



Junos[®] OS

Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide

Release

14.1



Published: 2014-08-18

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

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

Junos[®] OS Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide

14.1

Copyright © 2014, 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 <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

	About the Documentation	xiii
	Documentation and Release Notes	xiii
	Supported Platforms	xiii
	Using the Examples in This Manual	xiii
	Merging a Full Example	xiv
	Merging a Snippet	xiv
	Documentation Conventions	xv
	Documentation Feedback	xvii
	Requesting Technical Support	xvii
	Self-Help Online Tools and Resources	xvii
	Opening a Case with JTAC	xviii
Part 1	Overview	
Chapter 1	Junos Address Aware Network Addressing	3
	Junos Address Aware Network Addressing Overview	3
	Sample IPv6 Transition Scenarios	3
	Example 1: IPv4 Depletion with a Non-IPv6 Access Network	4
	Example 2: IPv4 Depletion with an IPv6 Access Network	4
	Example 3: IPv4 Depletion for Mobile Networks	5
Chapter 2	Carrier-Grade NAT Solutions	7
	Junos OS Carrier-Grade NAT Implementation Overview	7
	Port Control Protocol Overview	8
	Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card	9
	ALGs Available by Default for Junos OS Address Aware NAT	12
Chapter 3	Tunneling Solutions	15
	Tunneling Services for IPv4-to-IPv6 Transition Overview	15
	6to4 Overview	15
	Basic 6to4	16
	6to4 Anycast	16
	6to4 Provider-Managed Tunnels	17
	DS-Lite Softwires—IPv4 over IPv6	17
	6rd Softwires—IPv6 over IPv4	18
Part 2	Configuration	
Chapter 4	NAT Configuration Concepts	23
	Network Address Translation Configuration Overview	23

Chapter 5	NAT Configuration Tasks	25
	Configuring Static Source Translation in IPv4 Networks	25
	Configuring the NAT Pool and Rule	26
	Configuring the Service Set for NAT	27
	Configuring Trace Options	28
	Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range	29
	Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet	30
	Configuring Static Source Translation in IPv6 Networks	31
	Configuring the NAT Pool and Rule	31
	Configuring the Service Set for NAT	32
	Configuring Trace Options	33
	Configuring Static Destination Address Translation in IPv4 Networks	35
	Configuring Dynamic Address-Only Source Translation in IPv4 Networks	39
	Configuring Dynamic Source Address and Port Translation in IPv4 Networks	43
	Configuring Dynamic Source Address and Port Translation for IPv6 Networks	47
	Configuring Secured Port Block Allocation	49
	Configuring Deterministic Port Block Allocation	51
	Configuring Stateful NAT64	52
	Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT	54
	Configuring the DNS ALG Application	54
	Configuring the NAT Pool and NAT Rule	54
	Configuring the Service Set for NAT	58
	Configuring Trace Options	58
	Example: Assigning Addresses from a Dynamic Pool for Static Use	60
	Example: Configuring NAT for Multicast Traffic	61
	Rendezvous Point Configuration	61
	Router 1 Configuration	64
	Configuring Port Forwarding for Static Destination Address Translation	65
	Configuring Port Forwarding Without Destination Address Translation	68
	Example: Configuring Port Forwarding with Twice NAT	69
	Configuring Port Control Protocol	71
	Configuring PCP Server Options	71
	Configuring a PCP Rule	72
	Configuring a Service Set to Apply PCP	73
	SYSLOG Message Configuration	73
Chapter 6	Carrier-Grade NAT Complete Configuration Examples	75
	Example: Configuring Basic NAT44	75
	Example: NAPT Configuration for the MS-MPC	77
	Example: Configuring NAT-PT	82
	Example: Configuring Inline Network Address Translation - Interface-Service Service Set	96
	Port Control Protocol Configuration Examples	104
	Example: Configuring Port Control Protocol with NAPT44	104

Chapter 7	Carrier-Grade NAT Implementation Best Practices	111
	Carrier-Grade NAT Implementation: Best Practices	111
	Use APP and Round-Robin Address-Allocation	111
	Do Not Use EIM with SIP	112
	Do Not Use EIM with HTTP, DNS, or When Not Needed	112
	Define PBA Blocks Based on User Profiles	113
	Do Not Change the PBA Configuration on Running Systems	114
	Do Not Allocate Excessively Large NAT Pools	115
	Configure the System Log for PBA Only When Needed	115
	Use Redundant Service PIC (RSP) Interfaces for Failover	117
	Contain the Effects of Missing IP Fragments	118
	Do Not Use Configurations Prone to Routing Loops	118
Chapter 8	NAT Configuration Statements	121
	address (Services NAT Pool)	122
	address-allocation	123
	address-range	123
	allow-overlapping-nat-pools	124
	app-mapping-timeout	124
	application-sets (Services NAT)	125
	applications (Services NAT)	125
	cgn-pic	126
	destination-address	126
	destination-address-range	127
	destination-pool	127
	destination-port range	128
	destination-prefix	128
	destination-prefix-list	129
	destined-port	129
	deterministic-port-block-allocation	130
	dns-alg-pool	131
	dns-alg-prefix	131
	ei-mapping-timeout	132
	elf-flow-limit	132
	from (Services NAT)	133
	ipv6-multicast-interfaces	134
	mapping-refresh	134
	mapping-timeout	135
	match-direction	135
	no-translation	136
	overload-pool	136
	overload-prefix	137
	pool	138
	port	139
	port-forwarding	140
	port-forwarding-mappings	140
	ports-per-session	141
	rule	142
	rule-set	143

	secure-nat-mapping	143
	secured-port-block-allocation	144
	server (pcp)	145
	services (NAT)	146
	service-set (Services)	147
	source-address (NAT)	149
	source-address-range	149
	source-pool	150
	source-prefix	150
	source-prefix-list	151
	syslog	151
	translated-port	152
	term	153
	then	154
	translated	155
	translation-type	156
Chapter 9	Softwire Configuration Tasks	159
	Configuring a DS-Lite Softwire Concentrator	159
	Configuring a 6rd Softwire Concentrator	160
	Configuring Stateful Firewall Rules for 6rd Softwire	161
	Configuring Softwire Rules	161
	Configuring Service Sets for Softwire	162
Chapter 10	Softwire Configuration Examples	165
	Example: Basic DS-Lite Configuration	165
	Example: Basic 6rd Configuration	171
	Example: Configuring DS-Lite and 6rd in the Same Service Set	176
Chapter 11	6to4 Configuration	183
	Configuring a 6to4 Provider-Managed Tunnel	183
Chapter 12	Softwire Configuration Statements	187
	ds-lite	188
	rule (Softwire)	189
	rule-set (Softwire)	189
	softwire-concentrator	190
	softwire-options	191
	softwire-rules	191
	v6rd	192
Part 3	Administration	
Chapter 13	Monitoring CGN and Softwire Tunnels	195
	Monitoring CGN, Stateful Firewall, and Softwire Flows	195
	Monitoring Stateful Firewall Conversations	196
	Monitoring Global Stateful Firewall Statistics	196
	Monitoring NAT Pool Usage	197
	Monitoring Port Control Protocol Operations	197
	Monitoring Softwire Statistics	199
	Ping and Traceroute for DS-Lite	201

Chapter 14	Logging	203
	Log Generation	203
	Configuring NAT Session Logs	204
Chapter 15	High Availability and Load Balancing	207
	Inter-Chassis High Availability for MS-MIC and MS-MPC	207
	Inter-Chassis High Availability for Stateful Firewall and NAT44 Overview (MS-MIC, MS-MPC)	207
	Configuring Inter-Chassis High Availability for Stateful Firewall and NAT44 (MS-MPC, MS-MIC)	208
	Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall (MS-MIC, MS-MPC)	209
	High Availability and Load Balancing for 6rd Softwires	219
	Load Balancing a 6rd Domain Across Multiple Services PICs	219
	Example: Load Balancing a 6rd Domain Across Multiple Services PICs	219
	Configuring High Availability for 6rd Using 6rd Anycast	224
Chapter 16	Protecting Against Denial of Service Attacks	225
	Protecting CGN Devices Against Denial of Service (DOS) Attacks	225
	Mapping Refresh Behavior	225
	EIF Inbound Flow Limit	225
	DS-Lite Subnet Limitation	225
	DS-Lite Per Subnet Limitation Overview	226
	Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks	226
Chapter 17	Network Address Translation Operational Mode Commands	229
	clear services inline nat pool	230
	clear services inline nat statistics	231
	clear services nat flows	232
	clear services nat mappings	233
	clear services nat mappings app	235
	clear services nat mappings eim	236
	clear services nat mappings pcp	238
	clear services nat statistics	240
	show services inline nat pool	241
	show services inline nat statistics	242
	show services nat ipv6-multicast-interfaces	243
	show services nat pool	245
	show services nat mappings	249
	show services nat statistics	253
	show services pcp statistics	262
	show services software	265
	show services software flows	266
	show services software statistics	269
	show services stateful-firewall conversations	275
	show services stateful-firewall flows	279
	show services stateful-firewall statistics	285

