol—(Optional) Display information about one of the following application protocols:

- **ftp**—File Transfer Protocol
- **icmp**—Internet Control Message Protocol
- **pptp**—Point-to-Point Tunneling Protocol
- **rtsp**—Real-Time Streaming Protocol
- **sqlnet**—SQL *Net
- **tcp**—Transmission Control Protocol
- **traceroute**—Traceroute
- **tftp**—Trivial File Transfer Protocol
- **udp**—User Datagram Protocol

count—(Optional) Display a count of the matching entries.

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is from 0 through 65,535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, *interface-name* can be **ms-fpc/pic/port** or **rspnumber**. On J Series routers, *interface-name* is **ms-pim/0/port**.

limit *number*—(Optional) Maximum number of entries to display.

protocol *protocol*—(Optional) Display information about one of the following IP types:

- **number**—Numeric protocol value from 0 through 255
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **icmp6**—Internet Control Message Protocol version 6
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-over-IP encapsulation protocol
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Transmission Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 through 65,535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level

view

List of Sample Output

[show services sessions brief on page 980](#)

[show services sessions interface mams-5/0/0 extensive on page 980](#)

[show services sessions terse on page 983](#)

[show services sessions count on page 985](#)

Output Fields

Table 95 on page 966 lists the output fields for the **show services sessions** command. Output fields are listed in the approximate order in which they appear.

Table 96: show services sessions Output Fields

Field Name	Field Description
Interface	Name of the member interface (mams-) and the aggregated multiservices interface (ams) to which it belongs.
Session ID	Session ID that uniquely identifies the session.
ALG	Name of the application.
Flags	Session flag for the ALG: <ul style="list-style-type: none"> • 0x1—Found an existing session. • 0x2—Reached session or flow limit. • 0x3—No memory available for new sessions. • 0x4—No free session ID available.
IP Action	Flag indicating whether IP action has been set for the session.
Offload	Flag indicating whether the session has been offloaded to the Packet Forwarding Engine.
Asymmetric	Flag indicating whether the session is unidirectional.
Service set	Name of a service set. Individual empty service sets are not displayed.
Sessions Count	Number of sessions.
Flow or Flow Prot	Protocol used for this session.
Source	Source prefix of the flow in the format source-prefix:port . For ICMP flows, port information is not displayed.
Dest	Destination prefix of the flow. For ICMP flows, port information is not displayed.

Table 96: show services sessions Output Fields (*continued*)

Field Name	Field Description
State	Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without response. • Forward—Forward the packet in the flow without looking at it. • Reject—Drop all packets in the flow with response. • Watch—Inspect packets in the flow. • Bypass—Bypass packets in the flow. • Unknown—Unknown flow status.
Packet Direction	Direction of the flow: ingress (I), egress (O), or unknown.
Frm count	Number of frames in the flow.

Sample Output

show services sessions brief

```
user@host> show services sessions brief
```

```
mams-1/0/0 (ams0)
Service Set: napt_set, Session: 16777217, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
UDP      30.30.30.2:63    ->    40.40.40.2:63    Forward I      85689
UDP      40.40.40.2:63    ->    30.30.30.160:6000 Forward O      0
```

show services sessions interface mams-5/0/0 extensive

```
user@host> show services sessions interface mams-5/0/0 extensive
```

```
mams-1/0/0 (ams0)
Service Set: napt_set, Session: 16777235, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
NAT PPlugin Data:
  NAT Action: Translation Type - NAPT-44
  NAT source 30.30.30.62:63 -> 30.30.30.176:6003
UDP      30.30.30.62:63    ->    40.40.40.62:63    Forward I      1805
Byte count: 83030
Flow role: Initiator, Timeout: 0
```

```

UDP      40.40.40.62:63    ->    30.30.30.176:6003  Forward  O              0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

NAT PPlugin Data:
  NAT Action:    Translation Type - NAPT-44
  NAT source     30.30.30.57:63      ->    30.30.30.163:6003
UDP      30.30.30.57:63    ->    40.40.40.57:63    Forward  I              1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP      40.40.40.57:63    ->    30.30.30.163:6003  Forward  O              0
  Byte count: 0
  Flow role: Responder, Timeout: 0

[...output truncated...]
mams-1/1/0 (ams0)
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

NAT PPlugin Data:
  NAT Action:    Translation Type - NAPT-44
  NAT source     30.30.30.63:63      ->    30.30.30.165:6004
UDP      30.30.30.63:63    ->    40.40.40.63:63    Forward  I              1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP      40.40.40.63:63    ->    30.30.30.165:6004  Forward  O              0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

NAT PPlugin Data:
  NAT Action:    Translation Type - NAPT-44
  NAT source     30.30.30.60:63      ->    30.30.30.164:6004
UDP      30.30.30.60:63    ->    40.40.40.60:63    Forward  I              1805
  Byte count: 83030
  Flow role: Initiator, Timeout: 0
UDP      40.40.40.60:63    ->    30.30.30.164:6004  Forward  O              0
  Byte count: 0
  Flow role: Responder, Timeout: 0
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

```

[...output truncated...]

mams-5/0/0 (ams0)

Service Set: napt_set, Session: 16777225, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

NAT PPlugin Data:

NAT Action: Translation Type - NAPT-44
 NAT source 30.30.30.64:63 -> 30.30.30.168:6002
 UDP 30.30.30.64:63 -> 40.40.40.64:63 Forward I 1805
 Byte count: 83030
 Flow role: Initiator, Timeout: 0
 UDP 40.40.40.64:63 -> 30.30.30.168:6002 Forward O 0
 Byte count: 0
 Flow role: Responder, Timeout: 0

Service Set: napt_set, Session: 16777224, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

NAT PPlugin Data:

NAT Action: Translation Type - NAPT-44
 NAT source 30.30.30.56:63 -> 30.30.30.171:6001
 UDP 30.30.30.56:63 -> 40.40.40.56:63 Forward I 1805
 Byte count: 83030
 Flow role: Initiator, Timeout: 0
 UDP 40.40.40.56:63 -> 30.30.30.171:6001 Forward O 0
 Byte count: 0
 Flow role: Responder, Timeout: 0

Service Set: napt_set, Session: 16777223, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

[...output truncated...]

mams-5/1/0 (ams0)

Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no

NAT PPlugin Data:

NAT Action: Translation Type - NAPT-44
 NAT source 30.30.30.61:63 -> 30.30.30.172:6004
 UDP 30.30.30.61:63 -> 40.40.40.61:63 Forward I 1805
 Byte count: 83030
 Flow role: Initiator, Timeout: 0
 UDP 40.40.40.61:63 -> 30.30.30.172:6004 Forward O 0
 Byte count: 0
 Flow role: Responder, Timeout: 0


```
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
```

```
NAT PPlugin Data:
```

```
  NAT Action: Translation Type - NAPT-44
```

```
    NAT source      30.30.30.52:63      ->      30.30.30.175:6003
```

```
UDP      30.30.30.52:63      ->      40.40.40.52:63      Forward  I      1805
```

```
  Byte count: 83030
```

```
  Flow role: Initiator, Timeout: 0
```

```
UDP      40.40.40.52:63      ->      30.30.30.175:6003  Forward  O      0
```

```
  Byte count: 0
```

```
  Flow role: Responder, Timeout: 0
```

```
Service Set: napt_set, Session: 16777231, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
```

```
[...output truncated...]
```

show services sessions terse

```
user@router> show services sessions terse
```

```
mams-1/0/0 (ams0)
```

```
Service Set: napt_set, Session: 16777235, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
```

```
UDP      30.30.30.62:63      ->      40.40.40.62:63      Forward  I      2541
```

```
UDP      40.40.40.62:63      ->      30.30.30.176:6003  Forward  O      0
```

```
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
```

```
UDP      30.30.30.57:63      ->      40.40.40.57:63      Forward  I      2541
```

```
UDP      40.40.40.57:63      ->      30.30.30.163:6003  Forward  O      0
```

```
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
```

```
UDP      30.30.30.50:63      ->      40.40.40.50:63      Forward  I      2541
```

```
UDP      40.40.40.50:63      ->      30.30.30.162:6003  Forward  O      0
```

```
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
```

```
UDP      30.30.30.48:63      ->      40.40.40.48:63      Forward  I      2541
```

```
UDP      40.40.40.48:63      ->      30.30.30.161:6003  Forward  O      0
```

```
[...output truncated...]
```

```
mams-1/1/0 (ams0)
```

```
Service Set: napt_set, Session: 16777234, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
```

```
UDP      30.30.30.63:63      ->      40.40.40.63:63      Forward  I      2543
```

```
UDP      40.40.40.63:63      ->      30.30.30.165:6004  Forward  O      0
```

```

Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.60:63    ->    40.40.40.60:63    Forward  I            2543
UDP      40.40.40.60:63    ->    30.30.30.164:6004 Forward  O              0
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.59:63    ->    40.40.40.59:63    Forward  I            2543
UDP      40.40.40.59:63    ->    30.30.30.167:6003 Forward  O              0
Service Set: napt_set, Session: 16777231, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.58:63    ->    40.40.40.58:63    Forward  I            2543
UDP      40.40.40.58:63    ->    30.30.30.166:6003 Forward  O              0
[...output truncated...]
mams-5/0/0 (ams0)
Service Set: napt_set, Session: 16777225, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.64:63    ->    40.40.40.64:63    Forward  I            2543
UDP      40.40.40.64:63    ->    30.30.30.168:6002 Forward  O              0
Service Set: napt_set, Session: 16777224, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.56:63    ->    40.40.40.56:63    Forward  I            2543
UDP      40.40.40.56:63    ->    30.30.30.171:6001 Forward  O              0
Service Set: napt_set, Session: 16777223, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.55:63    ->    40.40.40.55:63    Forward  I            2543
UDP      40.40.40.55:63    ->    30.30.30.170:6001 Forward  O              0
Service Set: napt_set, Session: 16777222, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.51:63    ->    40.40.40.51:63    Forward  I            2543
UDP      40.40.40.51:63    ->    30.30.30.169:6001 Forward  O              0
[...output truncated...]
mams-5/1/0 (ams0)
Service Set: napt_set, Session: 16777233, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.61:63    ->    40.40.40.61:63    Forward  I            2544
UDP      40.40.40.61:63    ->    30.30.30.172:6004 Forward  O              0
Service Set: napt_set, Session: 16777232, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.52:63    ->    40.40.40.52:63    Forward  I            2545
UDP      40.40.40.52:63    ->    30.30.30.175:6003 Forward  O              0
Service Set: napt_set, Session: 16777231, ALG: none, Flags: 0x2000, IP Action: no,
  Offload: no, Asymmetric: no
UDP      30.30.30.47:63    ->    40.40.40.47:63    Forward  I            2545
UDP      40.40.40.47:63    ->    30.30.30.174:6003 Forward  O              0

```

```

Service Set: napt_set, Session: 16777230, ALG: none, Flags: 0x2000, IP Action: no,
Offload: no, Asymmetric: no
UDP      30.30.30.46:63    ->    40.40.40.46:63    Forward  I          2545
UDP      40.40.40.46:63    ->    30.30.30.173:6003 Forward  O           0
[...output truncated...]

```

show services sessions count

```
user@host> show services sessions count
```

Interface	Service set	Sessions count
mams-1/0/0	napt_set	19
mams-1/0/0	ssl	0
mams-1/1/0	napt_set	18
mams-1/1/0	ssl	0
mams-5/0/0	napt_set	9
mams-5/0/0	ssl	0
mams-5/1/0	napt_set	17
mams-5/1/0	ssl	0

show services sessions analysis

Syntax

```
show services sessions analysis
<interface interface-name>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series MS-MPC.
Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display session statistics.

Options

- none**—Display standard information about all session statistics.
- interface *interface-name***—(Optional) Display information about the specified interface.

Required Privilege Level

view

List of Sample Output

[show services sessions analysis interface on page 988](#)

Output Fields

[Table 97 on page 986](#) lists the output fields for the **show services sessions analysis** command. Output fields are listed in the approximate order in which they appear.

Table 97: show services sessions analysis Output Fields

Field Name	Field Description
Services PIC Name	FPC and PIC slots for the services PIC on which the sessions are running.
Session Analysis Statistics:	
Total Sessions Active	Total active sessions in the MS-PIC including TCP, UDP, ICMP and Softwires.
Total TCP Sessions Active	Total active TCP sessions in the MS-PIC.

Table 97: show services sessions analysis Output Fields (*continued*)

Field Name	Field Description
Total UDP Sessions Active	Total active UDP session in the MS-PIC.
Total Other Sessions Active	Total other active sessions in the MS-PIC including ICMP and softwires.
Total Predicted Sessions Active	Predicted sessions are created only by the ALG traffic using the L3/L4 information available.
Created Sessions per Second	Session setup rate at the time of running the command.
Deleted Sessions per Second	Session deletion rate at the time of running the command.
Peak Total Sessions Active	Highest number of active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total TCP Sessions Active	Highest number of active TCP sessions since the last PIC restart or since the last time session stats are flushed.
Peak Total UDP Sessions Active	Highest number of active UDP sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total Other Sessions Active	Highest number of other active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Created Sessions per Second	Maximum session setup rate observed since the last PIC restart or since the last time session statistics are flushed.
Peak Deleted Sessions per Second	Maximum session deletion rate observed since the last PIC restart or from the last time session statistics are flushed.
Packets received	Total number of packets received by the MS-PIC.
Packets transmitted	Total number of packets transmitted by the MS-PIC.
Slow path forward	Number of packets forwarded in the slow path (that is, after the successful rule match and session creation).
Slow path discard	Number of packets discarded before the session creation.

Table 97: show services sessions analysis Output Fields (*continued*)

Field Name	Field Description
Session Rate Data: Number of Samples	Number of samples used to calculate the session rate since the last PIC restart or since the last time session statistics are flushed.
Session Rate Distribution(sec)	
Session Operation :Creation	Number of sampling intervals during which a number of sessions in the indicated range were created during the current sampling period.
Session Operation :Deletion	Number of sampling intervals during which a number of sessions in the indicated range were deleted during the current sampling period.
Session Lifetime Distribution(sec):	Number of TCP, UDP, and HTTP sessions whose length was in the indicated range in seconds.

Sample Output

show services sessions analysis interface

user@host> **show services sessions analysis interface ms-5/1/0**

```

Services PIC Name:      ms-5/1/0

Session Analysis Statistics:

Total sessions Active           :0
Total TCP Sessions Active       :0
  Tcp sessions from gate        :0
  Tunneled TCP sessions         :0
  Regular TCP sessions          :0
  IPv4 active Session           :0
  IPv6 active Session           :0
Total UDP sessions Active       :0
  UDP sessions from gate        :0
  Tunneled UDP sessions         :0
  Regular UDP sessions          :0
  IPv4 active Session           :0
  IPv6 active Session           :0
Total Other sessions Active     :0
  IPv4 active Session           :0

```

```

IPv6 active Session           :0
Created sessions per Second   :0
Deleted sessions per Second   :0
Peak Total sessions Active    :0
Peak Total TCP sessions Active :0
Peak Total UDP sessions Active :0
Peak Total Other sessions Active :0
Peak Created Sessions per Second :0
Peak Deleted Sessions per Second :0
Packets received              :0
Packets transmitted           :0
Slow path forward             :0
Slow path discard             :0

```

Session Rate Data:

Number of Samples: 3518

Session Rate Distribution(sec)

Session Operation :Creation

```

400000+           :0
350001 - 400000   :0
300001 - 350000   :0
250001 - 300000   :0
200001 - 250000   :0
150001 - 200000   :0
50001  - 150000   :0
40001  - 50000    :0
30001  - 40000    :0
20001  - 30000    :0
10001  - 20000    :0
1001   - 10000    :0
1       - 1000     :0
          0       :3518

```

Session Operation :Deletion

```

400000+           :0
350001 - 400000   :0
300001 - 350000   :0
250001 - 300000   :0
200001 - 250000   :0
150001 - 200000   :0

```

```
50001 - 150000 :0
40001 - 50000 :0
30001 - 40000 :0
20001 - 30000 :0
10001 - 20000 :0
1001 - 10000 :0
1 - 1000 :0
0 :3518
```

Session Lifetime Distribution(sec):

	TCP	UDP	HTTP
240+	:0	0	0
120 - 240	:0	0	0
60 - 120	:0	0	0
30 - 60	:0	0	0
15 - 30	:0	0	0
5 - 15	:0	0	0
1 - 5	:0	0	0
0 - 1	:0	0	0

show services sessions analysis (USF)

Syntax

```
show services sessions analysis
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 19.3R2 on MX Series for USF.

Description

Display session statistics.

Options

none—Display standard information about all session statistics.

interface *interface-name*—(Optional) Display information about the specified services interface.

Required Privilege Level

view

List of Sample Output

[show services sessions analysis interface on page 993](#)

Output Fields

[Table 97 on page 986](#) lists the output fields for the **show services sessions analysis** command. Output fields are listed in the approximate order in which they appear.

Table 98: show services sessions analysis Output Fields

Field Name	Field Description
Services PIC Name	FPC and PIC slots for the services PIC on which the sessions are running.
Session Analysis Statistics:	
Total Sessions Active	Total active sessions in the services PIC, including TCP, UDP, ICMP and Softwires.
Total TCP Sessions Active	Total active TCP sessions in the services PIC.
Total UDP Sessions Active	Total active UDP session in the services PIC.

Table 98: show services sessions analysis Output Fields (*continued*)

Field Name	Field Description
Total Other Sessions Active	Total other active sessions in the services PIC, including ICMP and softwires.
Total Predicted Sessions Active	Predicted sessions are created only by the ALG traffic using the L3/L4 information available.
Created Sessions per Second	Session setup rate at the time of running the command.
Deleted Sessions per Second	Session deletion rate at the time of running the command.
Peak Total Sessions Active	Highest number of active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total TCP Sessions Active	Highest number of active TCP sessions since the last PIC restart or since the last time session stats are flushed.
Peak Total UDP Sessions Active	Highest number of active UDP sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Total Other Sessions Active	Highest number of other active sessions since the last PIC restart or since the last time session statistics are flushed.
Peak Created Sessions per Second	Maximum session setup rate observed since the last PIC restart or since the last time session statistics are flushed.
Peak Deleted Sessions per Second	Maximum session deletion rate observed since the last PIC restart or from the last time session statistics are flushed.
Packets received	Total number of packets received by the services PIC.
Packets transmitted	Total number of packets transmitted by the services PIC.
Slow path forward	Number of packets forwarded in the slow path (that is, after the successful rule match and session creation).
Slow path discard	Number of packets discarded before the session creation.
Session Rate Data: Number of Samples	Number of samples used to calculate the session rate since the last PIC restart or since the last time session statistics are flushed.