Part 4	Index	
	Index	297

List of Figures

Part 1	Overview	
Chapter 1	Junos Address Aware Network Addressing	3
	Figure 1: IPv4 Depletion Solution – IPv4 Access Network	4
	Figure 2: IPv4 Depletion Solution – IPv6 Access Network	5
Chapter 2	Carrier-Grade NAT Solutions	7
	Figure 3: Basic PCP NAPT44 Topology	8
	Figure 4: PCP with DS-Lite Plain Mode	9
	Figure 5: PCP with DS-Lite Tunnel Mode	9
Chapter 3	Tunneling Solutions	15
	Figure 6: 6rd Software Flow	18
Part 2	Configuration	
Chapter 5	NAT Configuration Tasks	25
	Figure 7: Configuring NAT for Multicast Traffic	61
Chapter 6	Carrier-Grade NAT Complete Configuration Examples	75
	Figure 8: Configuring DNS ALGs with NAT-PT Network Topology	83
	Figure 9: Deploy Inline NAT within L3VPN	97
	Figure 10: PCP with NAPT44	104
Chapter 10	Software Configuration Examples	165
	Figure 11: DS-Lite Topology	166
Part 3	Administration	
Chapter 15	High Availability and Load Balancing	207
	Figure 12: Inter-Chassis High Availability Topology	208
	Figure 13: Inter-Chassis High Availability Topology	210

List of Tables

	About the Documentation	xiii
	Table 1: Notice Icons	xv
	Table 2: Text and Syntax Conventions	xvi
Part 1	Overview	
Chapter 2	Carrier-Grade NAT Solutions	7
	Table 3: Carrier-Grade NAT—Feature Comparison by Platform	10
	Table 4: Carrier-Grade NAT Translation Types	11
	Table 5: ALGs Available by Default	12
Part 3	Administration	
Chapter 17	Network Address Translation Operational Mode Commands	229
	Table 6: clear services nat flows Output Fields	232
	Table 7: clear services nat mappings Output Fields	233
	Table 8: clear services nat mappings app Output Fields	235
	Table 9: clear services nat mappings eim Output Fields	236
	Table 10: clear services nat mappings pcp Output Fields	238
	Table 11: show services inline nat pool Output Fields	241
	Table 12: show services inline nat statistics Output Fields	242
	Table 13: show services nat ipv6-multicast-interfaces Output Fields	243
	Table 14: show services nat pool Output Fields	245
	Table 15: show services nat mappings Output Fields	250
	Table 16: show services nat statistics Output Fields	253
	Table 17: show services pcp statistics Output Fields	262
	Table 18: show-services-softwire Output Fields	265
	Table 19: show services softwire flows Output Fields	266
	Table 20: command-name Output Fields	269
	Table 21: show services stateful-firewall conversations Output Fields	277
	Table 22: show services stateful-firewall flows Output Fields	281
	Table 23: show services stateful-firewall statistics Output Fields	285

About the Documentation

- Documentation and Release Notes on page xiii
- Supported Platforms on page xiii
- Using the Examples in This Manual on page xiii
- Documentation Conventions on page xv
- Documentation Feedback on page xvii
- Requesting Technical Support on page xvii

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 <http://www.juniper.net/techpubs/>.

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 <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- M Series
- T Series
- MX Series

Using the Examples in This Manual

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

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

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

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

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

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

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

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

Documentation Conventions

Table 1 on page xv 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 xvi 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>
<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 metric>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop address; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	

GUI Conventions

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

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

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC 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 <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

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

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

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

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

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

PART 1

Overview

- [Junos Address Aware Network Addressing on page 3](#)
- [Carrier-Grade NAT Solutions on page 7](#)
- [Tunneling Solutions on page 15](#)

CHAPTER 1

Junos Address Aware Network Addressing

- [Junos Address Aware Network Addressing Overview on page 3](#)
- [Sample IPv6 Transition Scenarios on page 3](#)

Junos Address Aware Network Addressing Overview

In early 2011, the Internet Assigned Numbers Authority (IANA) allocated the last large block of IPv4 addresses. Now service providers and large enterprises, as well as cloud providers, e-tailers, and federal agencies, are evaluating technologies to help them avoid IPv4 address exhaustion and ensure uninterrupted subscriber and service growth.

Junos Address Aware Network Addressing is Juniper Networks' portfolio of IPv4 exhaustion avoidance, IPv4-IPv6 coexistence, and IPv6 transition technologies that include IPv6, v4/v6 dual stack, NAT44, NAT44(4), NAPT44, NAPT444, NAT-PT, NAT64, 6-to4-PMT, 6rd, and DS-Lite. These technologies help network operators improve subscriber and service scale, mitigate IPv4 address depletion, and pragmatically transition to IPv6 based on business requirements.

Junos Address Aware Network Addressing technologies are available on the following platforms:

- MultiServices Dense Port Concentrator (MS-DPC)
- MS-100, MS-400, and MS-500 MultiServices PICS
- MultiServices Modular Port Concentrator (MS-MPC) and MultiServices Modular Interface Card (MS-MIC)
- Modular Port Concentrator Types 1, 2, and 3 (inline NAT).

Sample IPv6 Transition Scenarios

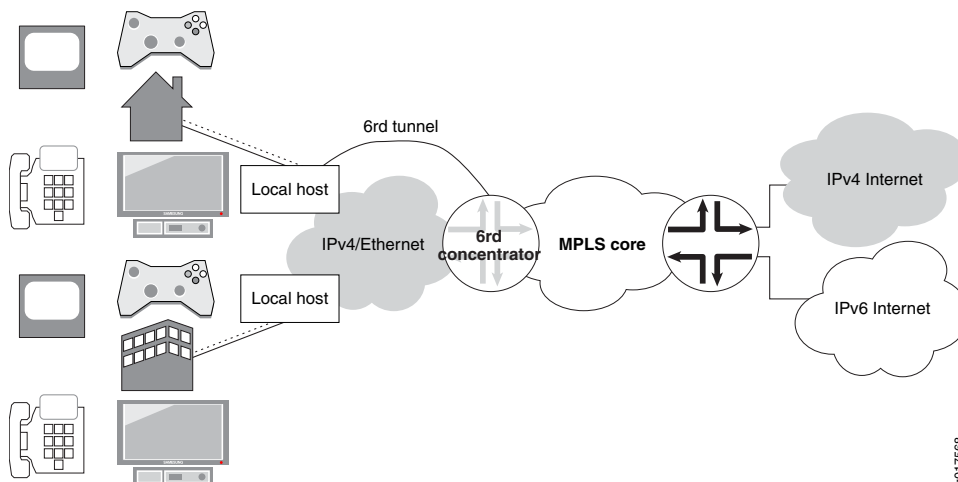
The Junos OS supports many IPv6 transition scenarios required by Junos OS customers. The following are selected examples:

- [Example 1: IPv4 Depletion with a Non-IPv6 Access Network on page 4](#)
- [Example 2: IPv4 Depletion with an IPv6 Access Network on page 4](#)
- [Example 3: IPv4 Depletion for Mobile Networks on page 5](#)

Example 1: IPv4 Depletion with a Non-IPv6 Access Network

Figure 1 on page 4 depicts a scenario in which the Internet service provider (ISP) has not significantly changed its IPv4 network. This approach enables IPv4 hosts to access the IPv4 Internet and IPv6 hosts to access the IPv6 Internet. A dual-stack host can be treated as an IPv4 host when it uses the IPv4 access service, and as an IPv6 host when it uses the IPv6 access service.

Figure 1: IPv4 Depletion Solution - IPv4 Access Network




Two new types of devices must be deployed in this approach: a dual-stack home gateway and a dual-stack carrier-grade Network Address Translation (NAT). The dual-stack home gateway integrates IPv4 forwarding and v6-over-v4 tunneling functions. It can also integrate a v4-v4 NAT function. The dual-stack carrier-grade NAT (CGN) integrates v6-over-v4 tunneling and carrier-grade v4-v4 NAT functions.

Example 2: IPv4 Depletion with an IPv6 Access Network

In the scenario shown in Figure 2 on page 5, the ISP network is IPv6-only.

Figure 2: IPv4 Depletion Solution - IPv6 Access Network



The dual-stack lite (DS-Lite) solution accommodates IPv6-only ISPs. The best business model for this approach is that the customer premises equipment (CPE) has integrated the functions for tunneling IPv6 to an IPv4 backbone, tunneling IPv4 to an IPv6 backbone, and can automatically detect which solution is required.

Not all customers of a given ISP must switch from IPv4 access to IPv6 access simultaneously; in fact, transition can be managed better by switching groups of customers (for example, all those connected to a single point of presence) on an incremental basis. Such an incremental approach should prove easier to plan, schedule, and execute than an across-the-board conversion.

Example 3: IPv4 Depletion for Mobile Networks

The complexity of mobile networks necessitates a flexible migration approach to ensure minimal disruption and maximum backward compatibility during transition. NAT64 can be used to enable IPv6 devices to communicate to IPv4 hosts without modifying the clients.