Table 98: show services sessions analysis Output Fields (continued)

Field Name	Field Description
Session Rate Distribution(sec)	
Session Operation :Creation	Number of sampling intervals during which a number of sessions in the indicated range were created during the current sampling period.
Session Operation:Deletion	Number of sampling intervals during which a number of sessions in the indicated range were deleted during the current sampling period.
Session Lifetime Distribution(sec):	Number of TCP, UDP, and HTTP sessions whose length was in the indicated range in seconds.

Sample Output

show services sessions analysis interface

user@host> show services sessions analysis interface vms-5/1/0

```

Services PIC Name:      vms-5/1/0

Session Analysis Statistics:

Total sessions Active           :0
Total TCP Sessions Active       :0
  Tcp sessions from gate       :0
  Tunneled TCP sessions        :0
  Regular TCP sessions         :0
  IPv4 active Session          :0
  IPv6 active Session          :0
Total UDP sessions Active       :0
  UDP sessions from gate       :0
  Tunneled UDP sessions        :0
  Regular UDP sessions         :0
  IPv4 active Session          :0
  IPv6 active Session          :0
Total Other sessions Active     :0
  IPv4 active Session          :0
  IPv6 active Session          :0
Created sessions per Second     :0
Deleted sessions per Second     :0

```

```

Peak Total sessions Active           :0
Peak Total TCP sessions Active       :0
Peak Total UDP sessions Active       :0
Peak Total Other sessions Active     :0
Peak Created Sessions per Second     :0
Peak Deleted Sessions per Second     :0
Packets received                     :0
Packets transmitted                   :0
Slow path forward                     :0
Slow path discard                     :0

```

Session Rate Data:

Number of Samples: 3518

Session Rate Distribution(sec)

Session Operation :Creation

```

400000+           :0
350001 - 400000   :0
300001 - 350000   :0
250001 - 300000   :0
200001 - 250000   :0
150001 - 200000   :0
50001  - 150000   :0
40001  - 50000    :0
30001  - 40000    :0
20001  - 30000    :0
10001  - 20000    :0
1001   - 10000    :0
1       - 1000     :0
0       :3518

```

Session Operation :Deletion

```

400000+           :0
350001 - 400000   :0
300001 - 350000   :0
250001 - 300000   :0
200001 - 250000   :0
150001 - 200000   :0
50001  - 150000   :0
40001  - 50000    :0
30001  - 40000    :0

```

```
20001 - 30000 :0
10001 - 20000 :0
1001 - 10000 :0
1 - 1000 :0
0 :3518
```

Session Lifetime Distribution(sec):

	TCP	UDP	HTTP
240+	:0	0	0
120 - 240	:0	0	0
60 - 120	:0	0	0
30 - 60	:0	0	0
15 - 30	:0	0	0
5 - 15	:0	0	0
1 - 5	:0	0	0
0 - 1	:0	0	0

show services sessions count

Syntax

```
show services sessions count
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display the count of matching entries.

Required Privilege Level

view

List of Sample Output

[show services sessions count on page 996](#)

Output Fields

Sample Output

show services sessions count

user@host> show services sessions count

Interface	Service set	Valid	Invalid	Pending	Other state
vms-0/2/0	ssl_interface_style1	1	0	0	0

show services sessions service-set

Syntax

```
show services sessions service-set service-set
```

Release Information

Command introduced in Junos OS release 19.3R2.

Description

Display table session entries for the specified service set.

Required Privilege Level

view

List of Sample Output

[show services sessions service-set on page 997](#)

Output Fields

Sample Output

```
show services sessions service-set
```

```
user@host> show services sessions service-set ss1_interface_style1
```

```
Session ID: 3, Service-set: ss1_interface_style1, Policy name: R11/7, Timeout: 30,
Valid
  In: 20.1.1.2/48102 --> 30.1.1.2/22;tcp, Conn Tag: 0x0, If: vms-0/2/0.16387, Pkts:
70, Bytes: 6257,
  Out: 30.1.1.2/22 --> 44.0.0.3/29071;tcp, Conn Tag: 0x0, If: vms-0/2/0.0, Pkts:
59, Bytes: 8193,
Total sessions: 1
```

show services sessions service-set

Syntax

```
show services sessions service-set
```

Release Information

Command introduced in Junos OS Release 19.3R2 for Next Gen Services on MX Series.

Description

Display the open and close sessions for a service-set.

Required Privilege Level

show services sessions service-set

```
user@host> show services sessions service-set service-set-name
```

```
Session ID: 268436944, Policy name: self-traffic-policy/1, Timeout: 554, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12253 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12253;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,

Session ID: 268436945, Policy name: self-traffic-policy/1, Timeout: 554, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12254 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12254;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,

Session ID: 268436946, Policy name: self-traffic-policy/1, Timeout: 596, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12255 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12255;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 1,
Bytes: 44,

Session ID: 268436947, Policy name: self-traffic-policy/1, Timeout: 554, Valid
Logical system: root-logical-system
```



```
In: 5.5.5.1/12256 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12256;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,

Session ID: 268436948, Policy name: self-traffic-policy/1, Timeout: 596, Valid
Logical system: root-logical-system
  In: 5.5.5.1/12257 --> 70.0.0.2/514;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 2,
Bytes: 84,
  Out: 70.0.0.2/514 --> 5.5.5.1/12257;tcp, Conn Tag: 0x0, If: .local..6, Pkts: 1,
Bytes: 44,
Total sessions: 5
```

show services sessions software

Syntax

```
show services sessions software  
interfaces interface-name
```

Release Information

Command introduced in Junos OS Release 20.2R1.

Description

Display session information for softwires.

Options

count —Display statistics and information on the number of softwires.

ds-lite—Display information about DS-Lite softwires.

Required Privilege Level

view

List of Sample Output

[show services sessions software on page 1000](#)

[show services sessions software count on page 1001](#)

[show services sessions software ds-lite on page 1001](#)

[show services sessions software ds-lite count on page 1002](#)

[show services sessions software ds-lite aftr on page 1002](#)

[show services sessions software ds-lite b4 on page 1003](#)

[show services sessions software ds-lite b4 <ip address> aftr <ip-address> on page 1003](#)

[Show services sessions software flow-details on page 1004](#)

Output Fields

Sample Output

show services sessions software

user@host> **show services sessions software**

```
Session ID: 536870913, Service-set: vms-sset10, Policy name:  
default-service-set-policy/32779, Timeout: 26, Valid  
Logical system: root-logical-system
```

```
In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.0, Pkts: 0, Bytes: 0,
```

```
Session ID: 536870914, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 26, Valid
Logical system: root-logical-system
```

```
    Software      2002:2010::1401:4      -> 2002:2010::1401:1
In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391,
Pkts: 1, Bytes: 70,
Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,
Total sessions: 2
```

show services sessions software count

show services sessions software count

user@host> **show services sessions software count**

Interface	Service set	Valid	Invalid	Pending
Other state				
vms-2/0/0	vms-sset10	1	0	0
0				
vms-2/0/0	vms-sset11			

show services sessions software ds-lite

show services sessions software ds-lite

user@host> **show services sessions software ds-lite**

```
Session ID: 536870913, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 26, Valid
Logical system: root-logical-system
```

```
In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
```

```
Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.0, Pkts: 0, Bytes: 0,
```

```
Session ID: 536870914, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 26, Valid
```

```
Logical system: root-logical-system
```

```
Software 2002:2010::1401:4 -> 2002:2010::1401:1
```

```
In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391,
Pkts: 1, Bytes: 70,
```

```
Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,
```

```
Total sessions: 2
```

show services sessions software ds-lite count

show services sessions software ds-lite count

```
user@host> show services sessions software ds-lite count
```

Interface	Service set	Valid	Invalid	Pending
Other state				
vms-2/0/0	vms-sset10	1	0	0
0				
vms-2/0/0	vms-sset11			

show services sessions software ds-lite aftr

show services sessions software ds-lite aftr

```
user@host> show services sessions software ds-lite aftr
```

```
Session ID: 536870913, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 6, Valid
```

```
Logical system: root-logical-system
```

```
In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
```

```
Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.0, Pkts: 0, Bytes: 0,
```

```

Session ID: 536870914, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 6, Valid
Logical system: root-logical-system
    Software      2002:2010::1401:4      -> 2002:2010::1401:1
    In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391,
    Pkts: 1, Bytes: 70,
    Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0,
    Pkts: 0, Bytes: 0,
Total sessions: 2

```

show services sessions software ds-lite b4

show services sessions software ds-lite b4

user@host> show services sessions software ds-lite b4

```

Session ID: 536870913, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 6, Valid
Logical system: root-logical-system
    In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
    vms-2/0/0.16391, Pkts: 1, Bytes: 110,
    Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If:
    vms-2/0/0.0, Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 6, Valid
Logical system: root-logical-system
    Software      2002:2010::1401:4      -> 2002:2010::1401:1
    In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391,
    Pkts: 1, Bytes: 70,
    Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0,
    Pkts: 0, Bytes: 0,
Total sessions: 2

```

show services sessions software ds-lite b4 <ip-address> aftr <ip-address>

show services sessions software ds-lite b4 <ip address> aftr <ip-address>

user@host> show services sessions software ds-lite b4 ip address aftr ip-address

```
Session ID: 536870913, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 6, Valid
Logical system: root-logical-system
  In: DSLITE 2002:2010::1401:4/1 --> 2002:2010::1401:1/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.16391, Pkts: 1, Bytes: 110,
  Out: DSLITE 2002:2010::1401:1/1 --> 2002:2010::1401:4/1;ipip, Conn Tag: 0x0, If:
vms-2/0/0.0, Pkts: 0, Bytes: 0,

Session ID: 536870914, Service-set: vms-sset10, Policy name:
default-service-set-policy/32779, Timeout: 6, Valid
Logical system: root-logical-system
  Softwire          2002:2010::1401:4          -> 2002:2010::1401:1
  In: 30.1.0.101/1024 --> 30.2.0.101/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.16391,
Pkts: 1, Bytes: 70,
  Out: 30.2.0.101/1024 --> 50.0.12.1/1024;udp, Conn Tag: 0x0, If: vms-2/0/0.0,
Pkts: 0, Bytes: 0,
Total sessions: 2
```

Show services sessions softwire flow-details

Show services sessions softwire flow-details

```
user@host> show services sessions softwire flow-details
```

Interface: vms-2/0/0, Service set: vms-sset10			
Softwire		Direction	Flow
count			
2002:2010::1401:4->2002:2010::1401:1	In	1	

show services sessions utilization

Syntax

```
show services sessions utilization
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 19.3R2.

Description

Display session utilization statistics.

Options

interface *interface-name*—Display session utilization statistics specific to the interface.

Required Privilege Level

view

List of Sample Output

[show services sessions utilization on page 1005](#)

Output Fields

Sample Output

show services sessions utilization

user@host> show services sessions utilization

	Session	%Memory	%Session-Memory	Setup	%Rate	Drop	Teardown
%CPU							
Interface	Count			Rate		Rate	Rate
vms-3/0/0	0	24.96	0.00	0			0
0.13	Green						

show services software

Syntax

```
show services software
```

Release Information

Command introduced in Junos OS Release 10.4.

count option added in Junos OS Release 11.2.

Support added for Next Gen Services in Junos OS Release 20.2 on the MX-SPC3 security services card.

Description

Display information about software services. Information is displayed on both 6rd and DS-Lite services.

Options

count *interface-name*—(Optional) Display the current software counts for a service set for both DS-Lite and 6rd.

count — (Optional) Display the number of created softwares.

Required Privilege Level

view

List of Sample Output

[show services software on page 1007](#)

[show services software count \(sp- interfaces\) on page 1007](#)

[show services softwares count \(vms- interfaces\) on page 1007](#)

Output Fields

[Table 99 on page 1006](#) lists the output fields for the **command-name** command. Output fields are listed in the approximate order in which they appear.

Table 99: show-services-software Output Fields

Field Name	Field Description	Level of Output
Interface	Interface for which information is displayed.	All levels
Service Set	Service set containing the software rules for the interface.	All levels
Software	Name of the software concentrator.	All levels
Direction	Direction of the flow.	All levels
Flow count	Number of flows.	All levels

Sample Output

show services software

user@host> show services software

Interface: sp-3/0/0, Service set: v6rd-dom1-dom3-service-set				
Software		Direction		Flow count
10.10.10.2	->	192.0.2.1	I	13

show services software count (sp- interfaces)

user@host> show services software count

Interface	Service set	DS-Lite	6RD
sp-0/0/0	dslite-svc-set1	2	0

show services softwires count (vms- interfaces)

user@host> show services software count

Interface	Service set	DS-Lite	6RD	MAPE
vms-2/0/0	vms-sset10	1	0	

show services software flows

Syntax

```
show services software flows
(<interface interface-name> <service-set service-set-name>|
count <interface interface-name> <service-set service-set-name>|
ds-lite <B4 b4-address> <AFTR aftr-address>|
v6rd <initiator initiator-ip-address><concentrator concentrator-ip-address>)
```

Release Information

Command introduced in Junos OS Release 10.2.

Support added for Next Gen Services in Junos OS Release 20.2

Description

Display statistics information about the software flows.

NOTE: Starting with Junos OS Release 14.1R4, the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions (**dslite-ipv6-prefix-length** attribute) is taken into account while the session count is calculated and displayed in the output of the show services software flows command. Until Junos OS Release 14.1R3, only IPv4 flows were counted and IPv6 flows were not considered for the statistics about software flows

Options

interface *interface-name*—(Optional) Display statistics information about the specified interface only.

service-set *service-set-name*—(Optional) Display statistics information about the specified service set only.

count <interface *interface-name*> <service-set *service-set-name*>|—(Optional) Display flow count information only, with optional filtering by interface and service set.

ds-lite <B4 *b4-address*> <AFTR *aftr-address*>|—(Optional) Display DS-Lite flow information, with optional filtering by B4 (software initiator) and AFTR (software concentrator).

v6rd <initiator *initiator-ip-address*><concentrator *concentrator-ip-address*>|—(Optional) Display v6rd flow information, with optional filtering by the software initiator and software concentrator.

Required Privilege Level

view

List of Sample Output

[show services software flows on page 1010](#)

[show services software flows count on page 1010](#)

[show services software flows ds-lite B4 on page 1010](#)

[show services software flows ds-lite AFTR on page 1011](#)

[services software flows ds-lite AFTR and B4 on page 1011](#)

[show services softwares software-types map-e on page 1011](#)

Output Fields

[Table 100 on page 1009](#) lists the output fields for the **show services software flows** command. Output fields are listed in the approximate order in which they appear.

Table 100: show services software flows Output Fields

Field Name	Field Description
Interface	Name of the interface.
Service set	Name of the service set.
Flow	Description of flow, including protocol input and output interface addresses.
State	Flow state. Value is: <ul style="list-style-type: none"> • Forward
Dir	Flow direction. Values are: <ul style="list-style-type: none"> • I—inbound • O—outbound
Frm count	Number of frames transferred.
NAT dest	NAT translation of the decapsulated address.
Software	For outbound flows, the address of the local software initiator (B4 for DS-Lite) is shown first, followed by the address of the software concentrator (AFTR for DS-Lite). For inbound flows, the address of the software concentrator is shown first, followed by the address of the software initiator.

Sample Output

show services software flows

```
user@host> show services software flows
```

```
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                     State      Dir      Frm count
TCP      200.200.200.2:80  ->  33.33.33.1:1066 Forward  O      2005418
  NAT dest      33.33.33.1:1066  ->  20.20.1.2:1025
  Software      1001::1      ->  2001::2
TCP      20.20.1.2:1025  ->  200.200.200.2:80 Forward  I      2007168
  NAT source      20.20.1.2:1025  ->  33.33.33.1:1066
  Software      2001::2      ->  1001::1
TCP      20.20.1.2:1025  ->  200.200.200.2:80 Forward  I      2635998
  NAT source      20.20.1.2:1025  ->  33.33.33.1:1065
  Software      2001::3      ->  1001::1
DS-LITE      2001::2      ->  1001::1 Forward  I      2008157
TCP      200.200.200.2:80  ->  33.33.33.1:1065 Forward  O      2637909
  NAT dest      33.33.33.1:1065  ->  20.20.1.2:1025
  Software      1001::1      ->  2001::3
DS-LITE      2001::3      ->  1001::1 Forward  I      2640499
```

show services software flows count

```
user@host> show services software flows count
```

```
Interface      Service set      Flow count
sp-0/0/0      dslite-svc-set1      6
```

show services software flows ds-lite B4

```
user@host> show services software flows ds-lite B4 2001::2
```

```
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                     State      Dir      Frm count
TCP      200.200.200.2:80  ->  33.33.33.1:1066 Forward  O      2884037
  NAT dest      33.33.33.1:1066  ->  20.20.1.2:1025
  Software      1001::1      ->  2001::2
TCP      20.20.1.2:1025  ->  200.200.200.2:80 Forward  I      2885884
  NAT source      20.20.1.2:1025  ->  33.33.33.1:1066
  Software      2001::2      ->  1001::1
DS-LITE      2001::2      ->  1001::1 Forward  I      2886821
```

show services software flows ds-lite AFTR

```
user@host> show services software flows ds-lite AFTR 1001::1
```

```
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                     State      Dir      Frm count
TCP      200.200.200.2:80    ->    33.33.33.1:1066 Forward  O      3359356
    NAT dest      33.33.33.1:1066    ->    20.20.1.2:1025
    Software      1001::1      ->    2001::2
TCP      20.20.1.2:1025    ->    200.200.200.2:80 Forward  I      3361235
    NAT source      20.20.1.2:1025    ->    33.33.33.1:1066
    Software      2001::2      ->    1001::1
TCP      20.20.1.2:1025    ->    200.200.200.2:80 Forward  I      4479810
    NAT source      20.20.1.2:1025    ->    33.33.33.1:1065
    Software      2001::3      ->    1001::1
DS-LITE      2001::2      ->    1001::1 Forward  I      3362168
TCP      200.200.200.2:80    ->    33.33.33.1:1065 Forward  O      4481520
    NAT dest      33.33.33.1:1065    ->    20.20.1.2:1025
    Software      1001::1      ->    2001::3
DS-LITE      2001::3      ->    1001::1 Forward  I      4484094
```

services software flows ds-lite AFTR and B4

```
user@host> show services software flows ds-lite AFTR 1001::1 B4 2001::2
```

```
Interface: sp-0/0/0, Service set: dslite-svc-set1
Flow                                     State      Dir      Frm count
TCP      200.200.200.2:80    ->    33.33.33.1:1066 Forward  O      3931026
    NAT dest      33.33.33.1:1066    ->    20.20.1.2:1025
    Software      1001::1      ->    2001::2
TCP      20.20.1.2:1025    ->    200.200.200.2:80 Forward  I      3932792
    NAT source      20.20.1.2:1025    ->    33.33.33.1:1066
    Software      2001::2      ->    1001::1
DS-LITE      2001::2      ->    1001::1 Forward  I      3933782
```

show services softwares software-types map-e

```
user@host> show services softwares software-types map-e mape-tun1
```

```
br-address 2001:db8:ffff::1/128; //Mandatory
rule r1 {
    ipv4-prefix 192.0.2.0/24; //Mandatory
    ipv6-prefix 2001:db8:0000::/40; //Mandatory
    ea-bits-length 16; //Mandatory
```

```
    psid-offset 4; //Mandatory
    psid-len 8;
}
version 3;
```

show services software statistics

Syntax

```
show services software statistics
<ds-lite>
<ds-lite>
<interface interface-name>
<v6rd>
```

Release Information

Command introduced in Junos OS Release 10.4.
 Support for Next Gen Services with the MX-SPC3 security services card added in Junos OS Release 20.2.