CHAPTER 2

Carrier-Grade NAT Solutions

- [Junos OS Carrier-Grade NAT Implementation Overview on page 7](#)
- [Port Control Protocol Overview on page 8](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 9](#)
- [ALGs Available by Default for Junos OS Address Aware NAT on page 12](#)

Junos OS Carrier-Grade NAT Implementation Overview

Junos OS enables you to implement and scale a Carrier-Grade Network Address Translation (CGNAT) solution based on the type of services interfaces used for your implementation:

- **MultiServices Denser Port Concentrator (MS-DPC)**—The layer 3 services package is used to configure NAT for MS-DPC adaptive services PICs. You must configure the layer-3 services package before implementing NAT on the MS-DPC. This solution provides the NAT functionality described in *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*.
- **MultiServices Modular Port Concentrator (MS-MPC) and MultiServices Modular Interface Card (MS-MIC)**—MS-MPCs and MS-MICs are pre-configured to enable configuration of carrier-grade NAT. This solution provides the NAT functionality described in *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*.
- **Inline NAT for Type 1, 2, and 3 Modular Port Concentrator (MPC Line Cards)**—Inline NAT leverages the services capabilities of TRIO-based MPC line cards, allowing a cost-effective implementation of NAT functionality on the data plane, as described in *Inline Network Address Translation Overview for MPC Types 1, 2, and 3*.

Related Documentation

- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 9](#)
- [Carrier-Grade NAT Implementation: Best Practices on page 111](#)
- [Example: Configuring Basic NAT44 on page 75](#)

Port Control Protocol Overview

The Port Control Protocol (PCP) provides a mechanism to control the forwarding of incoming packets by upstream devices such as NAT44, and firewall devices, and a mechanism to reduce application keep-alive traffic. PCP is designed to be implemented in the context of both Carrier-Grade NATs (CGNs) and small NATs (for example, residential NATs). PCP allows 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 Internet service provider. PCP allows 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 creating a mapping for incoming connections, it is necessary to inform remote computers about the IP address and port for the incoming connection. This is usually done in an application-specific manner.

PCP consists of the following components:

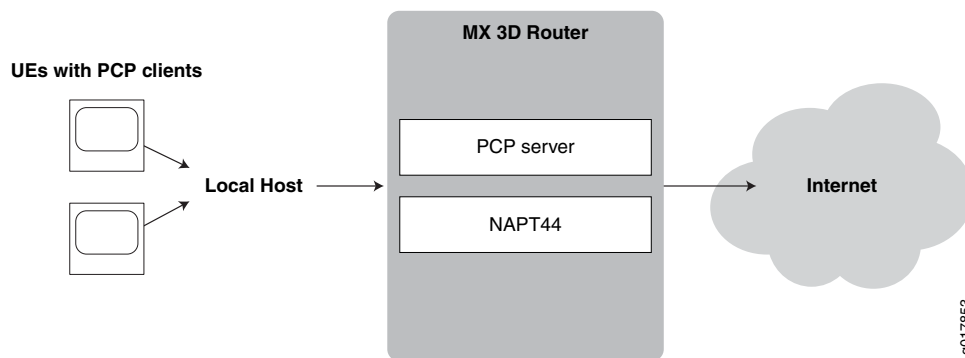
- PCP client—A host or gateway that issues PCP requests to a PCP server in order to obtain and control resources.
- PCP server—Typically a CGN gateway or co-located server that receives and processes PCP requests

Many NAT-friendly applications send frequent application-level messages to ensure their session are not be timed out by a NAT. These applications can reduce the frequency of such NAT keep-alive messages by using PCP to learn and influence the NAT mapping lifetime. This helps reduce bandwidth on the subscriber's access network, traffic to the server, and battery consumption on mobile devices.

The 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 UEs as shown in [Figure 3 on page 8](#).

Figure 3: Basic PCP NAPT44 Topology



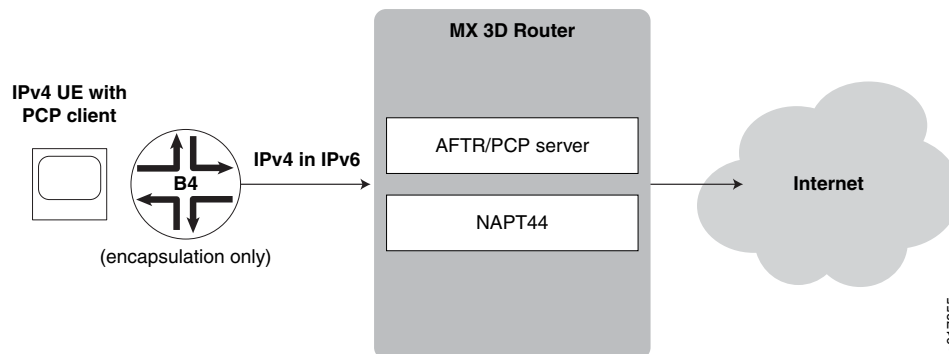
- Mapping of traffic containing PCP requests added by a router functioning as a DS-Lite software initiator (B4). This mode, known as *DS-Lite plain mode*, is shown in [Figure 4 on page 9](#)

Figure 4: PCP with DS-Lite Plain Mode



- Mapping of traffic containing PCP requests initiated directly by UEs and encapsulated by a router functioning as a DS-Lite software initiator (B4). This mode, known as *DS-Lite tunnel mode*, is shown in [Figure 5 on page 9](#).

Figure 5: PCP with DS-Lite Tunnel Mode



NOTE: The Junos OS does not support deterministic port block allocation for PCP-originated traffic.

Related Documentation

- [Configuring Port Control Protocol on page 71](#)

Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card

[Table 3 on page 10](#) summarizes feature differences among the Junos OS carrier-grade NAT implementations.

Table 3: Carrier-Grade NAT—Feature Comparison by Platform

Feature	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
Static Source NAT	yes	yes	yes
Dynamic Source NAT - Address Only	yes	yes	no
Dynamic Source NAT - NAT Port Translation with Secured Port Block Allocation	yes	no	no
Dynamic Source NAT - NAT Port Translation with Deterministic Port Block Allocation	yes	no	no
Static Destination NAT	yes	yes	yes <i>NOTE: Destination NAT can be implemented indirectly. See Inline Network Address Translation Overview for MPC Types 1, 2, and 3</i>
Twice NAT	yes	no	yes <i>NOTE: Twice NAT can be implemented indirectly. See Inline Network Address Translation Overview for MPC Types 1, 2, and 3</i>
NAPT - Preserve Parity and Port	yes	no	no
NAPT - EIM/EIF/APP	yes	yes	no
NAT64	yes	yes	no
NAT64 with APP/EIM/EIF	no	yes	no
NAT64 with ALGs	no	yes	no
<ul style="list-style-type: none"> • FTP • TFTP • SIP • RTSP • PPPT 			
DS-Lite	yes	no	no

Table 3: Carrier-Grade NAT—Feature Comparison by Platform (*continued*)

Feature	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
6rd	yes	no	no
Overload Pool/Overlap Address Across NAT Pool	yes	no	no
Port Control Protocol	yes	no	no
CGN-PIC	yes	no	no
AMS Support	no	yes	no

Table 4 on page 11 summarizes availability of translation types by type of line card.

Table 4: Carrier-Grade NAT Translation Types

Translation Type	MS-DPC MS-100 MS-400 MS-500	MS-MPC MS-MIC	MPC Types 1, 2, 3 <i>Inline NAT</i>
basic-nat44	yes	yes	yes
basic-nat66	yes	no	no
basic-nat-pt	yes	no	no
deterministic-napt44	yes	no	no
dnat-44	yes	yes	no
dynamic-nat44	yes	yes	no
napt-44	yes	yes	no
napt-66	yes	no	no
napt-pt	yes	no	no
stateful-nat64	yes	yes	no

Table 4: Carrier-Grade NAT Translation Types (*continued*)

Translation Type	MS-DPC		
	MS-100		
	MS-400	MS-MPC	MPC Types 1, 2, 3
	MS-500	MS-MIC	<i>Inline NAT</i>
<code>twice-basic-nat-44</code>	yes	no	no
<code>twice-dynamic-nat-44</code>	yes	no	no
<code>twice-dynamic-napt-44</code>	yes	no	no

Related Documentation • [Junos OS Carrier-Grade NAT Implementation Overview on page 7](#)

ALGs Available by Default for Junos OS Address Aware NAT

The following application-level gateways (ALGs) listed in [Table 5 on page 12](#) are supported for NAT processing on the listed platforms.

To view the implementation details (port, protocol, and so on) for these Junos OS default applications, locate the Junos OS Default ALG Name in the table and then look up the listed name in the **groups**. For example, for details about TFTP, look up **junos-tftp** as shown.