Description

Display information about software services.

Options

- ds-lite**—(Optional) Display only DS-Lite.
- interface *interface-name*** —(Optional) Name of the interface servicing the software. When you omit this option, data for all interfaces are shown.
- v6rd**—(Optional) Display only 6rd statistics.

Required Privilege Level

view

List of Sample Output

- [show services software statistics \(sp- interfaces\) on page 1017](#)
- [show services software statistics ds-lite \(sp- interfaces\) on page 1019](#)
- [show services software statistics \(vms- interfaces\) on page 1021](#)
- [show services software statistics ds-lite \(vms- interfaces\) on page 1022](#)

Output Fields

[Table 101 on page 1013](#) lists the output fields for the **command-name** command. Output fields are listed in the approximate order in which they appear.

Table 101: command-name Output Fields

Field Name	Field Description	Level of Output
Service PIC Name	Name of service PIC for which statistics are shown.	statistics

Table 101: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
Softwires Created	Number of softwires created.	statistics
Softwires Created for EIF/HP	Number of softwires created for endpoint-independent filtering (EIF) or hairpinning (HP).	statistics for ds-lite only
Softwires Deleted	Number of softwires deleted.	statistics
Softwires Flows Created	Number of flows created.	statistics
Softwires Flows Deleted	Number of flows deleted.	statistics
Slow Path Packets Processed	Number of packets processed as initial packets in a softwire session. These packets require a rule lookup and setting up of flows; this processing of an initial packet in a flow is called <i>the slow path</i> .	statistics
Slow Path Packets Processed for EIF/HP	Number of slow path EIF/HP packets processed.	statistics for ds-lite only
Fast Path Packets Processed	Number of packets processed that are not <i>slow path</i> .	statistics
Fast Path Encapsulated	Number of packets encapsulated in the fast path.	statistics
Softwire EIF Accept	Number of packets that matched an EIF entry that initiated the creation of a DS-Lite tunnel. The EIF entry was previously triggered by a DS-Lite packet.	statistics for ds-lite only
Rule Match Succeeded	Number of packets that matched a softwire rule.	statistics
Rule Match Failed	Number of packets that did not match any softwire rule.	statistics

Table 101: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
IPv6 Packets Fragmented	Number of packets fragmented by the services PIC.	statistics for ds-lite only
IPv4 Client Fragments	Number of IPv4 fragments received from the client end over the softwire tunnel destined to the server.	statistics for ds-lite only
IPv4 Server First Fragments	Number of IPv4 first fragments received from the server destined to go over the softwire tunnel to the client.	statistics for ds-lite only
IPv4 Server More Fragments	Number of IPv4 other fragments (excluding first and last fragment) received from the server destined to go over the softwire tunnel to the client.	statistics for ds-lite only
IPv4 Server Last Fragments	Number of IPv4 last fragments received from the server destined to go over the softwire tunnel to the client.	statistics for ds-lite only
ICMPv4 Packets sent	Number of ICMPv4 packets sent to the softwire concentrator.	statistics
ICMPv4 Error Packets sent	Number of ICMPv4 error packets sent to the softwire concentrator.	statistics
ICMPv6 Packets sent	Number of ICMPv6 packets sent to the softwire concentrator.	statistics
Dropped ICMPv6 packets destined to AFTR	Number of ICMPv6 packets dropped instead of sending to the softwire concentrator.	statistics
Softwire Creation Failed	Number of softwire creation failures.	statistics for ds-lite and 6rd
Softwire Creation Failed for EIF/HP	Number of softwire creation failures for EIF/HP.	statistics for ds-lite only
Flow Creation Failed	Number of flow creation failures.	statistics
Flow Creation Failed for EIF/HP	Number of flow creation failures for EIF/HP.	statistics for ds-lite only

Table 101: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
Flow Creation Failed - Retry	Number of flow creations retried after failure.	statistics
Slow Path Failed	Number of failures detected in the slow path.	statistics
Slow Path Failed - Retry	Number of times processing of a packet was reprocessed in the slow path.	statistics
Packet not IPv4-in-IPv6	Number of IPv4 packets not encapsulated in IPv6.	statistics for ds-lite only
IPv6 Fragmentation Error	Number of IPv6 packets with fragmentation errors.	statistics
Slow Path Failed-IPv6 Next Header Offset	Number of IPv6 header errors detected in slow path processing.	statistics for ds-lite only
Decapsulated Packet not IPv4	Number of packets without IPv4 inner header.	statistics for ds-lite only
Decap Failed - IPv6 Next Header Offset	Decapsulation failure due to an unexpected inner header.	statistics for ds-lite only
Decap Failed - IPv4 L3 Integrity	Decapsulation failure due to incorrect Layer 3 data, such as not an IP packet, bad source or destination address, checksum error, or protocol error.	statistics for ds-lite only
Decap Failed - IPv4 L4 Integrity	Decapsulation failure due to incorrect Layer 4 data, such as errors in TCP, UDP, or TCP headers.	statistics for ds-lite only
No Software ID	Number of times a software ID was not found.	statistics
No Flow Extension	Number of times flow extensions were not found.	statistics
ICMPv4 Dropped Packets	Number of ICMPv4 packets dropped.	statistics

Table 101: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
Packet not IPv6-in-IPv4	Number of IPv6 packets not encapsulated in IPv4.	statistics for v6rd only
Decapsulated Packet not IPv6	Number of packets without an IPv6 inner header.	statistics for v6rd only
Encapsulation Failed - No packet memory	Failed to encapsulate IPv6 packets in IPv4 due to low memory.	statistics for v6rd only
Flow limit exceeded	Flow not created because configured maximum flows per software is exceeded.	statistics
Session limit exceeded	Flow not created because configured maximum DS-Lite software sessions per IPv6 prefix is exceeded.	statistics for ds-lite only

Sample Output

show services software statistics (sp- interfaces)

user@host> **show services software statistics**

DS-Lite Statistics:

Service PIC Name: :sp-0/0/0

Statistics

Softwires Created	:0
Softwires Created for EIF/HP	:0
Softwires Deleted	:0
Softwires Flows Created	:0
Softwires Flows Deleted	:0
Slow Path Packets Processed	:0
SLow Path Packets Processed for EIF/HP	:0
Fast Path Packets Processed	:0
Fast Path Packets Encapsulated	:0

Softwire EIF Accept	:0
Rule Match Succeeded	:0
Rule Match Failed	:0
IPv6 Packets Fragmented	:0
IPv4 Client Fragments	:0
IPv4 Server First Fragments	:0
IPv4 Server More Fragments	:0
IPv4 Server Last Fragments	:0
ICMPv4 Packets sent	:0
ICMPv4 Error Packets sent	:0
ICMPv6 Packets sent	:0
Dropped ICMPv6 packets destined to AFTR	:0

Transient Errors

Flow Creation Failed - Retry	:0
FLOW Creation Failed - Retry for EIF/HP	:0
Slow Path Failed - Retry	:0

Errors

Softwire Creation Failed	:0
Softwire Creation Failed for EIF/HP	:0
Flow Creation Failed	:0
FLOW Creation Failed For EIF/HP	:0
Slow Path Failed	:0
Packet not IPv4-in-IPv6	:0
IPv6 Fragmentation Error	:0
Softwire Creation Failed - IPv6 Next Header Offset	:0
Decapsulated Packet not IPv4	:0
Decap Failed - IPv6 Next Header Offset	:0
Decap Failed - IPv4 L3 Integrity	:0
Decap Failed - IPv4 L4 Integrity	:0
No Softwire ID	:0
No Flow Extension	:0
Flow Limit Exceeded	:0

6rd Statistics:

Service PIC Name	:sp-0/0/0
------------------	-----------

Statistics

Softwires Created	:0
Softwires Deleted	:0
Softwires Flows Created	:0
Softwires Flows Deleted	:0
Slow Path Packets Processed	:0
Fast Path Packets Processed	:0
Fast Path Packets Encapsulated	:0
Rule Match Failed	:0
Rule Match Succeeded	:0

Transient Errors

Flow Creation Failed - Retry	:0
Slow Path Failed - Retry	:0

Errors

Softwire Creation Failed	:0
Flow Creation Failed	:0
Slow Path Failed	:0
Packet not IPv6-in-IPv4	:0
Slow Path Failed - IPv6 Next Header Offset	:0
Decapsulated Packet not IPv6	:0
Encapsulation Failed - No packet memory	:0
No Softwire ID	:0
No Flow Extension	:0
ICMPv4 Dropped Packets	:0

show services softwire statistics ds-lite (sp- interfaces)

user@host> **show services softwire statistics ds-lite**

DS-Lite Statistics:

Service PIC Name:	:sp-0/0/0
-------------------	-----------

Statistics

Softwires Created	:0
Softwires Created for EIF/HP	:0
Softwires Deleted	:0
Softwires Flows Created	:0
Softwires Flows Deleted	:0
Slow Path Packets Processed	:0
Slow Path Packets Processed for EIF/HP	:0
Fast Path Packets Processed	:0
Fast Path Packets Encapsulated	:0
Softwire EIF Accept	:0
Rule Match Succeeded	:0
Rule Match Failed	:0
IPv6 Packets Fragmented	:0
IPv4 Client Fragments	:0
IPv4 Server First Fragments	:0
IPv4 Server More Fragments	:0
IPv4 Server Last Fragments	:0
ICMPv4 Packets sent	:0
ICMPv4 Error Packets sent	:0
ICMPv6 Packets sent	:0
Dropped ICMPv6 packets destined to AFTR	:0

Transient Errors

Flow Creation Failed - Retry	:0
Flow Creation Failed - Retry for EIF/HP	:0
Slow Path Failed - Retry	:0

Errors

Softwire Creation Failed	:0
Softwire Creation Failed for EIF/HP	:0
Flow Creation Failed	:0
Flow Creation Failed For EIF/HP	:0
Slow Path Failed	:0
Packet not IPv4-in-IPv6	:0
IPv6 Fragmentation Error	:0
Softwire Creation Failed - IPv6 Next Header Offset	:0
Decapsulated Packet not IPv4	:0
Decap Failed - IPv6 Next Header Offset	:0

```

Decap Failed - IPv4 L3 Integrity      :0
Decap Failed - IPv4 L4 Integrity      :0
No Software ID                       :0
No Flow Extension                     :0
Flow Limit Exceeded                   :0
Session Limit Exceeded                :0

```

Sample Output

show services software statistics (vms- interfaces)

user@host> **show services software statistics**

```

vms-2/0/0
  Total Session Interest events      :3
  Total Session Destroy events       :2
  Total Session Public Request events :0
  Total Session Accepts              :1
  Total Session Discards             :0
  Total Session Ignores              :0
  Total Session extension alloc failures :0
  Total Session extension set failures :0
Software statistics
  Total Software sessions created     :1
  Total Software sessions deleted     :2
  Total Software sessions created for reverse packets :1
  Total Software session create failed for reverse pkts :0
  Total Software rule match success   :1
  Total Software rule match failed    :0
  Software session limit exceeded     :0
Software packet statistics
  Total Packets processed             :1
  Total packets encapsulated          :1
  Total packets decapsulated          :1
  Encapsulation errors                :0
  Decapsulation errors                :0
  Encapsulated pkts re-inject failures :0
  Decapsulated pkts re-inject failures :0
  DS-Lite ICMPv4 Echo replies sent    :0
  DS-Lite ICMPv4 TTL exceeded messages sent :0
  ICMPv6 ECHO request messages received destined to AFTR :0
  ICMPv6 ECHO reply messages sent from AFTR :0

```

```

    ICMPv6 ECHO requests to AFTR process failures           :0
    V6 untunnelled packets destined to AFTR dropped         :1
    Softwire policy add errors                               :0
    Softwire policy delete errors                           :0
    Softwire policy memory alloc failures                   :0
    Softwire Untunnelled packets ignored                    :0
    Softwire Misc errors
        DS-Lite ICMPv4 TTL exceed message process errors   :0

```

show services softwire statistics ds-lite (vms- interfaces)

user@host> show services softwire statistics ds-lite interface vms-2/0/0

```

vms-2/0/0
    Total Session Interest events           :3
    Total Session Destroy events             :2
    Total Session Public Request events      :0
    Total Session Accepts                   :1
    Total Session Discards                   :0
    Total Session Ignores                   :0
    Total Session extension alloc failures   :0
    Total Session extension set failures     :0
    Softwire statistics
        Total Softwire sessions created      :1
        Total Softwire sessions deleted      :2
        Total Softwire sessions created for reverse packets :1
        Total Softwire session create failed for reverse pkts :0
        Total Softwire rule match success    :1
        Total Softwire rule match failed     :0
        Softwire session limit exceeded      :0
    Softwire packet statistics
        Total Packets processed               :1
        Total packets encapsulated           :1
        Total packets decapsulated           :1
        Encapsulation errors                 :0
        Decapsulation errors                 :0
        Encapsulated pkts re-inject failures :0
        Decapsulated pkts re-inject failures :0
        DS-Lite ICMPv4 Echo replies sent     :0
        DS-Lite ICMPv4 TTL exceeded messages sent :0
        ICMPv6 ECHO request messages received destined to AFTR :0
        ICMPv6 ECHO reply messages sent from AFTR :0
        ICMPv6 ECHO requests to AFTR process failures :0
        V6 untunnelled packets destined to AFTR dropped :1

```


Softwire policy add errors	:0
Softwire policy delete errors	:0
Softwire policy memory alloc failures	:0
Softwire Untunnelled packets ignored	:0
Softwire Misc errors	
DS-Lite ICMPv4 TTL exceed message process errors	:0

show services stateful-firewall conversations

Syntax

```
show services stateful-firewall conversations
<brief | extensive | terse>
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<pgcp>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced before Junos OS Release 7.4.

pgcp option introduced in Junos OS Release 8.4.

Description

Display information about stateful firewall conversations.

Options

none—Display standard information about all stateful firewall conversations.

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

application-protocol *protocol*—(Optional) Display information about one of the following application protocols:

- **bootp**—Bootstrap protocol
- **dce-rpc**—Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—Domain Name System protocol
- **exec**—Exec
- **ftp**—File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol

- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port* or *rspnumber*.

limit *number*—(Optional) Maximum number of entries to display.

pgcp —(Optional) Display information about stateful firewall conversations for Packet Gateway Control Protocol (PGCP) flows.

protocol *protocol*—(Optional) Display information about one of the following IP types:

- **number**—Numeric protocol value from 0 to 255
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol

- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for the specific service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level

view

List of Sample Output

[show services stateful-firewall conversations on page 1028](#)

[show services stateful-firewall conversations destination-port on page 1028](#)

Output Fields

[Table 102 on page 1026](#) lists the output fields for the **show services stateful-firewall conversations** command. Output fields are listed in the approximate order in which they appear.

Table 102: show services stateful-firewall conversations Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set. Individual empty service sets are not displayed, but if no service set has any flows, a flow table header is printed for each service set.
Conversation	Information about a group of related flows. <ul style="list-style-type: none"> • ALG Protocol—Application-level gateway protocol. • Number of initiators—Number of flows that initiated a session. • Number of responders—Number of flows that responded in a session.

Table 102: show services stateful-firewall conversations Output Fields (*continued*)

Field Name	Field Description
Flow or Flow Prot	Protocol used for this flow.
Source	Source prefix of the flow, in the format <i>source-prefix-port</i> .
Destination	Destination prefix of the flow.
State	Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without response. • Forward—Forward the packet in the flow without looking at it. • Reject—Drop all packets in the flow with response. • Watch—Inspect packets in the flow.
Dir	Direction of the flow: input (I) or output (O).
Source NAT	Original and translated source IPv4 or IPv6 addresses are displayed if Network Address Translation (NAT) is configured on this particular flow or conversation.
Frm Count	Number of frames in the flow.
Destin NAT	Original and translated destination IPv4 or IPv6 addresses are displayed if NAT is configured on this particular flow or conversation.
Byte count	Number of bytes forwarded in the flow.
TCP established	Whether a TCP connection was established: Yes or No .
TCP window size	Negotiated TCP connection window size, in bytes.
TCP acknowledge	TCP acknowledgment sequence number.
TCP tickle	Whether TCP inquiry mode is on (enabled or disabled) and the time remaining to send the next inquiry, in seconds.
Master flow	Flow that initiated the conversation.
Tlmeout	Lifetime of the flow, in seconds.

Sample Output

show services stateful-firewall conversations

user@host> show services stateful-firewall conversations

```
Interface: sp-1/3/0, Service set: green
Conversation: ALG Protocol: any, Number of initiators: 1,
Number of responders: 1

Flow
Prot      Source                Dest                State      Dir    Frm count
TCP       10.58.255.50:33005->    10.58.255.178:23   Forward    I      13
    Source NAT      10.58.255.50:33005->    10.59.16.100:4000
    Destin NAT      10.58.255.178:23  ->                0.0.0.0:4000
Byte count:          918
TCP established, TCP window size: 65535, TCP acknowledge: 2502627025
TCP tickle enabled, 0 seconds,
Master flow, Timeout: 30 seconds
TCP       10.58.255.178:23  ->    10.59.16.100:4000 Forward    O      8
```

show services stateful-firewall conversations destination-port

user@host> show services stateful-firewall conversations destination-port 21

```
Interface: sp-0/3/0, Service set: svc_set_trust

Interface: sp-0/3/0, Service set: svc_set_untrust
Conversation: ALG protocol: ftp
    Number of initiators: 1, Number of responders: 1
Flow
TCP       10.50.10.2:2143  ->    10.50.20.2:21    Watch    O      0
TCP       10.50.20.2:21    ->    10.50.10.2:2143  Watch    I      0
TCP       10.50.20.2:21    ->    10.50.10.2:2143  Watch    I      0
```

show services stateful-firewall flow-analysis

Syntax

```
show services stateful-firewall flow-analysis
<interface interface-name>
```

Release Information

Command introduced in Junos OS Release 10.4R1.

Description

Display stateful firewall flow statistics.

Options

none—Display standard information about all stateful firewall flow statistics.

interface *interface-name*—(Optional) Display information about a particular interface.

Required Privilege Level

view

List of Sample Output

[show services stateful-firewall flow-analysis on page 1031](#)

[show services stateful-firewall flow-analysis interface sp-3/0/0 on page 1032](#)

Output Fields

[Table 97 on page 986](#) lists the output fields for the **show services stateful-firewall flow-analysis** command. Output fields are listed in the approximate order in which they appear.

Table 103: show services stateful-firewall flow-analysis Output Fields

Field Name	Field Description
Total Flows Active	Total active flows in the MS-PIC including TCP, UDP, ICMP and Softwires.
Total TCP Flows Active	Total active TCP flows in the MS-PIC.
Total UDP Flows Active	Total active UDP flows in the MS-PIC.
Total Other Flows Active	Total other active flows in the MS-PIC including ICMP and softwires.
Total Predicted Flows Active	Predicted flows are created only by the ALG traffic using the L3/L4 information available.

Table 103: show services stateful-firewall flow-analysis Output Fields (continued)

Field Name	Field Description
Created Flows per Second	Flow setup rate at the time of running the command.
Deleted Flows per Second	Flow deletion rate at the time of running the command.
Peak Total Flows Active	The highest number of active flows since the last PIC restart or since the last time flow statistics are flushed.
Peak Total TCP Flows Active	The highest number of active TCP flows since the last PIC restart or since the last time flow stats are flushed.
Peak Total UDP Flows Active	The highest number of active UDP flows since the last PIC restart or since the last time flow statistics are flushed.
Peak Total Other Flows Active	The highest number of other active flows since the last PIC restart or since the last time flow statistics are flushed.
Peak Created Flows per Second	The maximum flow setup rate observed since the last PIC restart or since the last time flow statistics are flushed.
Peak Deleted Flows per Second	The maximum flow deletion rate observed since the last PIC restart or from the last time flow statistics are flushed.
Average HTTP Flow Lifetime(ms)	Average HTTP Flow Lifetime in millisecond.
Packets received	The total number of packets received by the MS-PIC.
Packets transmitted	The total number of packets transmitted by the MS-PIC.
Slow path forward	The number of packets forwarded in the slow path (i.e. after the successful rule match and flow creation).
Slow path discard	The number of packets discarded before the flow creation.
Flow Rate Data: Number of Samples	The number of samples used to calculate the flow rate, since the last PIC restart or since the last time flow statistics are flushed.

Table 103: show services stateful-firewall flow-analysis Output Fields (*continued*)

Field Name	Field Description
Flow Rate Distribution(sec) Flow Operation :Creation Flow Operation :Deletion	Histogram of the samples used for flow rate calculation.
Flow Lifetime Distribution(sec):	Histogram of the samples used to calculate the flow life time in sec.

Sample Output

show services stateful-firewall flow-analysis

user@host> **show services stateful-firewall flow-analysis**

```

Services PIC Name: sp-3/0/0
Flow Analysis Statistics:
    Total Flows Active           :40
    Total TCP Flows Active       :0
    Total UDP Flows Active       :40
    Total Other Flows Active     :0
    Total Predicted Flows Active :0
    Created Flows per Second     :0
    Deleted Flows per Second     :0
    Peak Total Flows Active      :40
    Peak Total TCP Flows Active  :0
    Peak Total UDP Flows Active  :40
    Peak Total Other Flows Active :0
    Peak Created Flows per Second :20
    Peak Deleted Flows per Second :20
    Average HTTP Flow Lifetime(ms) :0
    Packets received             :48682539117
    Packets transmitted          :48682502703
    Slow path forward            :6550
    Slow path discard            :0
Flow Rate Data:
    Number of Samples: 19720
Flow Rate Distribution(sec)
Flow Operation :Creation

```

```

300000+      :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000 :0
40000 - 50000 :0
30000 - 40000 :0
20000 - 30000 :0
10000 - 20000 :0
1000 - 10000 :0
0 - 1000 :19720
Flow Operation :Deletion
300000+      :0
250000 - 300000 :0
200000 - 250000 :0
160000 - 200000 :0
150000 - 160000 :0
50000 - 150000 :0
40000 - 50000 :0
30000 - 40000 :0
20000 - 30000 :0
10000 - 20000 :0
1000 - 10000 :0
0 - 1000 :19720
Flow Lifetime Distribution(sec):
          TCP          UDP          HTTP
240+      :0          0          0
120 - 240 :0          0
60 - 120  :0          0
30 - 60   :0          0
15 - 30   :0          6530
5 - 15    :0          0
1 - 5     :0          0
0 - 1     :0          6530

```

Sample Output

show services stateful-firewall flow-analysis interface sp-3/0/0

user@host> **show services stateful-firewall flow-analysis interface sp-3/0/0**

Services PIC Name: sp-3/0/0

Flow Analysis Statistics:

Total Flows Active	:40
Total TCP Flows Active	:0
Total UDP Flows Active	:40
Total Other Flows Active	:0
Total Predicted Flows Active	:0
Created Flows per Second	:0
Deleted Flows per Second	:0
Peak Total Flows Active	:40
Peak Total TCP Flows Active	:0
Peak Total UDP Flows Active	:40
Peak Total Other Flows Active	:0
Peak Created Flows per Second	:20
Peak Deleted Flows per Second	:20
Average HTTP Flow Lifetime(ms)	:0
Packets received	:54696856768
Packets transmitted	:54696815873
Slow path forward	:7350
Slow path discard	:0

Flow Rate Data:

Number of Samples: 22139

Flow Rate Distribution(sec)

Flow Operation :Creation

300000+	:0
250000 - 300000	:0
200000 - 250000	:0
160000 - 200000	:0
150000 - 160000	:0
50000 - 150000	:0
40000 - 50000	:0
30000 - 40000	:0
20000 - 30000	:0
10000 - 20000	:0
1000 - 10000	:0
0 - 1000	:22139

Flow Operation :Deletion

300000+	:0
250000 - 300000	:0
200000 - 250000	:0
160000 - 200000	:0
150000 - 160000	:0
50000 - 150000	:0
40000 - 50000	:0

30000	-	40000	:	0
20000	-	30000	:	0
10000	-	20000	:	0
1000	-	10000	:	0
0	-	1000	:	22139
Flow Lifetime Distribution(sec):				
		TCP	UDP	HTTP
240+		:0	0	0
120 - 240		:0	0	
60 - 120		:0	0	
30 - 60		:0	0	
15 - 30		:0	7330	
5 - 15		:0	0	
1 - 5		:0	0	
0 - 1		:0	7330	

show services stateful-firewall flows

Syntax

```
show services stateful-firewall flows
<brief | extensive | summary | terse>
<application-protocol protocol>
<count>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced before Junos OS Release 7.4.

pgcp option introduced in Junos OS Release 8.4.

application-protocol option introduced in Junos OS Release 10.4.

Description

Display stateful firewall flow table entries. When the interface is used for software processing, the type of software concentrator (**DS-LITE** or **6rd**) is shown, and frame counts are provided.

Options

none—Display standard information about all stateful firewall flows.

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

application-protocol *application-protocol*—(Optional) Display information about one of the following application-level gateway (ALG) protocol types:

- **bootp**—Bootstrap protocol
- **dce-rpc**—Distributed Computing Environment (DCE) remote procedure call (RPC) protocol

NOTE: Use this option to select Microsoft Remote Procedure Call (MSRPC).

- **dce-rpc-portmap**—Distributed Computing Environment (DCE) remote procedure call (RPC) portmap protocol
- **dns**—Domain Name Service protocol

- **exec**—Remote execution protocol
- **ftp**—File Transfer Protocol
- **h323**—H.323 protocol
- **icmp**—Internet Control Message Protocol
- **iiop**—Internet Inter-ORB Protocol
- **ip**—Internet protocol
- **netbios**—NetBIOS protocol
- **netshow**—Netshow protocol
- **pptp**—Point-to-Point Tunneling Protocol
- **realaudio**—RealAudio protocol
- **rpc**—Remote Procedure Call protocol

NOTE: Use this option to select Sun Microsystems Remote Procedure Call protocol (SunRPC).

- **rpc-portmap**—Remote Procedure Call portmap protocol
- **rtsp**—Real-Time Streaming Protocol
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **talk**—Talk protocol
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

count—(Optional) Display a count of the matching entries.

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, *interface-name* can be ***ms-fpc/pic/port*** or ***rspnumber***.

limit *number*—(Optional) Maximum number of entries to display.

protocol *protocol*—(Optional) Display information about one of the following IP types:

- **number**—Numeric protocol value from 0 to 255
- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear services stateful-firewall flows](#) | [774](#)

List of Sample Output

[show services stateful-firewall flows on page 1038](#)

[show services stateful-firewall flows \(For Software Flows\) on page 1039](#)

[show services stateful-firewall flows brief on page 1039](#)

[show services stateful-firewall flows extensive on page 1040](#)

[show services stateful-firewall flows count on page 1040](#)

[show services stateful-firewall flows destination port on page 1040](#)

[show services stateful-firewall flows source port on page 1040](#)

[show services stateful-firewall flows \(Twice NAT\) on page 1041](#)

Output Fields

[Table 104 on page 1038](#) lists the output fields for the **show services stateful-firewall flows** command. Output fields are listed in the approximate order in which they appear.

Table 104: show services stateful-firewall flows Output Fields

Field Name	Field Description
Interface	Name of the interface.
Service set	Name of a service set. Individual empty service sets are not displayed. If no service set has any flows, a flow table header is displayed for each service set.
Flow Count	Number of flows in a session.
Flow or Flow Prot	Protocol used for this flow.
Source	Source prefix of the flow in the format <i>source-prefix:port</i> . For ICMP flows, port information is not displayed.
Dest	Destination prefix of the flow. For ICMP flows, port information is not displayed.
State	Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without response. • Forward—Forward the packet in the flow without looking at it. • Reject—Drop all packets in the flow with response. • Watch—Inspect packets in the flow.
Dir	Direction of the flow: input (I) or output (O). For any configured stateful firewall rule, the reverse flow is dynamically created, so you will see an input and an output flow.
Frm count	Number of frames in the flow. If this value is zero, then that flow does not yet exist.

Sample Output

show services stateful-firewall flows

On the MX Series router, both input (I) and output (O) flow entries appear, even if traffic only flows in one direction. This applies to both NAT and non-NAT cases.


```
user@host> show services stateful-firewall flows
```

```
Interface: ms-1/3/0, Service set: green

Flow
Prot      Source          Dest              State    Dir      Frm count
TCP       10.58.255.178:23  -> 10.59.16.100:4000 Forward  O
TCP       10.58.255.50:33005-> 10.58.255.178:23 Forward  I          1
Source NAT 10.58.255.50:33005-> 10.59.16.100:4000
Destin NAT 10.58.255.178:23  -> 0.0.0.0:4000
```

show services stateful-firewall flows (For Software Flows)

When a service set includes software processing, the following output format is used for the software flows:

```
user@host> show services stateful-firewall flows
```

```
Interface: sp-0/1/0, Service set: dslite-svc-set2

Flow
TCP       200.200.200.2:80  -> 44.44.44.1:1025 Forward  O          219942
NAT dest  44.44.44.1:1025  -> 20.20.1.4:1025
Software  2001::2          -> 1001::1
TCP       20.20.1.2:1025 -> 200.200.200.2:80 Forward  I          110244
NAT source 20.20.1.2:1025  -> 44.44.44.1:1024
Software  2001::2          -> 1001::1
TCP       200.200.200.2:80 -> 44.44.44.1:1024 Forward  O          219140
NAT dest  44.44.44.1:1024  -> 20.20.1.2:1025
Software  2001::2          -> 1001::1
DS-LITE   2001::2          -> 1001::1 Forward  I          988729
TCP       200.200.200.2:80 -> 44.44.44.1:1026 Forward  O          218906
NAT dest  44.44.44.1:1026  -> 20.20.1.3:1025
Software  2001::2          -> 1001::1
TCP       20.20.1.3:1025 -> 200.200.200.2:80 Forward  I          110303
NAT source 20.20.1.3:1025 -> 44.44.44.1:1026
Software  2001::2          -> 1001::1
TCP       20.20.1.4:1025 -> 200.200.200.2:80 Forward  I          110944
NAT source 20.20.1.4:1025 -> 44.44.44.1:1025
Software  2001::2          -> 1001::1
```

show services stateful-firewall flows brief

The output for the **show services stateful-firewall flows brief** command is identical to that for the **show services stateful-firewall flows** command. For sample output, see [show services stateful-firewall flows](#).

show services stateful-firewall flows extensive

```
user@host> show services stateful-firewall flows extensive
```

```
Interface: ms-0/3/0, Service set: ss_nat
Flow
count
TCP          16.1.0.1:2330  ->    16.49.0.1:21      Forward  I
8
  NAT source      16.1.0.1:2330    ->    16.41.0.1:2330
  NAT dest        16.49.0.1:21    ->    16.99.0.1:21
Byte count: 455, TCP established, TCP window size: 57344
TCP acknowledge: 3251737524, TCP tickle enabled, tcp_tickle: 0
Flow role: Master, Timeout: 720
TCP          16.99.0.1:21  ->    16.41.0.1:2330    Forward  O
5
  NAT source      16.99.0.1:21    ->    16.49.0.1:21
  NAT dest        16.41.0.1:2330  ->    16.1.0.1:2330
Byte count: 480, TCP established, TCP window size: 57344
TCP acknowledge: 463128048, TCP tickle enabled, tcp_tickle: 0
Flow role: Responder, Timeout: 720
```

show services stateful-firewall flows count

```
user@host> show services stateful-firewall flows count
```

Interface	Service set	Flow Count
ms-1/3/0	green	2

show services stateful-firewall flows destination port

```
user@host> show services stateful-firewall flows destination-port 21
```

```
Interface: ms-0/3/0, Service set: svc_set_trust
Flow
State Dir Frm count
Interface: ms-0/3/0, Service set: svc_set_untrust
Flow
State Dir Frm count
TCP          10.50.10.2:2143  ->    10.50.20.2:21      Watch   O          0
```

show services stateful-firewall flows source port

```
user@host> show services stateful-firewall flows source-port 2143
```

```
Interface: ms-0/3/0, Service set: svc_set_trust
```

```
Flow
```

```
State Dir Frm count
```

```
Interface: ms-0/3/0, Service set: svc_set_untrust
```

```
Flow
```

```
State Dir Frm count
```

```
TCP 10.50.10.2:2143 -> 10.50.20.2:21 Watch 0 0
```

show services stateful-firewall flows (Twice NAT)

```
user@host> show services stateful-firewall flows
```

```
Flow State Dir Frm count
```

```
UDP 40.0.0.8:23439 -> 80.0.0.1:16485 Watch I 20
```

```
NAT source 40.0.0.8:23439 -> 172.16.1.10:1028
```

```
NAT dest 80.0.0.1:16485 -> 192.16.1.10:22415
```

```
UDP 192.16.1.10:22415 -> 172.16.1.10:1028 Watch O 20
```

```
NAT source 192.16.1.10:22415 -> 80.0.0.1:16485
```

```
NAT dest 172.16.1.10:1028 -> 40.0.0.8:23439
```

show services stateful-firewall sip-call

Syntax

```
show services stateful-firewall sip-call
<brief | extensive | terse>
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced in Junos OS Release 7.4.

Description

Display stateful firewall Session Initiation Protocol (SIP) call information.

Options

count—(Optional) Display a count of the matching entries.

brief—(Optional) Display brief SIP call information.

extensive—(Optional) Display detailed SIP call information.

terse—(Optional) Display terse SIP call information.

application-protocol—(Optional) Display information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol

- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Display information for a particular destination port. The range of values is from 0 to 65535.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.

interface *interface-name*—(Optional) Display information about a particular adaptive services interface. On M Series and T Series routers, *interface-name* can be **sp-fpc/pic/port** or **rspnumber**.

limit *number*—(Optional) Maximum number of entries to display.

protocol—(Optional) Display information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP

- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level

view

RELATED DOCUMENTATION

| [clear services stateful-firewall sip-call](#) | [777](#)

List of Sample Output

[show services stateful-firewall sip-call](#) extensive on page 1045

Output Fields

[Table 105 on page 1044](#) lists the output fields for the **show services stateful-firewall sip-call** command. Output fields are listed in the approximate order in which they appear.

Table 105: show services stateful-firewall sip-call Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.
From	Initiator address.
To	Responder address.
Call ID	SIP call identification string.

Table 105: show services stateful-firewall sip-call Output Fields (*continued*)

Field Name	Field Description
Number of initiator flows	Number of control , contact , or media initiator flows.
Number of responder flows	Number of control , contact , or media responder flows.
<i>protocol</i>	Protocol used for this flow.
<i>source-prefix</i>	Source prefix of the flow in the format <i>source-prefix : port</i> .
<i>destination-prefix</i>	Destination prefix of the flow.
<i>state</i>	Status of the flow: <ul style="list-style-type: none"> • Drop—Drop all packets in the flow without a response. • Forward—Forward the packet in the flow without examining it. • Reject—Drop all packets in the flow with a response. • Unknown—Unknown status. • Watch—Inspect packets in the flow.
<i>direction</i>	Direction of the flow: input (I), output (O), or unknown (U).
<i>frame-count</i>	Number of frames in the flow.
Byte count	Number of bytes forwarded in the flow.
Flow role	Role of the flow that is under evaluation: Initiator , Master , Responder , or Unknown .
Timeout	Lifetime of the flow, in seconds.

Sample Output

show services stateful-firewall sip-call extensive

user@host> show services stateful-firewall sip-call extensive