TIP: The Junos OS provides the `junos-alg`, which enables other ALGs to function by handling ALG registrations, causing slow path packets to flow through registered ALGs, and transferring ALG events to the ALG plug-ins. The `junos-alg` ALG is automatically available on the MS-MPC and MS-MIC platforms and does not require further configuration.

```
user@host# show groups junos-defaults applications application junos-tftp
application-protocol tftp;
protocol udp;
destination-port 69;
```

Table 5: ALGs Available by Default

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
Basic TCP ALG	yes	yes	NOTE: Specific Junos ALGs are not supported. However, a feature called TCP tracker, available by default, performs segment ordering and retransmit and connection tracking, validations for TCP connections.
Basic UDP ALG	yes	yes	NOTE: TCP tracker performs limited integrity and validation checks for UDP.

Table 5: ALGs Available by Default (*continued*)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
BOOTP	yes	no	<ul style="list-style-type: none"> • junos-bootpc • junos-bootps
DCE RPC Services	yes	yes	<ul style="list-style-type: none"> • junos-dce-rpc-portmap • junos-dcerpc-endpoint-mapper-service • junos-dcerpc-msexchange-directory-nsp • junos-dcerpc-msexchange-directory-rfr • junos-dcerpc-msexchange-information-store
DNS	yes	yes	<ul style="list-style-type: none"> • junos-dns-tcp • junos-dns-udp
FTP	yes	yes	<ul style="list-style-type: none"> • junos-ftp
H323	yes	no	<ul style="list-style-type: none"> • junos-h323
ICMP	yes	yes NOTE: ICMP messages are handled by default, but PING ALG support is not provided.	<ul style="list-style-type: none"> • junos-icmp-all • junos-icmp-ping
IIOp	yes	no	<ul style="list-style-type: none"> • junos-iiop-java • junos-iiop-orbix
IP	yes	The TCP tracker, available by default on these platforms, performs limited integrity and validation checks.	<ul style="list-style-type: none"> • junos-ip
NETBIOS	yes	no	<ul style="list-style-type: none"> • junos-netbios-datagram • junos-netbios-name-tcp • junos-netbios-name-udp • junos-netbios-session
NETSHOW	yes	no	<ul style="list-style-type: none"> • junos-netshow
PPTP	yes	yes	<ul style="list-style-type: none"> • junos-pptp
REALAUDIO	yes	no	<ul style="list-style-type: none"> • junos-realaudio
Sun RPC and RPC Port Map Services	yes	yes	<ul style="list-style-type: none"> • junos-rpc-portmap-tcp • junos-rpc-portmap-udp
RTSP	yes	yes	<ul style="list-style-type: none"> • junos-rtsp

Table 5: ALGs Available by Default (*continued*)

ALG	MS-DPC	MS-MPC, MS-MIC	Junos OS Default ALG Name
SIP	yes	Yes	<ul style="list-style-type: none"> • junos-sip <p>The SIP callid is <i>not</i> translated in register messages.</p> <p>NOTE: SIP sessions are limited to 12 hours (720 minutes) for NAT processing on the MS-MIC and MS-MPC interface cards. There is no time limit for SIP sessions on the MS-DPC.</p>
SNMP	yes	No	<ul style="list-style-type: none"> • junos-snmp-get • junos-snmp-get-next • junos-snmp-response junos-snmp-trap
SQLNET	yes	yes	<ul style="list-style-type: none"> • junos-sqlnet
TFTP	yes	yes	<ul style="list-style-type: none"> • junos-tftp
Traceroute	yes	no	<ul style="list-style-type: none"> • junos-traceroute
Unix Remote Shell Service	yes	Yes	<ul style="list-style-type: none"> • junos-rsh
WINFrame	yes	No	<ul style="list-style-type: none"> • junos-citrix-winframe • junos-citrix-winframe-udp
TALK-UDP	No	Yes	<ul style="list-style-type: none"> • junos-talk-udp
MS RPC	No	Yes	<ul style="list-style-type: none"> • junos-rpc-portmap-tcp • junos-rpc-portmap-udp • junos-rpc-services-tcp • junos-rpc-services-udp

Related Documentation • *ALG Descriptions*

CHAPTER 3

Tunneling Solutions

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)

Tunneling Services for IPv4-to-IPv6 Transition Overview

The 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 softwares, making it a very light and scalable solution. When you use softwares, you need not maintain an interface infrastructure for each software, unlike a typical mesh of generic routing encapsulation (GRE) tunnels that would require 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 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 softwares.

Software addresses are not specifically configured under any physical or virtual interface. Therefore, the number of established softwares does not affect throughput, and scalability is independent of the number of interfaces. The scalability is only limited to the number of flows that the platform (services DPC or PIC) can support.

This topic contains the following sections:

- [6to4 Overview on page 15](#)
- [DS-Lite Softwares—IPv4 over IPv6 on page 17](#)
- [6rd Softwares—IPv6 over IPv4 on page 18](#)

6to4 Overview

- [Basic 6to4 on page 16](#)
- [6to4 Anycast on page 16](#)
- [6to4 Provider-Managed Tunnels on page 17](#)

Basic 6to4

6to4 is an Internet transition mechanism for migrating from IPv4 to IPv6, a system that allows IPv6 packets to be transmitted over an IPv4 network (generally the IPv4 Internet) without the need to configure explicit tunnels. 6to4 is described in RFC 3056, *Connection of IPv6 Domains via IPv4 Clouds*. 6to4 is especially relevant during the initial phases of deployment to full, native IPv6 connectivity, since IPv6 is not required on nodes between the host and the destination. However, it is intended only as a transition mechanism and is not meant to be used permanently.

6to4 can be used by an individual host, or by a local IPv6 network. When used by a host, it must have a global IPv4 address connected, and the host is responsible for the encapsulation of outgoing IPv6 packets and the decapsulation of incoming 6to4 packets. If the host is configured to forward packets for other clients, often a local network, it is then a router.

There are two kinds of 6to4 virtual routers: border routers and relay routers. A 6to4 border router is an IPv6 router supporting a 6to4 pseudointerface, and is normally the border router between an IPv6 site and a wide-area IPv4 network. A relay router is a 6to4 router configured to support transit routing between 6to4 addresses and pure native IPv6 addresses.

In order for a 6to4 host to communicate with the native IPv6 Internet, its IPv6 default gateway must be set to a 6to4 address which contains the IPv4 address of a 6to4 relay router. To avoid the need for users to set this up manually, the Anycast address of 192.88.99.1 has been allocated to send packets to a 6to4 relay router. Note that when wrapped in 6to4 with the subnet and hosts fields set to zero, this IPv4 address (192.88.99.1) becomes the IPv6 address 2002:c058:6301:: To ensure BGP routing propagation, a short prefix of 192.88.99.0/24 has been allocated for routes pointed at 6to4 relay routers that use this Anycast IP address. Providers willing to provide 6to4 service to their clients or peers should advertise the Anycast prefix like any other IP prefix, and route the prefix to their 6to4 relay.

Packets from the IPv6 Internet to 6to4 systems must be sent to a 6to4 relay router by normal IPv6 routing methods. The specification states that such relay routers must only advertise 2002::/16 and not subdivisions of it to prevent IPv4 routes from polluting the routing tables of IPv6 routers. From there they can then be sent over the IPv4 Internet to the destination.

6to4 Anycast

Router 6to4 assumes that 6to4 routers and relays are managed and configured cooperatively. In particular, 6to4 sites must configure a relay router to carry the outbound traffic, which becomes the default IPv6 router (except for 2002::/16). The objective of the Anycast variant, defined in RFC 3068, *An Anycast Prefix for 6to4 Relay Routers*, is to avoid the need for such configuration. This makes the solution available for small or domestic users, even those with a single host or simple home gateway instead of a border router. This is achieved by defining 192.88.99.1 as the default IPv4 address for a 6to4 relay, and 2002:c058:6301:: as the default IPv6 router prefix ("well-known prefix") for a 6to4 site.

RFC 6343, *Advisory Guidelines for 6to4 Deployment*, published in August 2011, identifies a wide range of problems associated with the use of unmanaged 6to4 Anycast relay routers.

6to4 Provider-Managed Tunnels

A solution to many problems associated with unmanaged Anycast 6to4 is presented in IETF informational draft draft-kuarsingh-v6ops-6to4-provider-managed-tunnel-02, *6to4 Provider-Managed Tunnels (PMT)*. That document, a “work in progress,” proposes a solution that allows providers to exercise greater control over the routing of 6to4 traffic.

Anycast 6to4 implies a default configuration for the user site. It does not require any particular user action. It does require an IPv4 Anycast route to be in place to a relay at 192.88.99.1. Traffic does not necessarily return to the same 6to4 gateway because of the the “well-known” 6to4 prefix used and advertised by all 6to4 traffic.

6to4 provider-managed tunnels (PMTs) facilitate the management of 6to4 tunnels using an Anycast configuration. 6to4 PMT enables service providers to improve 6to4 operation when network conditions provide suboptimal performance or break normal 6to4 operation. 6to4 PMT provides a stable provider prefix and forwarding environment by utilizing existing 6to4 relays with an added function of IPv6 prefix translation that controls the flow of return traffic.

The 6to4 managed tunnel model behaves like a standard 6to4 service between the customer IPv6 host or gateway and the 6to4 PMT relay (within the provider domain). The 6to4-PMT relay shares properties with 6rd (RFC5969) by decapsulating and forwarding embedded IPv6 flows, within an IPv4 packet, to the IPv6 Internet. The model provides an additional function which translates the source 6to4 prefix to a provider assigned prefix which is not found in 6rd (RFC5969) or traditional 6to4 operation. The 6to4-PMT relay provides a stateless (or stateful) mapping of the 6to4 prefix to a provider-supplied prefix by mapping the embedded IPv4 address in the 6to4 prefix to the provider prefix.

DS-Lite Softwires—IPv4 over IPv6

When an Internet service provider (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 (CE) 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.

DS-Lite is supported on Multiservices 100, 400, and 500 PICs on M Series routers and on MX Series routers equipped with Multiservices Dense Port Concentrator (DPCs).



NOTE: IPv6 Provider Edge (6PE), or MPLS-enabled IPv6, is available for ISPs with MPLS-enabled networks. These networks now can use Multiprotocol BGP (MP-BGP) to provide connectivity between the DS-Lite B4 and AFTR (or any two IPv6 nodes). DS-Lite properly handles encapsulation and decapsulation despite the presence of additional MPLS header information.

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.

6rd Softwires—IPv6 over IPv4

6rd softwire flow is shown in [Figure 6 on page 18](#).

Figure 6: 6rd Softwire Flow



The Junos OS supports a 6rd softwire concentrator on a services DPC or PIC to facilitate rapid deployment of IPv6 service to subscribers on native IPv4 CE WANs. IPv6 packets are encapsulated in IPv4 packets by a softwire initiator at the CE WAN. These packets are tunneled to a softwire concentrator residing on a multiservices DPC (branch relay). 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. All of these functions are performed in a single pass of the services PIC.

In the reverse path, IPv6 packets are sent to the Services DPC where they are encapsulated in IPv4 packets corresponding to the proper softwire and sent to the CE WAN.

The softwire concentrator creates softwires as the IPv4 packets are received from the CE WAN side or IPv6 packets are received from the Internet. A 6rd softwire on the Services DPC is identified by the 3-tuple containing the service set ID, CE softwire initiator IPv4 address, and softwire concentrator IPv4 address. IPv6 flows are also created for the encapsulated IPv6 payload, and are associated with the specific softwire that carried

them in the first place. When the last IPv6 flow associated with a softwire ends, the softwire is deleted. This simplifies configuration and there is no need to create or manage tunnel interfaces.

6rd is supported on Multiservices 100, 400, and 500 PICs on M Series and T Series routers, and on MX Series platforms equipped with Multiservices DPCs.

For more information on 6rd softwires, see RFC 5969, *IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification*.

**Related
Documentation**

- *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*
- [Configuring a 6rd Softwire Concentrator on page 160](#)
- [Configuring a DS-Lite Softwire Concentrator on page 159](#)
- [Configuring Softwire Rules on page 161](#)
- [Configuring Service Sets for Softwire on page 162](#)

PART 2

Configuration

- [NAT Configuration Concepts on page 23](#)
- [NAT Configuration Tasks on page 25](#)
- [Carrier-Grade NAT Complete Configuration Examples on page 75](#)
- [Carrier-Grade NAT Implementation Best Practices on page 111](#)
- [NAT Configuration Statements on page 121](#)
- [Softwire Configuration Tasks on page 159](#)
- [Softwire Configuration Examples on page 165](#)
- [6to4 Configuration on page 183](#)
- [Softwire Configuration Statements on page 187](#)

CHAPTER 4

NAT Configuration Concepts

- [Network Address Translation Configuration Overview on page 23](#)

Network Address Translation Configuration Overview

To configure network address translation (NAT), complete the following high-level steps:

1. Configure the source and destination addresses. For more information, see *Configuring Source and Destination Addresses Network Address Translation Overview*.
2. Define the addresses or prefixes, address ranges, and ports used for NAT. For more information, see *Configuring Pools of Addresses and Ports for Network Address Translation Overview*.
3. If applicable, configure the address pools for network address port translation (NAPT). For more information, see *Configuring Address Pools for Network Address Port Translation (NAPT) Overview*.
4. Configure the NAT rules. Within the rules, include match directions, match conditions, actions, and translation types. For more information, see *Network Address Translation Rules Overview*.
5. Configure service sets for NAT processing. Within each service set, define the interfaces for handling inbound and outbound traffic and a NAT rule or ruleset. For more information, see *Configuring Service Sets for Network Address Translation*.

Related Documentation

- [Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards](#)
- [Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 9](#)

CHAPTER 5

NAT Configuration Tasks

- [Configuring Static Source Translation in IPv4 Networks on page 25](#)
- [Configuring Static Source Translation in IPv6 Networks on page 31](#)
- [Configuring Static Destination Address Translation in IPv4 Networks on page 35](#)
- [Configuring Dynamic Address-Only Source Translation in IPv4 Networks on page 39](#)
- [Configuring Dynamic Source Address and Port Translation in IPv4 Networks on page 43](#)
- [Configuring Dynamic Source Address and Port Translation for IPv6 Networks on page 47](#)
- [Configuring Secured Port Block Allocation on page 49](#)
- [Configuring Deterministic Port Block Allocation on page 51](#)
- [Configuring Stateful NAT64 on page 52](#)
- [Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT on page 54](#)
- [Example: Assigning Addresses from a Dynamic Pool for Static Use on page 60](#)
- [Example: Configuring NAT for Multicast Traffic on page 61](#)
- [Configuring Port Forwarding for Static Destination Address Translation on page 65](#)
- [Configuring Port Forwarding Without Destination Address Translation on page 68](#)
- [Example: Configuring Port Forwarding with Twice NAT on page 69](#)
- [Configuring Port Control Protocol on page 71](#)

Configuring Static Source Translation in IPv4 Networks

To configure the translation type as **basic-nat44**, you must configure the NAT pool and rule, service set with service interface, and trace options. This topic includes the following tasks:

- [Configuring the NAT Pool and Rule on page 26](#)
- [Configuring the Service Set for NAT on page 27](#)
- [Configuring Trace Options on page 28](#)
- [Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range on page 29](#)
- [Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet on page 30](#)

Configuring the NAT Pool and Rule

To configure the NAT pool, rule, and term:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

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

2. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool name address address
```

In the following example, the pool name is **src_pool** and the address is **10.10.10.2/32**.

```
[edit services nat]
user@host# set pool src_pool address 10.10.10.2/32
```

3. Configure the NAT rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the NAT rule name is **rule-basic-nat44** and the match direction is **input**.

```
[edit services nat]
user@host# set rule rule-basic-nat44 match-direction input
```

4. Configure the source address in the **from** statement.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term term-name from from
```

In the following example, the term name is **t1** and the input condition is **source-address 3.1.1.2/32**.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 from source-address 3.1.1.2/32
```

5. Configure the NAT term action and properties of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 then term-action translated-property
```

In the following example, the term action is **translated** and the property of the translated traffic is **source-pool src_pool**.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 then translated source-pool src_pool
```

6. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat44 term t1 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat44**.

```
[edit services nat]
```

```
user@host# set rule rule-basic-nat44 term t1 then translated translation-type
basic-nat44
```

7. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services]
user@host# show
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat44 {
    match-direction input;
    term t1 {
      from {
        source-address {
          3.1.1.2/32;
        }
      }
      then {
        translated {
          source-pool src_pool;
          translation-type {
            basic-nat44;
          }
        }
      }
    }
  }
}
```



NOTE: If you don't configure a stateful firewall (SFW) rule for your traffic, then each packet is subjected to the following default stateful firewall rule:

- Allow any valid packets from inside to outside.
- Create forward and return flow based on packets 5-tuple.
- Allow only valid packets matching return flows from outside to inside.

The stateful firewall's packet validity checks are described in the *Stateful Firewall Anomaly Checking in Junos Network Secure Overview*. When a packets pass stateful firewall validity checking but are not matched by a NAT rule, they are not translated and may be forwarded if the NAT node has a valid route to the packets' destination IP addresses.

Configuring the Service Set for NAT

To configure the service set for NAT:

1. In configuration mode, go to the **[edit services]** hierarchy level.

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

2. Configure the service set.

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

In the following example, the service set name is **s1**.

```
[edit services]
user@host# edit service-set s1
```

3. For the **s1** service set, set the reference to the NAT rules configured at the **[edit services nat]** hierarchy level.

```
[edit services service-set s1]
user@host# set nat-rules rule-name
```

In the following example, the rule name is **rule-basic-nat44**.

```
[edit services service-set s1]
user@host# set nat-rules rule-basic-nat44
```

4. Configure the service interface.

```
[edit services service-set s1]
user@host# set interface-service service-interface service-interface-name
```

In the following example, the service interface name is **ms-1/2/0**.

```
[edit services service-set s1]
user@host# set interface-service service-interface ms-1/2/0
```



NOTE: If you have a Trio-based line card, you can configure an inline-services interface on that card:

```
[edit]
user@host# set interfaces si-0/0/0
[edit services service-set s1]
user@host# set interface-service service-interface si-0/0/0
```

5. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
    nat-rules rule-basic-nat44;
    interface-service {
        service-interface ms-1/2/0;
    }
}
```

Configuring Trace Options

To configure the trace options:

1. In configuration mode, go to the **[edit services adaptive-services-pics]** hierarchy level.