```
Interface: sp-0/3/0, Service set: test_sip_777
```

```

From: : 6507771234@10.200.100.1:0;000ff73ac89900021bb231dc-3ef68435
To: : 4085551234@10.200.100.1:0;0011bb65c2a30007777bd0fc-5748b749
Call ID: : 000ff73a-c8990004-0741adac-3e027c7e@10.20.70.2
Number of control initiator flows: : 1, Number of control responder flows:
: 1
UDP      10.20.70.2:50354 -> 10.200.100.1:5060 Watch I
2
    Byte count: 1112
    Flow role: Master, Timeout: 30
UDP      10.200.100.1:5060 -> 10.20.170.111:50354 Watch O
0
    Byte count: 0
    Flow role: Responder, Timeout: 30
UDP      0.0.0.0:0 -> 10.20.170.111:5060 Watch O
7
    Byte count: 2749
    Flow role: Responder, Timeout: 30
Number of contact initiator flows: 1, Number of contact responder flows: 1
UDP      0.0.0.0:0 -> 10.20.140.11:5060 Watch I
1
    Byte count: 409
    Flow role: Master, Timeout: 30
UDP      10.20.140.11:31864 -> 10.20.170.111:18808 Forward O
622
    Byte count: 124400
    Flow role: Master, Timeout: 30
UDP      0.0.0.0:0 -> 10.20.170.111:18809 Forward O
0
    Byte count: 0
    Flow role: Initiator, Timeout: 30
Number of media initiator flows: 4, Number of media responder flows: 0
UDP      10.20.70.2:18808 -> 10.20.140.11:31864 Forward I
628
    Byte count: 125600
    Flow role: Initiator, Timeout: 30
UDP      0.0.0.0:0 -> 10.20.140.11:31865 Forward I
0
    Byte count: 0
    Flow role: Initiator, Timeout: 30
0      0.0.0.0:0 -> 0.0.0.0:0 Unknown U
0
    Byte count: 0
    Flow role: Unknown, Timeout: 0

```



```
0          0.0.0.0:0    ->    0.0.0.0:0    Unknown  U
Interface: sp-0/3/0, Service set: test_sip_888
```

show services stateful-firewall sip-register

Syntax

```
show services stateful-firewall sip-register
<brief | extensive | terse>
<application-protocol protocol>
<destination-port destination-port>
<destination-prefix destination-prefix>
<interface interface-name>
<limit number>
<protocol protocol>
<service-set service-set>
<source-port source-port>
<source-prefix source-prefix>
```

Release Information

Command introduced in Junos OS Release 7.4.

Description

Display stateful firewall Session Initiation Protocol (SIP) register information.

Options

count—(Optional) Display a count of the matching entries.

brief—(Optional) Display brief SIP register information.

extensive—(Optional) Display detailed SIP register information.

terse—(Optional) Display terse SIP register information.

application-protocol—(Optional) Display information about one of the following application protocols:

- **bootp**—(SIP only) Bootstrap protocol
- **dce-rpc**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols
- **dce-rpc-portmap**—(SIP only) Distributed Computing Environment-Remote Procedure Call protocols portmap service
- **dns**—(SIP only) Domain Name System protocol
- **exec**—(SIP only) Exec
- **ftp**—(SIP only) File Transfer Protocol
- **h323**—H.323 standards
- **icmp**—Internet Control Message Protocol

- **iiop**—Internet Inter-ORB Protocol
- **login**—Login
- **netbios**—NetBIOS
- **netshow**—NetShow
- **realaudio**—RealAudio
- **rpc**—Remote Procedure Call protocol
- **rpc-portmap**—Remote Procedure Call protocol portmap service
- **rtsp**—Real-Time Streaming Protocol
- **shell**—Shell
- **sip**—Session Initiation Protocol
- **snmp**—Simple Network Management Protocol
- **sqlnet**—SQLNet
- **tftp**—Trivial File Transfer Protocol
- **traceroute**—Traceroute
- **winframe**—WinFrame

destination-port *destination-port*—(Optional) Display information for a particular destination port.

destination-prefix *destination-prefix*—(Optional) Display information for a particular destination prefix.
The range of values is from 0 to 65535.

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be *sp-fpc/pic/port* or *rspnumber*.

limit *number*—(Optional) Maximum number of entries to display.

protocol—(Optional) Display information about one of the following IP types:

- **ah**—IPsec Authentication Header protocol
- **egp**—An exterior gateway protocol
- **esp**—IPsec Encapsulating Security Payload protocol
- **gre**—A generic routing encapsulation protocol
- **icmp**—Internet Control Message Protocol
- **igmp**—Internet Group Management Protocol
- **ipip**—IP-within-IP Encapsulation Protocol
- **ipv6**—IPv6 within IP

- **ospf**—Open Shortest Path First protocol
- **pim**—Protocol Independent Multicast protocol
- **rsvp**—Resource Reservation Protocol
- **sctp**—Stream Control Protocol
- **tcp**—Transmission Control Protocol
- **udp**—User Datagram Protocol

service-set *service-set*—(Optional) Display information for a particular service set.

source-port *source-port*—(Optional) Display information for a particular source port. The range of values is from 0 to 65535.

source-prefix *source-prefix*—(Optional) Display information for a particular source prefix.

Required Privilege Level

view

RELATED DOCUMENTATION

| [clear services stateful-firewall sip-register](#) | 780

List of Sample Output

[show services stateful-firewall sip-register](#) extensive on page 1051

Output Fields

[Table 106 on page 1050](#) lists the output fields for the **show services stateful-firewall sip-register** command. Output fields are listed in the approximate order in which they appear.

Table 106: show services stateful-firewall sip-register Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.
SIP Register	Register information header.
Protocol	Protocol used for this flow.
Registered IP	Register IP address.

Table 106: show services stateful-firewall sip-register Output Fields (*continued*)

Field Name	Field Description
Port	Register port number.
Expiration timeout	Configured lifetime, in seconds.
Timeout remaining	Lifetime remaining, in seconds.
From	Initiator address.
To	Responder address.
Call ID	SIP call identification string.

Sample Output

show services stateful-firewall sip-register extensive

user@host> **show services stateful-firewall sip-register extensive**

```
Interface: sp-0/3/0, Service set: test_sip_777
```

```
SIP Register: Protocol: UDP, Registered IP: 10.20.170.111, Port: 5060, Acked
Expiration timeout: 36000, Timeout remaining: 35544
From: : 6507771234@10.200.100.1:0;
To: : 6507771234@10.200.100.1:0;
Call ID: : 000ff73a-c8990002-23b1d942-2ba1f91f@10.20.70.2
```

```
Interface: sp-0/3/0, Service set: test_sip_888
```

```
SIP Register: Protocol: UDP, Registered IP: 10.20.170.112, Port: 5060, Acked
Expiration timeout: 36000, Timeout remaining: 35549
From: : 8881234@10.200.100.1:0;
To: : 8881234@10.200.100.1:0;
Call ID: : 00112096-81fc0002-23b38905-7cb41f62@10.20.71.2
```

show services stateful-firewall statistics

Syntax

```
show services stateful-firewall statistics
<application-protocol protocol>
<brief | detail | extensive | summary>
<interface interface-name>
<service-set service-set>
```

Release Information

Command introduced before Junos OS Release 7.4.

Description

Display stateful firewall statistics.

Options

none—Display standard information about all stateful firewall statistics.

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

interface *interface-name*—(Optional) Display information about a particular interface. On M Series and T Series routers, the *interface-name* can be **ms-fpc/pic/port** or **rspnumber**.

service-set *service-set*—(Optional) Display information about a particular service set.

Required Privilege Level

view

RELATED DOCUMENTATION

[clear services stateful-firewall statistics](#) | 783

List of Sample Output

[show services stateful-firewall statistics extensive on page 1060](#)

Output Fields

[Table 107 on page 1053](#) lists the output fields for the **show services stateful-firewall statistics** command. Output fields are listed in the approximate order in which they appear.

Table 107: show services stateful-firewall statistics Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of a service set.
New flows	Rule match counters for new flows: <ul style="list-style-type: none"> • Rule Accepts—New flows accepted. • Rule Discards—New flows discarded. • Rule Rejects—New flows rejected.
Existing flow types packet counters	Rule match counters for existing flows: <ul style="list-style-type: none"> • Accepts—Match existing forward or watch flow. • Drop—Match existing discard flow. • Rejects—Match existing reject flow.
Hairpinning Counters	Hairpinning counters: <ul style="list-style-type: none"> • Slow Path Hairpinned Packets—Slow path packets that were hairpinned back to the internal network. • Fast Path Hairpinned Packets—Fast path packets that were hairpinned back to the internal network.
Drops	Drop counters: <ul style="list-style-type: none"> • IP option—Packets dropped in IP options processing. • TCP SYN defense—Packets dropped by SYN defender. • NAT ports exhausted—Hide mode. The router has no available Network Address Translation (NAT) ports for a given address or pool. • Sessions dropped due to subscriber flow limit—Sessions dropped because the subscriber's flow limit was exceeded.
Errors	Total errors, categorized by protocol: <ul style="list-style-type: none"> • IP—Total IP version 4 errors. • TCP—Total Transmission Control Protocol (TCP) errors. • UDP—Total User Datagram Protocol (UDP) errors. • ICMP—Total Internet Control Message Protocol (ICMP) errors. • Non-IP packets—Total non-IPv4 errors. • ALG—Total application-level gateway (ALG) errors

Table 107: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
IP Errors	<p>IPv4 errors:</p> <ul style="list-style-type: none"> • IP packet length inconsistencies—IP packet length does not match the Layer 2 reported length. • Minimum IP header length check failures—Minimum IP header length is 20 bytes. The received packet contains less than 20 bytes. • Reassembled packet exceeds maximum IP length—After fragment reassembly, the reassembled IP packet length exceeds 65,535. • Illegal source address 0—Source address is not a valid address. Invalid addresses are, loopback, broadcast, multicast, and reserved addresses. Source address 0, however, is allowed to support BOOTP and the destination address 0xffffffff. • Illegal destination address 0—Destination address is not a valid address. The address is reserved. • TTL zero errors—Received packet had a time-to-live (TTL) value of 0. • Illegal IP protocol number (0 or 255)—IP protocol is 0 or 255. • Land attack—IP source address is the same as the destination address. • Non-IPv4 packets—Packet was not IPv4. (Only IPv4 is supported.) • Bad checksum—Packet had an invalid IP checksum. • Illegal IP fragment length—Illegal fragment length. All fragments (other than the last fragment) must have a length that is a multiple of 8 bytes. • IP fragment overlap—Fragments have overlapping fragment offsets. • IP fragment reassembly timeout—Some of the fragments for an IP packet were not received in time, and the reassembly handler dropped partial fragments. • IP fragment limit exceeded: 0—Fragments that exceeded the limit. • Unknown: 0—Unknown fragments.

Table 107: show services stateful-firewall statistics Output Fields *(continued)*

Field Name	Field Description
TCP Errors	

Table 107: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	<p>TCP protocol errors:</p> <ul style="list-style-type: none"> • TCP header length inconsistencies—Minimum TCP header length is 20 bytes, and the IP packet received does not contain at least 20 bytes. • Source or destination port number is zero—TCP source or destination port is zero. • Illegal sequence number and flags combinations — Dropped because of TCP errors, such as an illegal sequence number, which causes an illogical combination of flags to be set. • SYN attack (multiple SYN messages seen for the same flow)—Multiple SYN packets received for the same flow are treated as a SYN attack. The packets might be retransmitted SYN packets and therefore valid, but a large number is cause for concern. • First packet not a SYN message—First packets for a connection are not SYN packets. These packets might originate from previous connections or from someone performing an ACK/FIN scan. • TCP port scan (TCP handshake, RST seen from server for SYN)—In the case of a SYN defender, if an RST (reset) packet is received instead of a SYN/ACK message, someone is probably trying to scan the server. This behavior can result in false alarms if the RST packet is not combined with an intrusion detection service (IDS). • Bad SYN cookie response—SYN cookie generates a SYN/ACK message for all incoming SYN packets. If the ACK received for the SYN/ACK message does not match, this counter is incremented. • TCP reconstructor sequence number error—This counter is incremented in the following cases: The TCP seqno is 0 and all the TCP flags are also 0. The TCP seqno is 0 and FIN/PSH/URG TCP flags are set. • TCP reconstructor retransmissions—This counter is incremented for the retransmitted packets during connection 3-way handshake. • TCP partially opened connection timeout (SYN)—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder. • TCP partially opened connection timeout (SYN-ACK)—This counter is incremented when the SYN Defender is enabled and the 3-way handshake is not completed within the SYN DEFENDER TIMEOUT. The connection will be closed and resources will be released by sending RST to the responder.

Table 107: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> • TCP partially closed connection reuse—Not supported. • TCP 3-way error - client sent SYN+ACK—A SYN/ACK should be sent by the server on receiving a SYN. This counter is incremented when the first message received from the initiator is SYN+ACK. • TCP 3-way error - server sent ACK—ACK should be sent by the client on receiving a SYN/ACK from the server. This counter is incremented when the ACK is received from the Server instead of from the Client. • TCP 3-way error - SYN seq number retransmission mismatch—This counter is incremented when the SYN is received again with a different sequence number from the first SYN sequence number. • TCP 3-way error - RST seq number mismatch—A reset could be received from either side. The server could send a RST on receiving a SYN or the client could send a RST on receiving SYN/ACK. This counter is incremented when the RST is received either from the client or server with a non-matching sequence number. • TCP 3-way error - FIN received—This counter is incremented when the FIN is received during the 3-way handshake. • TCP 3-way error - invalid flags (PSH, URG, ECE, CWR)—This counter is incremented when any of the PSH, URG, ECE, or CWR flags were received during the 3-way handshake. • TCP 3-way error - SYN recvd but no client flows—This counter is incremented when SYN is received but not from the connection initiator. The counter is not incremented in the case of simultaneous open, when the SYN is received in both the directions. • TCP 3-way error - first packet SYN+ACK—The first packet received was SYN+ACK instead of SYN. • TCP 3-way error - first packet FIN+ACK—The first packet received was FIN+ACK instead of SYN. • TCP 3-way error - first packet FIN—The first packet received was FIN instead of SYN. • TCP 3-way error - first packet RST—The first packet received was RST instead of SYN. • TCP 3-way error - first packet ACK—The first packet received was ACK instead of SYN. • TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR)—The first packet received had invalid flags. • TCP Close error - no final ACK—This counter is incremented when ACK is not received after the FINs are received from both directions.

Table 107: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	<ul style="list-style-type: none"> • TCP Resumed Flow—Plain ACKs create flows if rule match permits, and these are classified as TCP Resumed Flows. This counter is incremented in the case of a TCP Resumed Flow.
UDP Errors	<p>UDP protocol errors:</p> <ul style="list-style-type: none"> • IP data length less than minimum UDP header length (8 bytes)—Minimum UDP header length is 8 bytes. The received IP packets contain less than 8 bytes. • Source or destination port is zero—UDP source or destination port is 0. • UDP port scan (ICMP error seen for UDP flow)—ICMP error is received for a UDP flow. This could be a genuine UDP flow, but it is counted as an error.
ICMP Errors	<p>ICMP protocol errors:</p> <ul style="list-style-type: none"> • IP data length less than minimum ICMP header length (8 bytes)—ICMP header length is 8 bytes. This counter is incremented when received IP packets contain less than 8 bytes. • ICMP error length inconsistencies—Minimum length of an ICMP error packet is 48 bytes, and the maximum length is 576 bytes. This counter is incremented when the received ICMP error falls outside this range. • Duplicate ping sequence number—Received ping packet has a duplicate sequence number. • Mismatched ping sequence number—Received ping packet has a mismatched sequence number. • No matching flow—No matching existing flow was found for the ICMP error.

Table 107: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
ALG errors	<p>Accumulation of all the application-level gateway protocol (ALG) drops counted separately in the ALG context:</p> <ul style="list-style-type: none"> • BOOTP—Bootstrap protocol errors • DCE-RPC—Distributed Computing Environment-Remote Procedure Call protocols errors • DCE-RPC portmap—Distributed Computing Environment-Remote Procedure Call protocols portmap service errors • DNS—Domain Name System protocol errors • Exec—Exec errors • FTP—File Transfer Protocol errors • H323—H.323 standards errors • ICMP—Internet Control Message Protocol errors • IIOp—Internet Inter-ORB Protocol errors • Login—Login errors • NetBIOS—NetBIOS errors • Netshow—NetShow errors • Real Audio—RealAudio errors • RPC—Remote Procedure Call protocol errors • RPC portmap—Remote Procedure Call protocol portmap service errors • RTSP—Real-Time Streaming Protocol errors • Shell—Shell errors • SIP—Session Initiation Protocol errors • SNMP—Simple Network Management Protocol errors • SQLNet—SQLNet errors • TFTP—Trivial File Transfer Protocol errors • Traceroute—Traceroute errors

Table 107: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
Drop Flows	<ul style="list-style-type: none"> • Maximum Ingress Drop flows allowed--Maximum number of ingress flow drops allowed. • Maximum Egress Drop flows allowed--Maximum number of egress flow drops allowed. • Current Ingress Drop flows--Current number of ingress flow drops. • Current Egress Drop flows--Current number of egress flow drops. • Ingress Drop Flow limit drops count--Number of ingress flow drops due to maximum number of ingress flow drops being exceeded. • Egress Drop Flow limit drops count--Number of egress flow drops due to maximum number of egress flow drops being exceeded.

Sample Output

show services stateful-firewall statistics extensive

user@host> **show services stateful-firewall statistics extensive**

```

Interface: ms-1/3/0
Service set: interface-svc-set
New flows:
  Rule Accepts: 907, Rule Discards: 0, Rule Rejects: 0
Existing flow types packet counters:
  Accepts: 3535, Drop: 0, Rejects: 0
Hairpinning counters:
  Slow Path Hairpinned Packets: 0, Fast Path Hairpinned Packets: 0
Drops:
  IP option: 0, TCP SYN defense: 0
  NAT ports exhausted: 0, Sessions dropped due to subscriber flow limit: 0
Errors:
  IP: 0, TCP: 0
  UDP: 0, ICMP: 0
  Non-IP packets: 0, ALG: 0
IP errors:
  IP packet length inconsistencies: 0
  Minimum IP header length check failures: 0
  Reassembled packet exceeds maximum IP length: 0
  Illegal source address: 0
  Illegal destination address: 0

```