```
[edit]
user@host# edit services adaptive-services-pics
```

2. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

3. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

```
[edit]
user@host# show services
service-set s1 {
  nat-rules rule-basic-nat44;
  interface-service {
    service-interface ms-1/2/0;
  }
}
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat44 {
    match-direction input;
    term t1 {
      from {
        source-address {
          3.1.1.2/32;
        }
      }
      then {
        translated {
          source-pool src_pool;
          translation-type {
            basic-nat44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

Sample Configuration - Static Source NAT Using a Static Pool With An Address Prefix And An Address Range

```
[edit services nat]
```

```
pool p1 {
  address 30.30.30.252/30;
  address-range low 20.20.20.1 high 20.20.20.2;
}
rule r1 {
  match-direction input;
  term {
    from {
      source-address {
        10.10.10.252/30;
      }
    }
    then {
      translated {
        source-pool p1;
        translation-type basic-nat44;
      }
    }
  }
}
```

Sample Configuration - Static Source Nat for One-To-One Mapping Between a Private Subnet and a Public Subnet

```
[edit]
user@host# show services
service-set s1 {
  nat-rules rule-basic-nat44;
  interface-service {
    service-interface ms-1/2/0;
  }
}
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat44 {
    match-direction input;
    term t1 {
      from {
        source-address {
          3.1.1.2/32;
        }
      }
      then {
        translated {
          source-pool src_pool;
          translation-type {
            basic-nat44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```



```
}
}
```

Configuring Static Source Translation in IPv6 Networks

To configure the translation type as **basic-nat66**, you must configure the NAT pool and rule, service set with service interface, and trace options. This topic includes the following tasks:

- [Configuring the NAT Pool and Rule on page 31](#)
- [Configuring the Service Set for NAT on page 32](#)
- [Configuring Trace Options on page 33](#)

Configuring the NAT Pool and Rule

To configure the NAT pool, rule, and term:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

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

2. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool name address address
```

In the following example, the pool name is **src_pool** and the address is **10.10.10.2/32**.

```
[edit services nat]
user@host# set pool src_pool address 10.10.10.2/32
```

3. Configure the NAT rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the rule name is **rule-basic-nat66** and the match direction is **input**.

```
[edit services nat]
user@host# set rule rule-basic-nat66 match-direction input
```

4. Configure the source address in the **from** statement.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term term-name from from
```

In the following, the term name is **t1** and the input condition is **source-address 10:10:10::0/96**.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 from source-address 10:10:10::0/96
```

5. Configure the NAT term action and properties of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 then term-action translated-property
```

In the following example, the term action is **translated** and the property of the translated traffic is **source-pool src_pool**.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 then translated source-pool src_pool
```

6. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat66**.

```
[edit services nat]
user@host# set rule rule-basic-nat66 term t1 then translated translation-type
basic-nat66
```

7. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
nat {
  pool src_pool {
    address 10.10.10.2/32;
  }
  rule rule-basic-nat66 {
    match-direction input;
    term t1 {
      from {
        source-address {
          10:10:10::0/96;
        }
      }
      then {
        translated {
          source-pool src_pool;
          translation-type {
            basic-nat66;
          }
        }
      }
    }
  }
}
```

Configuring the Service Set for NAT

To configure the service set for NAT:

1. In configuration mode, go to the **[edit services]** hierarchy level.

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

2. Configure the service set.

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

In the following example, the service set name is **s1**.

```
[edit services]
user@host# edit service-set s1
```

3. For the **s1** service set, set the reference to the NAT rules configured at the **[edit services nat]** hierarchy level.

```
[edit services service-set s1]
user@host# set nat-rules rule-name
```

In the following example, the rule name is **rule-basic-nat66**.

```
[edit services service-set s1]
user@host# set nat-rules rule-basic-nat66
```

4. Configure the service interface.

```
[edit services service-set s1]
user@host# set interface-service service-interface service-interface-name
```

In the following example, the service interface name is **sp-1/2/0**.

```
[edit services service-set s1]
user@host# set interface-service service-interface sp-1/2/0
```

5. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-basic-nat66;
  interface-service {
    service-interface sp-1/2/0;
  }
}
```

Configuring Trace Options

To configure the trace options at the **[edit services adaptive-services-pics]** hierarchy level:

1. In configuration mode, go to the **[edit services adaptive-services-pics]** hierarchy level.

```
[edit]
user@host# edit services adaptive-services-pics
```

2. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

3. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

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

```
adaptive-services-pics {  
  traceoptions {  
    flag all;  
  }  
}
```

The following example configures the translation type as **basic-nat66**.

```
[edit]  
user@host# show services  
service-set s1 {  
  nat-rules rule-basic-nat66;  
  interface-service {  
    service-interface sp-1/2/0;  
  }  
}  
nat {  
  pool src_pool {  
    address 10.10.10.2/32;  
  }  
  rule rule-basic-nat66 {  
    match-direction input;  
    term t1 {  
      from {  
        source-address {  
          10:10:10::0/96;  
        }  
      }  
      then {  
        translated {  
          source-pool src_pool;  
          translation-type {  
            basic-nat66;  
          }  
        }  
      }  
    }  
  }  
}  
}  
adaptive-services-pics {  
  traceoptions {  
    flag all;  
  }  
}
```

Configuring Static Destination Address Translation in IPv4 Networks

In IPv4 networks, destination address translation is a mechanism used to implement address translation for destination traffic without port mapping. To use destination address translation, the size of the pool address space must be greater than or equal to the destination address space. You must specify a name for the **destination-pool** statement, which can contain multiple addresses, ranges, or prefixes, as long as the number of NAT addresses in the pool is larger than the number of destination addresses in the **from** statement.

To configure destination address translation in IPv4 networks:

1. In configuration mode, go to the **[edit services]** hierarchy level.

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

2. Configure the service set and the NAT rule.

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

In the following example, the name of the service set is **s1** and the name of the NAT rule is **rule-dnat44**.

```
[edit services]
user@host# set service-set s1 nat-rules rule-dnat44
```

3. Go to the **[interface-service]** hierarchy level of the service set.

```
[edit services]
user@host# edit service-set s1 interface-service
```

4. Configure the service interface.

```
[edit services service-set s1 interface-service]
user@host# set service-interface service-interface-name
```

In the following example, the name of the service interface is **ms-0/1/0**.



NOTE: If the service interface is not present in the router, or the specified interface is not functional, the following command can result in an error.

```
[edit services service-set s1 interface-service]
user@host# set service-interface ms-0/1/0
```

5. Go to the **[edit services nat]** hierarchy level. Issue the following command from the top of the services hierarchy, or use the **top** keyword.

```
[edit services service-set s1]
user@host# top edit services nat
```

6. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, **dest-pool** is used as the pool name and **4.1.1.2** as the address.

```
user@host# set pool dest-pool address 4.1.1.2
```

7. Configure the rule, match direction, term, and destination address.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name from
destination-address address
```

In the following example, the name of the rule is **rule-dnat44**, the match direction is **input**, the name of the term is **t1**, and the address is **20.20.20.20**.

```
[edit services nat]
user@host# set rule rule-dnat44 match-direction input term t1 from
destination-address 20.20.20.20
```

8. Go to the **[edit services nat rule rule-dnat44 term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-dnat44 term t1
```

9. Configure the destination pool and the translation type.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool-name translation-type
translation-type
```

In the following example, the destination pool name is **dest-pool**, and the translation type is **dnat-44**.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool translation-type dnat-44
```

10. Go to the **[edit services adaptive-services-pics]** hierarchy level. In the following command, the **top** keyword ensures that the command is run from the top of the hierarchy.

```
[edit services nat rule rule-dnat44 term t1]
user@host# top edit services adaptive-services-pics
```

11. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is configured as **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

12. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-dnat44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
```

```

pool dest-pool {
    address 4.1.1.2/32;
}
rule rule-dnat44 {
    match-direction input;
    term t1 {
        from {
            destination-address {
                20.20.20.20/32;
            }
        }
        then {
            translated {
                destination-pool dest-pool;
                translation-type {
                    dnat-44;
                }
            }
        }
    }
}
}
adaptive-services-pics {
    traceoptions {
        flag all;
    }
}

```

The following example configures the translation type as **dnat-44**.