```
TTL zero errors: 0, Illegal IP protocol number (0 or 255): 0
Land attack: 0
Non-IPv4 packets: 0, Bad checksum: 0
Illegal IP fragment length: 0
IP fragment overlap: 0
IP fragment reassembly timeout: 0
IP fragment limit exceeded:0
Unknown: 0
TCP errors:
  TCP header length inconsistencies: 0
  Source or destination port number is zero: 0
  Illegal sequence number and flags combination: 0
  SYN attack (multiple SYN messages seen for the same flow): 0
  First packet not a SYN message: 0
  TCP port scan (TCP handshake, RST seen from server for SYN): 0
  Bad SYN cookie response: 0
  TCP reconstructor sequence number error: 0
  TCP reconstructor retransmissions: 0
  TCP partially opened connection timeout (SYN): 0
  TCP partially opened connection timeout (SYN-ACK): 0
  TCP partially closed connection reuse: 0
  TCP 3-way error - client sent SYN+ACK: 0
  TCP 3-way error - server sent ACK: 0
  TCP 3-way error - SYN seq number retransmission mismatch: 0
  TCP 3-way error - RST seq number mismatch: 0
  TCP 3-way error - FIN received: 0
  TCP 3-way error - invalid flags (PSH, URG, ECE, CWR): 0
  TCP 3-way error - SYN recvd but no client flows: 0
  TCP 3-way error - first packet SYN+ACK: 0
  TCP 3-way error - first packet FIN+ACK: 0
  TCP 3-way error - first packet FIN: 0
  TCP 3-way error - first packet RST: 0
  TCP 3-way error - first packet ACK: 0
  TCP 3-way error - first packet invalid flags (PSH, URG, ECE, CWR): 0
  TCP Close error - no final ACK: 0
  TCP Resumed Flow: 0
UDP errors:
  IP data length less than minimum UDP header length (8 bytes): 0
  Source or destination port is zero: 0
  UDP port scan (ICMP error seen for UDP flow): 0
ICMP errors:
  IP data length less than minimum ICMP header length (8 bytes): 0
  ICMP error length inconsistencies: 0
  Duplicate ping sequence number: 0
```

```
Mismatched ping sequence number: 0
No matching flow: 0
ALG errors:
  BOOTP: 0, DCE-RPC: 0, DCE-RPC portmap: 0
  DNS: 0, Exec: 0, FTP: 0
  H323: 0, ICMP: 0, IIOP: 0
  Login: 0, NetBIOS: 0, Netshow: 0
  Real Audio: 0, RPC: 0, RPC portmap: 0
  RTSP: 0, Shell: 0, SIP: 0
  SNMP: 0, SQLNet: 0, TFTP: 0
  Traceroute: 0
```

```
Drop Flows:
  Maximum Ingress Drop flows allowed: 20
  Maximum Egress Drop flows allowed: 20
  Current Ingress Drop flows: 0
  Current Egress Drop flows: 0
  Ingress Drop Flow limit drops count: 0
  Egress Drop Flow limit drops count: 0
```

****If max-drop-flows is not configured, the following is shown****

```
Drop Flows:
  Maximum Ingress Drop flows allowed: Default
  Maximum Egress Drop flows allowed: Default
```


show services stateful-firewall statistics application-protocol sip

Syntax

```
show services stateful-firewall application-protocol sip
```

Release Information

Command introduced in Junos OS Release 7.4.

Description

Display stateful firewall Session Initiation Protocol (SIP) statistics.

Options

This command has no options.

Required Privilege Level

view

List of Sample Output

[show services stateful-firewall statistics application-protocol-sip on page 1065](#)

Output Fields

[Table 108 on page 1063](#) lists the output fields for the **show services stateful-firewall statistics application-protocol-sip** command. Output fields are listed in the approximate order in which they appear.

Table 108: show services stateful-firewall statistics application-protocol-sip Output Fields

Field Name	Field Description
Interface	Name of an adaptive services interface.
Service set	Name of the service set flow.
ALG	Name of the application-layer gateway.
Active SIP call count	Number of active SIP calls.
Active SIP registration count	Number of active SIP registrations.
REGISTER	Number of new, invalid, and retransmitted register requests sent to the SIP registrar.
INVITE	Number of new, invalid, and retransmitted invite messages sent by user agent clients.
ReINVITE	Number of new, invalid, and retransmitted reinvite messages sent by user agent clients.

Table 108: show services stateful-firewall statistics application-protocol-sip Output Fields (*continued*)

Field Name	Field Description
ACK	Number of new, invalid, and retransmitted ACK messages received (in response to a SIP Call Invite message).
BYE	Number of new, invalid, and retransmitted requests to terminate SIP dialogues.
CANCEL	Number of new, invalid, and retransmitted SIP request cancellations.
SUBSCRIBE	Number of new, invalid, and retransmitted SIP requests to subscribe for event notifications.
NOTIFY	Number of new, invalid, and retransmitted event notifications in SIP dialogues.
OPTIONS	Number of new, invalid, and retransmitted requests to query SIP capabilities.
INFO	Number of new, invalid, and retransmitted requests carrying application-level information.
UPDATE	Number of new, invalid, and retransmitted SIP dialogue updates.
REFER	Number of new, invalid, and retransmitted requests to the recipient to contact a third party.
Provisional responses	Number of new, invalid, and retransmitted responses from the user agent server to indicate the progress of a SIP transaction.
OK responses to INVITES	OK responses sent from the user agent clients to user agent servers in response to Invite messages. The server can then return an ACK message.
OK responses to non-INVITES	OK responses to SIP messages other than an Invite message.
Redirection responses	Responses from the user agent server to a user agent client requesting the client to contact a different SIP uniform resource identifier (URI).
Request failure responses	Responses that indicate a definite failure from a particular server. The client must not retry the same request without modification after receiving this response.
Server failure responses	Responses that indicate a server failure.
Global failure responses	Responses that indicate a server has definitive information about a particular user, not just the particular instance indicated in the Request URI.
Invalid responses	Responses that are invalid.

Table 108: show services stateful-firewall statistics application-protocol-sip Output Fields (*continued*)

Field Name	Field Description
Response (all) retransmits	Retransmissions of all responses.
Parser	Syntax errors, content errors, and unknown methods counted by the message parser.

Sample Output

show services stateful-firewall statistics application-protocol-sip

user@host> **show services stateful-firewall statistics application-protocol sip**

```
Interface: sp-0/3/0
Service set: test_sip_777, ALG: SIP
Active SIP call count: 0, Active SIP registration count: 1
```

	New	Invalid	Retransmit
REGISTER	2		
INVITE	1		0
ReINVITE	1		
ACK	1	0	0
BYE	0	0	
CANCEL	0	0	
SUBSCRIBE	0	0	
NOTIFY	0	0	
OPTIONS	0	0	
INFO	0	0	
UPDATE	0	0	
REFER	0	0	

```
Provisional responses (18x): 1, OK responses to INVITEs: 2
OK responses to non-INVITEs: 2, Redirection (3xx) responses: 0
Request failure (4xx) responses: 0, Server failure (5xx) responses: 0
Global failure (6xx) responses: 0, Invalid responses: 0
Response (all) retransmits: 0
Parser:
  Syntax errors: 0, Content errors: 0, Unknown methods: 0
Service set: test_sip_888, ALG: SIP
Active SIP call count: 0, Active SIP registration count: 1
```

	New	Invalid	Retransmit
REGISTER	2		
INVITE	0		0

ReINVITE	0		
ACK	0	0	0
BYE	0	0	
CANCEL	0	0	
SUBSCRIBE	0	0	
NOTIFY	0	0	
OPTIONS	0	0	
INFO	0	0	
UPDATE	0	0	
REFER	0	0	

Provisional responses (18x): 0, OK responses to INVITEs: 0
OK responses to non-INVITEs: 2, Redirection (3xx) responses: 0
Request failure (4xx) responses: 0, Server failure (5xx) responses: 0
Global failure (6xx) responses: 0, Invalid responses: 0
Response (all) retransmits: 0
Parser:
Syntax errors: 0, Content errors: 0, Unknown methods: 0

show services subscriber analysis

Syntax

```
show services subscriber analysis
<interface interface-name>
```

Release Information

Statement introduced in Junos OS Release 17.1 on MX Series MS-MPC.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display information about the number of active subscribers on the services PIC.

Options

none—Display standard information about all active subscribers on the PIC.

interface *interface-name*—(Optional) Display information about the specified interface.

Required Privilege Level

view

List of Sample Output

[show services subscriber analysis interface on page 1068](#)

Output Fields

[Table 109 on page 1067](#) lists the output fields for the **show services subscriber analysis** command. Output fields are listed in the approximate order in which they appear.

Table 109: show services subscriber analysis Output Fields

Field Name	Field Description
Services PIC Name	Name of an adaptive services interface.
Subscriber Analysis Statistics:	
Total Subscribers Active	Total number of subscribers currently active on the service PIC.
Created Subscribers per Second	Rate at which subscribers are currently being created on the service PIC.
Deleted Subscribers per Second	Rate at which subscribers are currently being deleted on the service PIC.

Table 109: show services subscriber analysis Output Fields (*continued*)

Field Name	Field Description
Peak Total Subscribers Active	Highest number of subscribers that were active during the lifetime of the service PIC.
Peak Created Subscribers per Second	Highest rate at which subscribers were being created during the lifetime of the service PIC.
Peak Deleted Subscribers per Second	Highest rate at which subscribers were being deleted during the lifetime of the service PIC.
Number of Samples	Number of samples during the current sampling period lifetime.
Subscriber Rate Distribution(sec)	
Subscriber Operation: Creation	Number of sampling intervals during which a number of subscribers in the indicated range were created during the current sampling period.
Subscriber Operation: Deletion	Number of sampling intervals during which a number of subscribers in the indicated range were deleted during the current sampling period.

Sample Output

show services subscriber analysis interface

```
user@host> show services subscriber analysis interface ms-5/1/0
```

```
Services PIC Name:      ms-5/1/0
```

```
Subscriber Analysis Statistics:
```

```

Total Subscribers Active           :0
Created Subscribers per Second     :0
Deleted Subscribers per Second     :0
Peak Total Subscribers Active      :0
Peak Created Subscribers per Second :0
Peak Deleted Subscribers per Second :0
```

```
Subscriber Rate Data:
```

```
Number of Samples: 3916
```

Subscriber Rate Distribution(sec)

Subscriber Operation :Creation

400000+		:0
350001	- 400000	:0
300001	- 350000	:0
250001	- 300000	:0
200001	- 250000	:0
160001	- 200000	:0
150001	- 160000	:0
50001	- 150000	:0
40001	- 50000	:0
30001	- 40000	:0
20001	- 30000	:0
10001	- 20000	:0
1001	- 10000	:0
1	- 1000	:0
	0	:3916

Subscriber Operation :Deletion

400000+		:0
350001	- 400000	:0
300001	- 350000	:0
250001	- 300000	:0
200001	- 250000	:0
160001	- 200000	:0
150001	- 160000	:0
50001	- 150000	:0
40001	- 50000	:0
30001	- 40000	:0
20001	- 30000	:0
10001	- 20000	:0
1001	- 10000	:0
1	- 1000	:0
	0	:3916

show services tcp-log

Syntax

```
show services tcp-log
```

Release Information

Command introduced in Junos OS Release 19.3R2 for Next Gen Services on MX Series.

Description

Display the specified TCP log.

Required Privilege Level

RELATED DOCUMENTATION

Sample Output

show services tcp-log

```
user@host> show services tcp-log
```

```
user@hst> show services tcp-log log1
Interface: vms-1/0/0

State: Reconnect-In-Progress
      5.5.5.1 -> 70.0.0.2 : 514
```


show services traffic-load-balance statistics

Syntax

```
show services traffic-load-balance statistics
<extensive>
<group group-name>
<instance instance-name>
<num-instances number>
<real-service real-service-name>
<summary>
<virtual-service virtual-service-name>
```

Release Information

Statement introduced in Junos OS Release 16.1 on MX Series.

num-instances option added in Junos OS Release 16.1R6 and 18.2R1 on MX Series.

Support added in Junos OS 19.3R2 for Next Gen Services with the MX-SPC3 services card.

Description

The basic form of the command displays the list of real servers associated with this group and traffic statistics, including packet count and byte count

Options

none—Display information about the load-balancing statistics in brief.

extensive—(Optional) Display extensive information about the traffic load-balancing statistics.

group *group-name*—(Optional) Display load-balancing statistics for a specified group of load-balancer servers.

instance *instance-name*—(Optional) Display load-balancing statistics for a specific traffic load balancer (TLB) instance.

num-instances *number*—(Optional) Display load-balancing statistics for a specified number of TLB instances.

real-service *real-service-name*—(Optional) Display load-balancing statistics for a specified load balancer serve.

summary—(Optional) Display summary information about the traffic load-balancing statistics.

virtual-service *virtual-service-name*—(Optional) Display load-balancing statistics for a specified TLB virtual service.

Required Privilege Level

view

List of Sample Output

[show services traffic-load-balance statistics on page 1079](#)

[show services traffic-load-balance statistics extensive on page 1080](#)

[show services traffic-load-balance statistics summary on page 1084](#)

Output Fields

Table 110 on page 1072 lists the output fields for the **show services traffic-load-balance statistics** command. Output fields are listed in the approximate order in which they appear.

Table 110: show services traffic-load-balance statistics Output Fields

Field Name	Field Description	Level of Output
Traffic load balance instance name	Name of the traffic load balancer (TLB) instance that contains the load-distribution-related configuration settings.	All levels
Multi services interface name	<p>Name of the services interface used for the TLB instance to provide one-to-one redundancy for server health monitoring.</p> <p>For MS-MPC services card, this is the name of the aggregated multiservices (AMS) interface or “ms-slot/pic/port”.</p> <p>For Next Gen Services and the MX-SPC3 services card, this is the name of the VMS interface or “vms-slot/pic/port”.</p>	All levels
Interface state	<p>Inter-process communications (IPC) status between the TLB daemon (traffic-dird) and the health checking daemon (net-monitor).</p> <ul style="list-style-type: none"> DOWN UP 	All levels
Interface type	Logical interface type.	All levels
Route hold timer	Time that the programmed VIP routes are kept intact after connectivity between traffic-dird and net-monitor daemons is lost. If connectivity is not reestablished within this time, all the VIP routes are withdrawn.	All levels
Traffic load balance virtual svc name	Name of the virtual service for the TLB instance. The virtual service provides an address that is associated with the group of servers to which traffic is directed.	none extensive
Virtual service	Name of the virtual service for the TLB instance. The virtual service provides an address that is associated with the group of servers to which traffic is directed.	summary

Table 110: show services traffic-load-balance statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Routing instance name	Name of the routing instance used for the virtual service.	none extensive
IP address	IP address of the virtual service.	none extensive
Address	IP address of the virtual service.	summary
Sts	Operational state of the virtual service.	summary
Packet Sent	Number of packets originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	summary
Byte Sent	Number of bytes originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	summary
Packet Recv	Number of packets returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	summary
Byte Recv	Number of bytes returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	summary
Virtual service mode	Virtual service processing mode. <ul style="list-style-type: none"> ● layer-2-direct-server-return—Virtual service is in transparent mode with Layer 2 direct server return (DSR) ● direct-server-return—Virtual service is in transparent mode with Layer 3 direct server return (DSR) ● translated—Virtual service is in translated mode. 	none extensive
Traffic load balance group name	Server group name used for the virtual service.	none extensive
Health check interface subunit	Number of the subunit of the multiservice interface used for health checking.	none extensive
Traffic load balance group down count	Number of times the status of the TLB server group was down.	extensive

Table 110: show services traffic-load-balance statistics Output Fields (continued)

Field Name	Field Description	Level of Output
Protocol	Virtual service protocol, either tcp or udp. In translated mode, packets destined to the virtual service IP address+port number+protocol are load balanced and then replaced by the real service IP address and server listening port number.	none extensive
Port Number	Virtual service port number. In translated mode, packets destined to the virtual service IP address+port number+protocol are load balanced and then replaced by the real service IP address and server listening port number.	none extensive
Server Listening Port Number	Real service port number that replaces the virtual service port number. In translated mode, packets destined to the virtual service IP address+port number+protocol are load balanced and then replaced by the real service IP address and server listening port number.	none extensive
Demux Nexthop index	Index number of the demultiplexing next hop for the virtual service. Index number is unique for a VIP, routing-instance, and protocol combination. The demultiplexing next hop is responsible for port-based demultiplexing of traffic to the load-balancing next hop for session distribution.	none extensive
DFW client-id	Client connection identifier assigned to the TLB daemon (traffic-dird) by the firewall daemon (dfwd) when the daemons are successfully connected.	extensive
Traffic load balance group warmup time	Time, in seconds, that passes after the traffic-dird daemon comes up until the traffic-dird programs the distribution table on the Packet Forwarding Engine.	extensive
Traffic load balance group auto-rejoin	Indicates whether the option that allows a server to rejoin the group automatically when it comes up is enabled or not.	extensive
Route metric	Routing metric assigned to the virtual service. A lower metric makes a route more preferred.	extensive
Virtual service down count	Number of times the status of the virtual service was down.	extensive
Traffic load balance hash method	Hash key parameter used for load balancing. Hash keys supported in the ingress direction are protocol, source IP address, and destination IP address.	extensive

Table 110: show services traffic-load-balance statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Nexthop index	Index number of the next-hop for the virtual service. A group of servers function as a pool for next-hop session distribution.	none extensive
Up time	Period of time for which the virtual service is up, in the format <i>number-of-days hh:mm:ss</i> .	none extensive
Real Server Up count	Starting in Junos OS Release 16.1R6 and 18.2R1, number of real servers that are up for the specified virtual service or server group.	none
Real Server Down count	Starting in Junos OS Release 16.1R6 and 18.2R1, number of real servers that are down for the specified virtual service or server group.	none
Total packet sent count	Number of packets originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	none extensive
Total byte sent count	Number of bytes originating from the clients that the TLB instance virtual service processes for load balancing to next-hop servers.	none extensive
Total packet received count	Number of packets returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	none extensive
Total byte received count	Number of bytes returning from the next-hop servers that the TLB instance virtual service processes and forwards to the clients.	none extensive
Network monitoring profile count	Number of network monitoring profiles that are used to monitor the health of servers used in TLB session distribution.	extensive
Active real service count	Number of real services that are functional and active.	extensive
Total real service count	Total number of real services in different states.	extensive
Network monitoring profile index	Unique index number associated with the network monitoring profile. Network monitoring profiles are used to monitor the health of servers used in TLB session distribution.	extensive

Table 110: show services traffic-load-balance statistics Output Fields (continued)

Field Name	Field Description	Level of Output
Network monitoring profile name	Name configured for the network monitoring profile.	extensive
Probe type	Probe type used to examine the health of servers. TLB supports ICMP, TCP, and HTTP health check probes to monitor the health of servers in a group.	extensive
Probe interval	Frequency, in number of seconds, at which health check probes are sent.	extensive
Probe failure retry count	Number of failure retries, after which the real service is tagged as down.	extensive
Probe recovery retry count	Number of successful retries after which the real service is tagged as up.	extensive
Real service	Name of the TLB server (also referred to as real service). The name is the identifier for a server to which sessions can be distributed using the server distribution table in conjunction with the session distribution API.	none
Address	IP address of the configured real service.	none
Sts	Operational state of the TLB server.	none
Packet Sent	Number of packets originating from the clients that the TLB instance virtual service sends to the real service.	none
Byte Sent	Number of bytes originating from the clients that the TLB instance virtual service sends to the real service next-hop server.	none
Packet Recv	Number of packets returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	none
Byte Recv	Number of bytes returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	none
Traffic load balance real svc name	Name of the real service used for traffic load-balancing.	extensive

Table 110: show services traffic-load-balance statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Routing instance name	Name of the routing instance on which the real service is configured.	extensive
IP address	IP address of the configured real service.	extensive
Traffic load balance group name	Name of the server group for real service.	extensive
Admin state	Administrative state of the real service, such as Up or Down .	extensive
Oper state	Operational state of the real service, such as Up or Down .	extensive
Network monitoring probe up count	Number of probes for which the status of the server whose health is checked is observed to be up. If a server group is configured for dual health check, a real service is declared to be UP only if both health-check probes are simultaneously UP; otherwise a real service declared to be DOWN.	extensive
Network monitoring probe down count	Number of probes for which the status of the server whose health is checked is observed to be down.	extensive
Total rejoin event count	Number of events that caused a server that was previously down and later operational to rejoin a group of real services for load-balancing.	extensive
Total up event count	Number of TLB events that identified a virtual service or real service to be up.	extensive
Total down event count	Number of TLB events that identified a virtual service or real service to be down.	extensive
Real Service packet sent count	Number of packets originating from the clients that the TLB instance virtual service sends to the real service.	extensive
Real Service byte sent count	Number of bytes originating from the clients that the TLB instance virtual service sends to the real service next-hop server.	extensive

Table 110: show services traffic-load-balance statistics Output Fields (*continued*)

Field Name	Field Description	Level of Output
Real Service packet received count	Number of packets returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	extensive
Real Service byte received count	Number of bytes returning from the real service next-hop server that the TLB instance virtual service processes and forwards to the clients.	extensive
Total probe sent	Number of health-monitoring probes sent from the TLB health check daemon.	extensive
Total probe success	Number of health-monitoring probes sent from the TLB health check daemon that were successful.	extensive
Total probe fail	Number of health-monitoring probes attempted to be sent from the TLB health check daemon that failed.	extensive
Total probe sent fail	Number of health-monitoring probes attempted to be sent from the TLB health check daemon that were unsuccessfully initiated.	extensive
Probe state	Status of the health-check probe, such as Up or Down .	extensive
Probe sent	Number of health-check probe requests transmitted from the TLB health check daemon.	extensive
Probe success	Number of successful health-check probe requests transmitted from the TLB health check daemon.	extensive
Probe fail	Number of failed health-check probe requests transmitted from the TLB health check daemon.	extensive
Probe sent failed	Number of times the TLB health check daemon was unable to initiate transmission of a extensive health-check probe.	extensive
Probe consecutive success	Number of health-check probe requests transmitted from the TLB health check daemon that were consecutively successful.	extensive
Probe consecutive fail	Number of health-check probe requests transmitted from the TLB health check daemon that failed for two successive times.	extensive

Sample Output

show services traffic-load-balance statistics

user@host> **show services traffic-load-balance statistics**

```

Traffic load balance instance name      : lb1
Multi services interface name          : ms-3/0/0
Interface state                        : UP
Interface type                         : Multi services
Route hold timer                      : 180
Active real service count              : 0
Total real service count              : 100
Traffic load balance virtual svc name  : v1
IP address                            : 0.0.0.0
Virtual service mode                   : Layer-2 based Direct Server Return mode
Routing instance name                  : internal-client-vrf
Traffic load balance group name        : g1
Health check interface subunit         : 40
Demux Nexthop index                   : N/A
Nexthop index                         : 840
Up time                               : 2d 19:09
Real Server Up count                   : 1
Real Server Down count                 : 1
Total packet sent count                : 0
Total byte sent count                  : 0

```

Real service	Address	Sts	Packet	Sent	Byte	Sent	Packet	Recv	Byte	Recv
r11	203.0.113.11	UP	0		0		0		0	
r10	203.0.113.10	UP	0		0		0		0	

```

Traffic load balance virtual svc name  : v2
IP address                            : 192.0.2.11
Virtual service mode                   : Translate mode
Routing instance name                  : msp-tproxy-forwarding1
Traffic load balance group name        : g2
Health check interface subunit         : 50
Protocol                              : tcp
Port number                           : 8080
Server Listening Port Number            : 8084
Demux Nexthop index                   : 536
Nexthop index                         : 539
Up time                               : 2d 19:06
Total packet sent count                : 0
Total byte sent count                  : 0

```

```

Total packet received count      : 0
Total byte received count        : 0
Real service    Address      Sts  Packet Sent  Byte Sent  Packet Recv  Byte Recv
r12             203.0.113.12  UP   0           0           0           0
r13             203.0.113.13  UP   0           0           0           0

```

show services traffic-load-balance statistics extensive

user@host> show services traffic-load-balance statistics extensive

```

Traffic Load Balance General Information
    DFW client-id                : 39

Traffic load balance instance name : lb1
Multi services interface name      : ms-3/0/0
Interface state                    : UP
Interface type                     : Multi services
Route hold timer                   : 180
Active real service count          : 0
Total real service count           : 100
Traffic load balance virtual svc name : v1
IP address                         : 0.0.0.0
Virtual service mode                : Layer-2 based Direct Server Return mode
Routing instance name              : internal-client-vrf
Traffic load balance group name    : g1
Traffic load balance group warmup time: 15
Traffic load balance group auto-rejoin: TRUE
Health check interface subunit     : 40
Traffic load balance group down count : 1
Route metric                       : 1
Virtual service down count         : 1
Traffic load balance hash method   : source
Network monitoring profile count   : 1
Active real service count          : 2
Total real service count           : 2
Demux Nexthop index               : N/A
Nexthop index                      : 840
Up time                            : 2d 19:09
Total packet sent count            : 0
Total byte sent count              : 0
Total packet received count        : 0
Total byte received count          : 0

```

```

Network monitoring profile index      : 1
Network monitoring profile name      : prof1
Probe type                           : ICMP
Probe interval                       : 5
Probe failure retry count            : 5
Probe recovery retry count           : 3

Traffic load balance real svc name   : r11
Routing instance name                : server-vrf10
IP address                           : 203.0.113.11
Traffic load balance group name      : g1
Admin state                          : UP
Oper state                           : UP
Network monitoring probe up count    : 1
Network monitoring probe down count  : 0
Total rejoin event count             : 0
Total up event count                 : 1
Total down event count               : 0
Real Service packet sent count       : 0
Real Service byte sent count         : 0
Total probe sent                     : 47939
Total probe success                   : 47918
Total probe fail                     : 21
Total probe sent failed              : 0
Network monitoring profile index      : 1
Network monitoring profile name      : prof1
Probe type                           : ICMP
Probe state                          : UP
Probe sent                           : 47939
Probe success                         : 47918
Probe fail                           : 21
Probe sent failed                    : 0
Probe consecutive success             : 10090
Probe consecutive fail                : 0

Traffic load balance real svc name   : r10
Routing instance name                : server-vrf10
IP address                           : 203.0.113.10
Traffic load balance group name      : g1
Admin state                          : UP
Oper state                           : UP
Network monitoring probe up count    : 1
Network monitoring probe down count  : 0
Total rejoin event count             : 0

```

```

Total up event count          : 1
Total down event count        : 0
Real Service packet sent count      : 0
Real Service byte sent count      : 0
Total probe sent              : 47939
Total probe success            : 47917
Total probe fail              : 22
Total probe sent failed        : 0
Network monitoring profile index    : 1
Network monitoring profile name    : prof1
Probe type                     : ICMP
Probe state                    : UP
Probe sent                     : 47939
Probe success                   : 47917
Probe fail                     : 22
Probe sent failed              : 0
Probe consecutive success       : 10090
Probe consecutive fail          : 0

Traffic load balance virtual svc name : v2
IP address                          : 192.0.2.11
Virtual service mode                : Translate mode
Routing instance name               : msp-tproxy-forwarding1
Traffic load balance group name     : g2
Traffic load balance group warmup time: 15
Traffic load balance group auto-rejoin: TRUE
Health check interface subunit      : 50
Traffic load balance group down count : 1
Protocol                           : tcp
Port number                        : 8080
Server Listening Port Number         : 8084
Route metric                        : 1
Virtual service down count          : 1
Traffic load balance hash method    : source-destination
Network monitoring profile count    : 1
Active real service count           : 2
Total real service count            : 2
Demux Nexthop index                 : 536
Nexthop index                       : 539
Up time                            : 2d 19:07
Total packet sent count             : 0
Total byte sent count               : 0
Total packet received count         : 0
Total byte received count           : 0

```

```

Network monitoring profile index      : 1
Network monitoring profile name      : prof1
Probe type                           : ICMP
Probe interval                       : 5
Probe failure retry count            : 5
Probe recovery retry count           : 3

Traffic load balance real svc name   : r12
Routing instance name                : server-vrf10
IP address                           : 203.0.113.12
Traffic load balance group name      : g2
Admin state                          : UP
Oper state                           : UP
Network monitoring probe up count    : 1
Network monitoring probe down count  : 0
Total rejoin event count             : 0
Total up event count                 : 1
Total down event count               : 0
Real Service packet sent count       : 0
Real Service byte sent count         : 0
Real Service packet received count   : 0
Real Service byte received count     : 0
Total probe sent                     : 47939
Total probe success                   : 47916
Total probe fail                     : 23
Total probe sent failed              : 0
Network monitoring profile index     : 1
Network monitoring profile name      : prof1
Probe type                           : ICMP
Probe state                          : UP
Probe sent                           : 47939
Probe success                         : 47916
Probe fail                           : 23
Probe sent failed                    : 0
Probe consecutive success             : 10089
Probe consecutive fail               : 0

Traffic load balance real svc name   : r13
Routing instance name                : server-vrf10
IP address                           : 203.0.113.13
Traffic load balance group name      : g2
Admin state                          : UP
Oper state                           : UP
Network monitoring probe up count    : 1

```

```

Network monitoring probe down count    : 0
Total rejoin event count              : 0
Total up event count                  : 1
Total down event count                 : 0
Real Service packet sent count         : 0
Real Service byte sent count           : 0
Real Service packet received count     : 0
Real Service byte received count       : 0
Total probe sent                      : 47939
Total probe success                    : 47910
Total probe fail                      : 29
Total probe sent failed                : 0
Network monitoring profile index       : 1
Network monitoring profile name        : prof1
Probe type                            : ICMP
Probe state                           : UP
Probe sent                            : 47939
Probe success                         : 47910
Probe fail                            : 29
Probe sent failed                     : 0
Probe consecutive success              : 6283
Probe consecutive fail                 : 0

```

show services traffic-load-balance statistics summary

```
user@host> show services traffic-load-balance statistics summary
```

```

Traffic load balance instance name    : tlb_sdg
Multi services interface name         : ms-8/3/0
Interface state                       : UP
Interface type                        : Multi services
Route hold timer                      : 180
Active real service count             : 0
Total real service count              : 100
Virtual service    Address            Sts Packet Sent Byte Sent   Packet Recv Byte
Recv
DNS-VIP1-TCP       198.51.100.1   Up  13182260    709736171  11951566   732469940
DNS-VIP1-UDP       198.51.100.1   Up  2683203     163675383  2683101    262943898
HTTP-80-ADDRESS-VIP 203.0.113.156 Up  363080548   25152313876 282072340
280409712450
HTTP-8080-ADDR-VIP 203.0.113.157 Up  363198700   25318638843 282030640
280388777065

```

Secure-Ent-443-VIP 203.0.113.158 Up 30561467 3012763619 28007583
3992807922

Simple-Ent-80-VIP 203.0.113.159 Up 155857682 11558785554 89649255
79217609518

Traffic load balance instance name : tlb_sdg_v6

Multi services interface name : ms-8/3/0

Interface state : UP

Interface type : Multi services

Route hold timer : 180

Virtual service	Address	Sts	Packet Sent	Byte Sent	Packet Recv	Byte Recv
-----------------	---------	-----	-------------	-----------	-------------	-----------

DNS-VIP1-TCP-V6	2001:db8:a::300	Up	25118146	1829085032	24172053	2088425092
-----------------	-----------------	----	----------	------------	----------	------------

DNS-VIP1-UDP-V6	2001:db8:a::300	Up	1318497	108116747	1319249	386274267
-----------------	-----------------	----	---------	-----------	---------	-----------

HTTP-80-ADDR-VIP-V6	2001:db8:a::100	Up	368696950	33051271152	282178604	287789935055
---------------------	-----------------	----	-----------	-------------	-----------	--------------

HTTP-8080-ADD-VIP-V6	2001:db8:a::100	Up	368797597	33217998028	281989122	287768684085
----------------------	-----------------	----	-----------	-------------	-----------	--------------

Sec-Ent-443-VIP-V6	2001:db8:a::200	Up	0662649	3622545250	28080924	4531356641
--------------------	-----------------	----	---------	------------	----------	------------

show services web-filter dns-resolution profile

Syntax

```
show services web-filter dns-resolution profile profile-name <template template-name>
<fpc-slot fpc-slot pic-slot pic-slot>
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display URL filter domain name system (DNS) resolution information.

URL filtering resolves the disallowed domains. The total number of domains are divided into chunks of 50 domains per chunk. The **filter term** in the command output is the name of a chunk.

Options

fpc-slot *fpc-slot* pic-slot *pic-slot*—(Optional) Specify the FPC and PIC for which you want URL filter information displayed.

profile *profile-name*—Specify the profile for which you want URL filter information displayed.

template *template-name*—(Optional) Specify the template for which you want URL filter information displayed.

Required Privilege Level

view

RELATED DOCUMENTATION

[show services web-filter dns-resolution-statistics profile template | 1090](#)

[show services web-filter statistics profile | 1102](#)

[Configuring URL Filtering | 335](#)

List of Sample Output

[show services web-filter dns-resolution profile on page 1087](#)

Output Fields

[Table 111 on page 1087](#) lists the output fields for the **show services web-filter dns-resolution profile** command. Output fields are listed in the approximate order in which they appear.

Table 111: show services web-filter dns-resolution profile Output Fields

Field Name	Field Description
Profile	Name of profile.
Template	Name of template.
Filter Term	Name of the domains chunk. All domains are divided into chunks of 50 domains per chunk.
IPv4 Address Count	The number of IPv4 addresses resolved for all domains under the filter term.
IPv6 Address Count	The number of IPv6 addresses resolved for all domains under the filter term.
Domain Name	Name of domain.
IPv4 Records	Listing of IPv4 addresses.
IPv6 Records	Listing of IPv6 addresses.

Sample Output

show services web-filter dns-resolution profile

user@host> **show services web-filter dns-resolution profile p1**

```

URL filtering DNS resolution:
Profile: p1
Template: t1

1). Filter Term: URLF_t1_0004

    IPv4 Address Count: 20
    IPv6 Address Count: 20

1 ). Domain Name: www.example.com

    IPv4 Records:
        31.13.77.36
        31.13.76.68

```

IPv6 Records:

2a03:2880:f122:83:face:b00c:0:25de
2a03:2880:f111:83:face:b00c:0:25de

2). Domain Name: www.youtube.com

IPv4 Records:

216.58.193.78
216.58.194.206

IPv6 Records:

2607:f8b0:400a:800::200e
2607:f8b0:4005:809::200e

3). Domain Name: www.netflix.com

IPv4 Records:

50.112.200.248
52.10.96.2
52.25.242.211
52.39.87.182
52.38.44.92
52.36.125.176
52.40.2.42
52.42.184.64
52.5.80.199
52.206.203.18
52.5.231.14
52.21.94.89
52.71.118.87
52.201.133.109
52.71.122.233
52.203.136.33

IPv6 Records:

2620:108:700f::342a:b840
2620:108:700f::3644:fc64
2620:108:700f::3459:2ce1
2620:108:700f::3459:c025
2620:108:700f::3459:f556
2620:108:700f::3459:c5c5
2620:108:700f::3644:c2a0
2620:108:700f::342a:df11

```
2406:da00:ff00::3404:d29c
2406:da00:ff00::3415:a86e
2406:da00:ff00::3415:fda4
2406:da00:ff00::3414:91d2
2406:da00:ff00::3403:73dd
2406:da00:ff00::22c7:d016
2406:da00:ff00::3400:290b
2406:da00:ff00::3213:c65f
```

show services web-filter dns-resolution-statistics profile template

Syntax

```
show services web-filter dns-resolution-statistics profile profile-name template template-name
(extensive | summary)
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display URL filter domain name system (DNS) resolution statistics.

Options

(extensive | summary)—Specify the level of detail of information you want displayed.

profile *profile-name*—Specify the profile for which you want URL filter information displayed.

template *template-name*—Specify the template for which you want URL filter information displayed.

Required Privilege Level

view

RELATED DOCUMENTATION

[show services web-filter dns-resolution profile | 1086](#)

[show services web-filter statistics profile | 1102](#)

[Configuring URL Filtering | 335](#)

List of Sample Output

[show services web-filter dns-resolution-statistics profile template summary on page 1092](#)

[show services web-filter dns-resolution-statistics profile template extensive on page 1093](#)

Output Fields

[Table 112 on page 1091](#) lists the output fields for the **show services web-filter dns-resolution-statistics profile template** command. Output fields are listed in the approximate order in which they appear.

Table 112: show services web-filter dns-resolution-statistics profile template Output Fields

Field Name	Field Description	Level of Detail
Profile	Name of profile.	all
Template	Name of template.	all
DNS start time	Start time of the DNS resolution.	summary
Next DNS start time	Start time of the next DNS resolution.	summary
Number of resolved A addresses	Number of resolved IPv4 addresses.	summary
Number of resolved AAAA addresses	Number of resolved IPv6 addresses.	summary
Number of unresolved A addresses	Number of unresolved IPv4 addresses.	summary
Number of unresolved AAAA addresses	Number of unresolved IPv6 addresses.	summary
Number of resolved A domains	Number of resolved IPv4 domains.	summary
Number of resolved AAAA domains	Number of resolved IPv6 domains.	summary
Number of unresolved A domains	Number of unresolved IPv4 domains.	summary
Number of unresolved AAAA domains	Number of unresolved IPv6 domains.	summary
Number of requests sent	Number of DNS requests sent.	summary
Number of responses received	Number of DNS responses received.	summary
Domain Name	Name of domain.	extensive

Table 112: show services web-filter dns-resolution-statistics profile template Output Fields (*continued*)

Field Name	Field Description	Level of Detail
IPv4 Address information	<p>IPv4 address information includes the following fields:</p> <ul style="list-style-type: none"> • DNS server IP—IPv4 address of DNS server. • Req Sent—Number of DNS requests sent. • Resp Received—Number of DNS responses received. • DNS retries—Number of times no DNS response was received and so retried. 	extensive
IPv6 Address information	<p>IPv6 address information includes the following fields:</p> <ul style="list-style-type: none"> • DNS server IP—IPv6 address of DNS server. • Req Sent—Number of DNS requests sent. • Resp Received—Number of DNS responses received. • DNS retries—Number of times no DNS response was received and so retried. 	extensive

Sample Output

show services web-filter dns-resolution-statistics profile template summary

user@host> **show services web-filter dns-resolution-statistics profile1 template t1 summary**

URL filtering DNS resolution statistics:

Profile: p1

Template: t1

DNS start time : May 01 16:40:24 PDT

Next DNS start time : May 01 17:40:24 PDT

Number of resolved A domains : 114