```

[edit services]
user@host# show
service-set s1 {
    nat-rules rule-dnat44;
    interface-service {
        service-interface ms-0/1/0;
    }
}
nat {
    pool dest-pool {
        address 4.1.1.2/32;
    }
    rule rule-dnat44 {
        match-direction input;
        term t1 {
            from {
                destination-address {
                    20.20.20.20/32;
                }
            }
            then {
                translated {
                    destination-pool dest-pool;
                    translation-type {
                        dnat-44;
                    }
                }
            }
        }
    }
}
}

```

```
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

In the following configuration, **term1** configures source address translation for traffic from any private address to any public address. The translation is applied for all services. **term2** performs destination address translation for Hypertext Transfer Protocol (HTTP) traffic from any public address to the server's virtual IP address. The virtual server IP address is translated to an internal IP address.

```
[edit services nat]
rule my-nat-rule {
  match-direction input;
  term my-term1 {
    from {
      source-address private;
      destination-address public;
    }
    then {
      translated {
        source-pool my-pool; # pick address from a pool
        translation-type napt-44; # dynamic NAT with port translation
      }
    }
  }
  term my-term2 {
    from {
      destination-address 192.168.137.3; # my server's virtual address
      application http;
    }
    then {
      translated {
        destination-pool nat-pool-name;
        translation-type dnat-44; # static destination NAT
      }
    }
  }
}
```

The following configuration performs NAT using the destination prefix **20.20.10.0/32** without defining a pool.

```
[edit services nat]
rule src-nat {
  match-direction input;
  term t1 {
    from {
      destination-address 10.10.10.10/32;
    }
    then {
      translation-type dnat44;
      destination-prefix 20.20.10.0/24;
    }
  }
}
```


}

Configuring Dynamic Address-Only Source Translation in IPv4 Networks

In IPv4 networks, dynamic address translation (dynamic NAT) is a mechanism to dynamically translate the destination traffic without port mapping. To use dynamic NAT, you must specify a source pool name, which includes an address configuration.

To configure dynamic NAT in IPv4 networks:

1. In configuration mode, go to the **[edit services]** hierarchy level.

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

2. Configure the service set and NAT rule.

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

In the following example, the name of the service set is **s1**, and the name of the NAT rule is **rule-dynamic-nat44**.

```
[edit services]
user@host# set service-set s1 nat-rules rule-dynamic-nat44
```

3. Go to the **[interface-service]** hierarchy level for the service set.

```
[edit services]
user@host# edit service-set s1 interface-service
```

4. Configure the service interface.

```
[edit services service-set s1 interface-service]
user@host# set service-interface service-interface-name
```

In the following example, the name of the service interface is **ms-0/1/0**.



NOTE: If the service interface is not present in the router, or the specified interface is not functional, the following command can result in an error.

```
[edit services service-set s1 interface-service]
user@host# set service-interface ms-0/1/0
```

5. Go to the **[edit services nat]** hierarchy level. Issue the following command from the top of the services hierarchy, or use the **top** keyword.

```
[edit services service-set s1 interface-service]
user@host# top edit services nat
```

6. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, the name of the pool is **source-dynamic-pool**, and the address is **10.10.10.0**.

```
[edit services nat]
```

```
user@host# set pool source-dynamic-pool address 10.10.10.0
```

7. Configure the rule, match direction, term, and source address.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name from
source-address address
```

In the following example, the name of the rule is **rule-dynamic-nat44**, the match direction is **input**, the name of the term is **t1**, and the source address is **3.1.1.0**.

```
[edit services nat]
user@host# set rule rule-dynamic-nat44 match-direction input term t1 from
source-address 3.1.1.0
```

8. Go to the **[edit rule rule-dynamic-nat-44 term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-dynamic-nat44 term t1
```

9. Configure the source pool and the translation type.

```
[edit services nat rule rule-dynamic-nat44 term t1]
user@host# set then translated source-pool src-pool-name translation-type
translation-type
```

In the following example, the name of the source pool is **source-dynamic-pool** and the translation type is **dynamic-nat44**.

```
[edit services nat rule rule-dynamic-nat44 term t1]
user@host# set then translated source-pool source-dynamic-pool translation-type
dynamic-nat44
```

10. Go to the **[edit services adaptive-services-pics]** hierarchy level. In the following command, the **top** keyword ensures that the command is run from the top of the hierarchy.

```
[edit services nat rule rule-dynamic-nat44 term t1]
user@host# top edit services adaptive-services-pics
```

11. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is configured as **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

12. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-dynamic-nat44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
  pool source-dynamic-pool {
```

```

        address 10.1.1.0/24;
    }
    rule rule-dynamic-nat44 {
        match-direction input;
        term t1 {
            from {
                source-address {
                    3.1.1.0/24;
                }
            }
            then {
                translated {
                    destination-pool source-dynamic-pool;
                    translation-type {
                        dynamic-nat44;
                    }
                }
            }
        }
    }
}
adaptive-services-pics {
    traceoptions {
        flag all;
    }
}

```

The following example configures the translation type as **dynamic-nat44**.

```

[edit services]
user@host# show
service-set s1 {
    nat-rules rule-dynamic-nat44;
    interface-service {
        service-interface ms-0/1/0;
    }
}
nat {
    pool source-dynamic-pool {
        address 10.1.1.0/24;
    }
    rule rule-dynamic-nat44 {
        match-direction input;
        term t1 {
            from {
                source-address {
                    3.1.1.0/24;
                }
            }
            then {
                translated {
                    destination-pool source-dynamic-pool;
                    translation-type {
                        dynamic-nat44;
                    }
                }
            }
        }
    }
}
adaptive-services-pics {

```

```
    traceoptions {  
        flag all;  
    }  
}
```

The following configuration specifies that NAT is not performed on incoming traffic from the source address **192.168.20.24/32** by providing a NAT rule term **t0** that configures **no-translation**. Dynamic NAT is performed on all other incoming traffic, as configured by term **t1** of the NAT rule.

```
[edit services nat]  
pool my-pool {  
    address-range low 10.10.10.1 high 10.10.10.16;  
    port-automatic;  
}  
rule src-nat {  
    match-direction input;  
    term t0 {  
        from {  
            source-address 192.168.20.24/32;  
        }  
        then {  
            no-translation;  
        }  
    }  
    term t1 {  
        then {  
            translated {  
                translation-type dynamic-nat44;  
                source-pool my-pool;  
            }  
        }  
    }  
}
```

The following configuration performs NAT using the source prefix **20.20.10.0/24** without defining a pool.

```
[edit services nat]  
rule src-nat {  
    match-direction input;  
    term t1 {  
        then {  
            translation-type dynamic-nat44;  
            source-prefix 20.20.10.0/24;  
        }  
    }  
}
```

The following configuration performs NAT using the destination prefix **20.20.10.0/32** without defining a pool.

```
[edit services nat]  
rule src-nat {  
    match-direction input;  
    term t1 {  
        from {
```

```

        destination-address 10.10.10.10/32;
    then {
        translation-type dnat44;
        destination-prefix 20.20.10.0/24;
    }
}
}
}

```

Configuring Dynamic Source Address and Port Translation in IPv4 Networks

Network Address Port Translation (NAPT) is a method by which many network addresses and their TCP/UDP ports are translated into a single network address and its TCP/UDP ports. This translation can be configured in both IPv4 and IPv6 networks. This section describes the steps for configuring NAPT in IPv4 networks.

To configure NAPT, you must configure a rule at the **[edit services nat]** hierarchy level for dynamically translating the source IPv4 addresses.

To configure the NAPT in IPv4 networks:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```

[edit]
user@host# edit services

```

2. Configure the service set and NAT rule.

```

[edit services]
user@host# set service-set service-set-name nat-rules rule-name

```

In the following example, the name of the service set is **s1** and the name of the NAT rule is **rule-napt-44**.

```

[edit services]
user@host# set service-set s1 nat-rules rule-napt-44

```

3. Go to the **[interface-service]** hierarchy level of the service set.

```

[edit services]
user@host# edit service-set s1 interface-service

```

4. Configure the service interface.

```

[edit services service-set s1 interface service]
user@host# set service-interface service-interface-name

```

In the following example, the name of the service interface is **ms-0/1/0**.



NOTE: If the service interface is not present in the router, or the specified interface is not functional, the following command can result in an error.

```

[edit services service-set s1 interface service]
user@host# set service-interface ms-0/1/0

```

5. Go to the **[edit services nat]** hierarchy level. Issue the command from the top of the services hierarchy, or use the **top** keyword.

```
[edit services service-set s1 interface service]
user@host# top edit services nat
```

6. Configure the NAT pool with an address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, the name of the pool is **napt-pool** and the address is **10.10.10.0**.

```
[edit services nat]
user@host# set pool napt-pool address 10.10.10.0
```

7. Configure the port.

```
[edit services nat]
user@host# set pool pool-name port port-type
```

In the following example, the port type is selected as **automatic**.

```
[edit services nat]
user@host# set pool napt-pool port automatic
```

8. Configure the rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the name of the rule is **rule-napt-44** and the match direction is **input**.

```
[edit services nat]
user@host# set rule rule-napt-44 match-direction input
```

9. Configure the term, the action for the translated traffic, and the translation type.

```
[edit services nat]
user@host# set rule rule-name term term-name then translated translated-action
translation-type translation-type
```

In the following example, the name of the term is **t1**, the action for the translated traffic is **translated**, the name of the source pool is **napt-pool**, and the translation type is **napt-44**.

```
[edit services nat]
user@host# set rule rule-napt-44 match-direction input term t1 then translated
source-pool napt-pool translation-type napt-44
```