```

Number of resolved AAAA domains      : 114
Number of unresolved A domains       : 0
Number of unresolved AAAA domains    : 0
Number of requests sent               : 246
Number of responses received          : 228

```

show services web-filter dns-resolution-statistics profile template extensive

user@host> **show services web-filter dns-resolution-statistics profile p1 template t1 extensive**

URL filtering DNS resolution statistics:

Profile: p1

Template: t1

1) Domain Name: www.facebook.com

IPv4 Address information:

```

DNS server IP      8.8.8.8
Req Sent           20
Resp Received      20
DNS retries        0

```

IPv4 Address information:

```

DNS server IP      172.29.131.60
Req Sent           21
Resp Received      20
DNS retries        0

```

IPv6 Address information:

```

DNS server IP      8.8.8.8
Req Sent           25
Resp Received      20
DNS retries        0

```

IPv6 Address information:

```

DNS server IP      172.29.131.60
Req Sent           24
Resp Received      20
DNS retries        0

```

2) Domain Name: www.youtube.com

IPv4 Address information:

```

DNS server IP      8.8.8.8
Req Sent           21

```

```

Resp Received      20
DNS retries        0

```

IPv4 Address information:

```

DNS server IP      172.29.131.60
Req Sent           21
Resp Received      20
DNS retries        0

```

IPv6 Address information:

```

DNS server IP      8.8.8.8
Req Sent           21
Resp Received      20
DNS retries        0

```

IPv6 Address information:

```

DNS server IP      172.29.131.60
Req Sent           21
Resp Received      20
DNS retries        0

```

3) Domain Name: www.netflix.com

IPv4 Address information:

```

DNS server IP      8.8.8.8
Req Sent           21
Resp Received      20
DNS retries        0

```

IPv4 Address information:

```

DNS server IP      172.29.131.60
Req Sent           21
Resp Received      20
DNS retries        0

```

IPv6 Address information:

```

DNS server IP      8.8.8.8
Req Sent           21
Resp Received      20
DNS retries        0

```

IPv6 Address information:

```

DNS server IP      172.29.131.60
Req Sent           21

```


Resp Received	20
DNS retries	0

show services web-filter secintel-policy status

Syntax

```
show services web-filter secintel-policy status
profile profile-name
template template-name
```

Release Information

Statement introduced before Junos OS Release 18.4.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480, and MX960 with the MX-SPC3 services card.

Description

Display the IPv4 and IPv6 count per threat level received from the C&C feed from Policy Enforcer. It also displays the count of the number of terms used in the implicit filter per threat level.

Options

profile-name—Name of the profile

template-name—Name of the template

Required Privilege Level

view

RELATED DOCUMENTATION

[security-intelligence](#) | 665

List of Sample Output

[show services web-filter secintel-policy status on page 1096](#)

Sample Output

show services web-filter secintel-policy status

user@host> show services web-filter secintel-policy status profile

```
URL Filtering SecIntel Policy Status:
Profile      : Profile1
```

```

C&C DB File : /var/db/url-filterd/urlf_si_cc_db.txt
Policy State: Ready
DB File Change Time : Tue Nov 27 11:01:10 2018
DB File Load Time   : Tue Nov 27 11:01:38 2018
C&C Prefix Count    : IPv4: 11093      IPv6: 5
Filters:

```

Threat level	Action	v4 Term Count	IPv4	v6 Term Count	IPv6
1	ACCEPT	23	1129	1	2
2	ACCEPT	11	1444	0	0
3	ACCEPT	6	996	0	0
4	ACCEPT	7	564	0	0
5	ACCEPT	7	451	0	0
6	ACCEPT	4	126	0	0
7	LOG	5	175	0	0
8	DROP AND LOG	4	396	1	1
9	ACCEPT	2	164	0	0
10	ACCEPT	33	5601	1	2

Sample Output

```

user@host> show services web-filter secintel-policy-status profile Profile1 url-filter-template
template200

```

```

Template      : template200
  C&C DB File : /var/db/url-filterd/urlf_si_ip_white_list_db.txt
  Policy State: NA
  DB File Change Time : NA
  DB File Load Time   : NA
  C&C Prefix Count    : IPv4: 0      IPv6: 0

  C&C DB File : /var/db/url-filterd/urlf_si_ip_black_list_db.txt
  Policy State: NA
  DB File Change Time : NA
  DB File Load Time   : NA
  C&C Prefix Count    : IPv4: 0      IPv6: 0

  C&C DB File : /var/db/url-filterd/urlf_si_ip_custom_db.txt
  Policy State: Ready

```

DB File Change Time : Tue Feb 04 15:22:20 2020

DB File Load Time : Tue Feb 04 15:24:29 2020

C&C Prefix Count : IPv4: 16 IPv6: 0

Filters:

Count	Threat level	Action	v4 Term Count	IPv4	v6 Term
0	0	ACCEPT AND SAMPLE	0	0	0
0	255	DROP AND SAMPLE	0	0	0
0	1	DROP AND SAMPLE	1	11	0
0	2	ACCEPT	0	0	0
0	3	DROP AND SAMPLE	1	1	0
0	4	DROP AND SAMPLE	1	1	0
0	5	ACCEPT	0	0	0
0	6	ACCEPT	1	1	0
0	7	ACCEPT	1	1	0
0	8	DROP AND SAMPLE	0	0	0
0	9	ACCEPT	1	1	0
0	10	DROP AND SAMPLE	0	0	0
0					

show services web-filter statistics dns-filter-template

Syntax

```
show services web-filter statistics dns-filter-template template-name
```

Release Information

Description

Display statistics for DNS request filtering and URL filtering for the specified filter profile.

Options

dns-filter-template *template-name*—(Optional) Display statistics for the specified DNS filter template.

Required Privilege Level

view

RELATED DOCUMENTATION

- [DNS Request Filtering for Disallowed Website Domains | 311](#)
- [Configuring URL Filtering | 335](#)

List of Sample Output

[show services web-filter statistics dns-filter-template on page 1100](#)

Output Fields

[Table 113 on page 1099](#) lists the output fields for the **show services web-filter statistics profile** command. Output fields are listed in the approximate order in which they appear.

Table 113: show services web-filter statistics profile Output Fields

Field Name	Field Description
UDP DNS	Number of UDP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.
TCP DNS	Number of TCP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.

Sample Output

show services web-filter statistics dns-filter-template

user@host> **show services web-filter statistics dns-filter-template DNS_CUSTOMER-A**

```

    DNS filtering counters:

UDP DNS A req count           : 0
UDP DNS A resp count          : 0
UDP DNS A log only count      : 0
UDP DNS AAAA req count        : 0
UDP DNS AAAA resp count       : 0
UDP DNS AAAA log only count   : 0
UDP DNS MX req count          : 0
UDP DNS MX resp count         : 0
UDP DNS MX log only count     : 0
UDP DNS CNAME req count       : 0
UDP DNS CNAME resp count      : 0
UDP DNS CNAME log only count  : 0
UDP DNS SRV req count         : 0
UDP DNS SRV resp count        : 0
UDP DNS SRV resp count        : 0

UDP DNS SRV resp count        : 0
+ UDP DNS SRV Resp No Err count : 0
+ UDP DNS SRV Resp Resp Refused Err count : 0
UDP DNS SRV log only count    : 0
UDP DNS TXT req count         : 0
UDP DNS TXT resp count        : 0
UDP DNS TXT log only count    : 0
+ UDP DNS TXT Resp No Err count : 0
+ UDP DNS TXT Resp Resp Refused Err count : 0
UDP DNS ANY req count         : 0
UDP DNS ANY resp count        : 0
UDP DNS ANY log only count    : 0
UDP DNS MISC req count        : 0
UDP DNS MISC log only count   : 0
TCP DNS A req count           : 0
TCP DNS A resp count          : 0
TCP DNS A log only count      : 0
TCP DNS AAAA req count        : 0
TCP DNS AAAA resp count       : 0
TCP DNS AAAA log only count   : 0

```

```
TCP DNS MX req count           : 0
TCP DNS MX resp count          : 0
TCP DNS MX log only count      : 0
TCP DNS CNAME req count        : 0
TCP DNS CNAME resp count       : 0
TCP DNS CNAME log only count   : 0
TCP DNS SRV req count          : 0
TCP DNS SRV resp count         : 0
TCP DNS SRV log only count     : 0
+ TCP DNS SRV Resp No Err count : 0
+ TCP DNS SRV Resp Resp Refused Err count : 0

TCP DNS TXT req count          : 0
TCP DNS TXT resp count         : 0
TCP DNS TXT log only count     : 0
+ TCP DNS SRV Resp No Err count : 0
+ TCP DNS SRV Resp Resp Refused Err count : 0

TCP DNS ANY req count          : 0
TCP DNS ANY resp count         : 0
TCP DNS ANY log only count     : 0
TCP DNS MISC req count         : 0
TCP DNS MISC log only count    : 0
```

show services web-filter statistics profile

Syntax

```
show services web-filter statistics profile profile-name
<dns-filter-template template-name>
<dns-filter-term term-name>
<fpc-slot fpc-slot pic-slot pic-slot>
<url-filter-template template-name>
```

Release Information

Command introduced in Junos OS Release 18.3R1.

Support added in Junos OS Release 19.3R2 for Next Gen Services on MX Series routers MX240, MX480 and MX960 with the MX-SPC3 services card.

Description

Display statistics for DNS request filtering and URL filtering for the specified filter profile.

Options

dns-filter-template *template-name*—(Optional) Display statistics for the specified DNS filter template.

dns-filter-term *term-name*—(Optional) Display statistics for the specified term in the DNS filter template.

fpc-slot *fpc-slot* pic-slot *pic-slot*—(Optional) Display statistics for the specified services PIC.

profile *profile-name*—Display statistics for the specified filter profile.

url-filter-template *template-name*—(Optional) Display statistics for the specified URL filter template.

Required Privilege Level

view

RELATED DOCUMENTATION

[DNS Request Filtering for Disallowed Website Domains | 311](#)

[Configuring URL Filtering | 335](#)

List of Sample Output

[show services web-filter statistics profile dns-filter-template on page 1104](#)

[show services web-filter statistics profile on page 1105](#)

Output Fields

[Table 113 on page 1099](#) lists the output fields for the **show services web-filter statistics profile** command. Output fields are listed in the approximate order in which they appear.

Table 114: show services web-filter statistics profile Output Fields

Field Name	Field Description
UDP Counters	Number of UDP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.
TCP Counters	Number of TCP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.
Accept	Action counters for accepted packets for URL filtering.
Custom page	Action counters for custom page sent to recipient for URL filtering.
Http scode	Action counters for HTTP status code response for URL filtering.
Redirect url	Action counters for redirect URL response for URL filtering.
TCP reset	Action counters for TCP reset for URL filtering. Connection is closed.
Bypass session count	Number of sessions not blocked by URL filtering because the match criteria was not met for URL filtering.
IPV4 Disable IP Blocking	Action counters for IPv4 packets that were accepted because filtering is disabled for HTTP traffic that contains an embedded IP address belonging to a disallowed domain name in the URL filter database.
IPV6 Disable IP Blocking	Action counters for IPv6 packets that were accepted because filtering is disabled for HTTP traffic that contains an embedded IP address belonging to a disallowed domain name in the URL filter database.
session count	The session of activity that a user with a unique IP address spends on a website during a specified period of time for URL filtering. A session, in this case, would be the packets going to the service PIC from the Packet Forwarding Engine and then back to the service PIC.
uplink packet count	Number of packets going from the Packet Forwarding Engine to the service PIC for URL filtering.
uplink bytes	Number of bytes passing uplink for URL filtering.
downlink packet count	Number of packets going from the service PIC to the service Packet Forwarding Engine for URL filtering.
downlink bytes	Number of bytes passing downlink for URL filtering.

Table 114: show services web-filter statistics profile Output Fields (*continued*)

Field Name	Field Description
UDP DNS	Number of UDP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.
TCP DNS	Number of TCP DNS requests, responses, and log only responses for DNS request filtering for queries of types A, AAAA, MX, CNAME, SRV, TXT, and MISC.

Sample Output

show services web-filter statistics profile dns-filter-template

user@host> **show services web-filter statistics profile pdns dns-filter-template tdn**

Query Type	Requests	Responses	Log only
------------	----------	-----------	----------

UDP Counters:

A	0	0	0
AAAA	0	0	0
MX	0	0	0
CNAME	0	0	0
SRV	0	0	0
TXT	0	0	0
MISC	0	0	0

TCP Counters:

A	0	0	0
AAAA	0	0	0
MX	0	0	0
CNAME	0	0	0
SRV	0	0	0
TXT	0	0	0
MISC	0	0	0

Sample Output

show services web-filter statistics profile

user@host> **show services web-filter statistics profile Profile1**

URL filtering action counters:

Accept session count	: 0
Accept uplink packet count	: 0
Accept uplink bytes	: 0
Accept downlink packet count	: 0
Accept downlink bytes	: 0

Custom page session count	: 0
Custom page uplink packet count	: 0
Custom page uplink bytes	: 0
Custom page downlink packet count	: 0
Custom page downlink bytes	: 0

Http scode session count	: 0
Http scode uplink packet count	: 0
Http scode uplink bytes	: 0
Http scode downlink packet count	: 0
Http scode downlink bytes	: 0

Redirect url session count	: 0
Redirect url uplink packet count	: 0
Redirect url uplink bytes	: 0
Redirect url downlink packet count	: 0
Redirect url downlink bytes	: 0

Tcp reset session count	: 0
Tcp reset uplink packet count	: 0
Tcp reset uplink bytes	: 0
Tcp reset downlink packet count	: 0
Tcp reset downlink bytes	: 0

Bypass session count	: 0
----------------------	-----

IPV4 Disable IP Blocking Sessions	: 0
IPV4 Disable IP Blocking uplink packets	: 0
IPV4 Disable IP Blocking uplink bytes	: 0
IPV4 Disable IP Blocking downlink packets	: 0

```

IPV4 Disable IP Blocking downlink bytes      : 0
IPV6 Disable IP Blocking Sessions            : 0
IPV6 Disable IP Blocking uplink packets      : 0
IPV6 Disable IP Blocking uplink bytes        : 0
IPV6 Disable IP Blocking downlink packets    : 0
IPV6 Disable IP Blocking downlink bytes      : 0

```

DNS filtering counters:

```

UDP DNS A req count                          : 0
UDP DNS A resp count                         : 0
UDP DNS A log only count                     : 0
UDP DNS AAAA req count                       : 0
UDP DNS AAAA resp count                     : 0
UDP DNS AAAA log only count                  : 0
UDP DNS MX req count                         : 0
UDP DNS MX resp count                       : 0
UDP DNS MX log only count                   : 0
UDP DNS CNAME req count                     : 0
UDP DNS CNAME resp count                   : 0
UDP DNS CNAME log only count                 : 0
UDP DNS SRV req count                       : 0
UDP DNS SRV resp count                     : 0
UDP DNS SRV log only count                  : 0
UDP DNS TXT req count                       : 0
UDP DNS TXT resp count                     : 0
UDP DNS TXT log only count                  : 0
UDP DNS ANY req count                       : 0
UDP DNS ANY resp count                     : 0
UDP DNS ANY log only count                  : 0
UDP DNS MISC req count                      : 0
UDP DNS MISC log only count                 : 0
TCP DNS A req count                         : 0
TCP DNS A resp count                       : 0
TCP DNS A log only count                    : 0
TCP DNS AAAA req count                     : 0
TCP DNS AAAA resp count                    : 0
TCP DNS AAAA log only count                 : 0
TCP DNS MX req count                       : 0
TCP DNS MX resp count                      : 0
TCP DNS MX log only count                  : 0
TCP DNS CNAME req count                    : 0
TCP DNS CNAME resp count                   : 0
TCP DNS CNAME log only count               : 0

```

TCP DNS SRV req count	: 0
TCP DNS SRV resp count	: 0
TCP DNS SRV log only count	: 0
TCP DNS TXT req count	: 0
TCP DNS TXT resp count	: 0
TCP DNS TXT log only count	: 0
TCP DNS ANY req count	: 0
TCP DNS ANY resp count	: 0
TCP DNS ANY log only count	: 0
TCP DNS MISC req count	: 0
TCP DNS MISC log only count	: 0

show system unified-services status

Syntax

```
show system unified-services status
```

Release Information

Command introduced in Junos OS Release 19.3R1 on MX Series.

Description

Determine whether Next Gen Services is enabled or disabled on the MX.

Required Privilege Level

List of Sample Output

[show system unified-services status on page 1108](#)

Output Fields

When you enter this command, you are provided feedback on the status of your request.

Sample Output

show system unified-services status

```
user@host> show system unified-services status
```

One of the following four messages appears:

```
Enabled
```

```
Unified Services : Upgrade staged , please reboot with 'request system reboot' to  
enable unified services.
```

```
Disabled
```

```
Unified Services : Upgrade staged , please reboot with 'request system reboot' to  
disable unified services.
```