10. Go to the **[edit services adaptive-services-pics]** hierarchy level. In the command, the **top** keyword ensures that the command is run from the top of the hierarchy.

```
[edit services nat]
user@host# top edit services adaptive-services-pics
```

11. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is configured as **all**.

```
[edit services adaptive-services-pics]
```

```
user@host# set traceoptions flag all
```

12. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-napt-44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
  pool napt-pool {
    address 10.10.10.0/32;
    port {
      automatic;
    }
  }
  rule rule-napt-44 {
    match-direction input;
    term t1 {
      then {
        translated {
          source-pool napt-pool;
          translation-type {
            napt-44;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following example configures the translation type as **napt-44**.

```
[edit services]
user@host# show
service-set s1 {
  nat-rules rule-napt-44;
  interface-service {
    service-interface ms-0/1/0;
  }
}
nat {
  pool napt-pool {
    address 10.10.10.0/32;
    port {
      automatic;
    }
  }
  rule rule-napt-44 {
    match-direction input;
    term t1 {
      then {
        translated {
```

```

        source-pool napt-pool;
        translation-type {
            napt-44;
        }
    }
}
}
}
}
}
adaptive-services-pics {
    traceoptions {
        flag all;
    }
}
}

```

Dynamic Address Translation to a Small Pool with Fallback to NAT

The following configuration shows dynamic address translation from a large prefix to a small pool, translating a /24 subnet to a pool of 10 addresses. When the addresses in the source pool (**src-pool**) are exhausted, NAT is provided by the NAPT overload pool (**pat-pool**).

```
[edit services nat]
pool src-pool {
    address-range low 192.16.2.1 high 192.16.2.10;
}
pool pat-pool {
    address-range low 192.16.2.11 high 192.16.2.12;
    port automatic;
    rule myrule {
        match-direction input;
        term myterm {
            from {
                source-address 10.150.1.0/24;
            }
            then {
                translated {
                    source-pool src-pool;
                    overload-pool pat-pool;
                    translation-type napt-44;
                }
            }
        }
    }
}
```

Dynamic Address Translation with Small Pool

The following configuration shows dynamic address translation from a large prefix to a small pool, translating a /24 subnet to a pool of 10 addresses. Sessions from the first 10 host sessions are assigned an address from the pool on a first-come, first-served basis, and any additional requests are rejected. Each host with an assigned NAT can participate in multiple sessions.

```
[edit services nat]
pool my-pool {
  address-range low 10.10.10.1 high 10.10.10.10;
}
rule src-nat {
  match-direction input;
  term t1 {
```



```

    from {
        source-address 192.168.1.0/24;
    }
    then {
        translated {
            translation-type dynamic-nat44;
            source-pool my-pool;
        }
    }
}
}

```

Configuring Dynamic Source Address and Port Translation for IPv6 Networks

Network Address Port Translation (NAPT) is a method by which many network addresses and their TCP/UDP ports are translated into a single network address and its TCP/UDP ports. This translation can be configured in both IPv4 and IPv6 networks. This section describes the steps for configuring NAPT in IPv6 networks. For information about configuring NAPT in IPv4 networks, see [“Configuring Dynamic Source Address and Port Translation in IPv4 Networks” on page 43](#).

To configure NAPT, you must configure a rule at the **[edit services nat]** hierarchy level for dynamically translating the source IPv6 addresses.

To configure NAPT in IPv6 networks:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```

[edit]
user@host# edit services nat

```

2. Define the pool of IPv6 source addresses that must be used for dynamic translation. For NAPT, also specify port numbers when configuring the source pool.

```

[edit services nat]
user@host# set pool pool name address IPv6 source addresses
user@host# set pool pool name port source ports

```

For example:

```

[edit services nat]
user@host# set pool IPV6-NAPT-Pool address 2002::1/96
user@host# set pool IPV6-NAPT-Pool port automatic

```

3. Define a NAT rule for translating the source addresses. To do this, set the **match-direction** statement of the rule as **input**. In addition, define a term that uses **napt-66** as the translation type for translating the addresses of the pool defined in the previous step.

```

[edit services nat]
user@host# set rule rule name match-direction input
user@host# set rule rule name term term name then translated source-pool pool name
user@host# set rule rule name term term name then translated translation-type napt-66

```

For example:

```

[edit services nat]

```

```

user@host# set rule IPV6-NAPT-Rule match-direction input
user@host# set rule IPV6-NAPT-Rule term t1 then translated source-pool
    IPV6-NAPT-Pool
user@host# set rule IPV6-NAPT-Rule term t1 then translated translation-type napt-66

```

4. Enter the **up** command to navigate to the **[edit services]** hierarchy level.

```

[edit services nat]
user@host# up

```

5. Define a service set to specify the services interface that must be used, and reference the NAT rule implemented for NAPT translation.

```

[edit services]
user@host# set service-set service-set name interface-service service interface
    services interface
user@host# set service-set service-set name nat-rules rule name

```

For example:

```

[edit services]
user@host# set service-set IPV6-NAPT-ServiceSet interface-service service-interface
    ms-0/1/0
user@host# set service-set IPV6-NAPT-ServiceSet nat-rules IPV6-NAPT-Rule

```

6. Define the trace options for the adaptive services PIC.

```

[edit services]
user@host# set adaptive-services-pics traceoptions flag tracing parameter

```

For example:

```

[edit services]
user@host# set adaptive-services-pics traceoptions flag all

```

The following example configures dynamic source (address and port) translation or NAPT for an IPv6 network.

```

[edit services]
user@host# show
    service-set IPV6-NAPT-ServiceSet {
        nat-rules IPV6-NAPT-Rule;
        interface-service {
            service-interface ms-0/1/0;
        }
    }
    nat {
        pool IPV6-NAPT-Pool {
            address 2002::1/96;
            port automatic;
        }
        rule IPV6-NAPT-Rule {
            match-direction input;
            term term1 {
                then {
                    translated {
                        source-pool IPV6-NAPT-Pool;
                        translation-type {
                            napt-66;
                        }
                    }
                }
            }
        }
    }

```

```

    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
}

```

Configuring Secured Port Block Allocation

To configure secured port block allocation:

1. At the `[edit services nat pool poolname]` hierarchy level, create a pool.

```
user@host# edit services nat pool poolname
```

For example:

```
user@host# edit services nat pool pba-pool1
```

2. Define the range of addresses to be translated, specifying the upper and lower limits of the range or an address prefix that describes the range.

```
[edit services nat pool pba-pool1]
user@host# set address-range low address high address
```

Or

```
user@host# set address address-prefix
```

For example:

```
[edit services nat pool pba-pool1]
user@host# set address 203.0.113.0/24
```

3. Define the range of ports to be used in the translation, or use automatic port assignment by the Junos OS. You can optionally specify random assignment of ports (sequential assignment is the default).

```
[edit services nat pool pba-pool1]
user@host# set port range low address high address random
```

Or

```
user@host# set port automatic random-allocation
```

For example:

```
[edit services nat pool pba-pool1]
user@host# set port range low 256 high 511 random
```

Or

```
[edit services nat pool pba-pool1]
user@host# set port automatic random-allocation
```



NOTE: When you configure a port range, the range should be a multiple of the port block-size value (see Step 4). When the NAT pool port range is *not* a multiple of the port block-size value, the number of ports or port-blocks that are effectively available for use is less than the configured number of ports and port-blocks. The port block allocation mechanism uses ports in the range 0 through 1023 of a NAT address.

When you configure automatic assignment of ports, the available port range for allocation is 1024 through 65535. Automatic allocation can result in no ports being available for use. Use the [show services nat pool](#) command on the Routing Engine after you configure the port block allocation method to determine the number of ports and port blocks available for allocation to users.

4. Configure secured port block allocation. Specify **active-block-timeout**, **block-size**, and **max-blocks-per-address**, or accept the default values for those options.

```
[edit services nat pool pba-pool1]
user@host# set secured-port-block-allocation active-block-timeout
active-block-timeout block-size block-size max-blocks-per-address
max-blocks-per-address
```

For example:

```
[edit services nat pool pba-pool1]
user@host# set secured-port-block-allocation active-block-timeout 120 block-size
256 max-blocks-per-address 12
```



NOTE: In order for secured-port-block-allocation configuration changes to take effect, you must reboot the services PIC whenever you change any of the following nat pool options:

- *pool-name*
- address or address-range
- port range
- port secured-port-block-allocation block-size
- port secured-port-block-allocation max-blocks-per-address.
- port secured-port-block-allocation active-block-timeout.
- from hierarchy in the nat rule

Related
Documentation

- [Network Address Translation Configuration Overview on page 23](#)

Configuring Deterministic Port Block Allocation

To configure deterministic port block allocation:

1. At to the `[edit services nat pool poolname]` hierarchy level, create a pool.

```
user@host# edit services nat pool poolname
```

For example:

```
user@host# edit services nat pool pba-pool2
```

2. Define the range of addresses to be translated, specifying the upper and lower limits of the range or an address prefix that describes the range.

```
[edit services nat pool pba-pool1]
user@host# set address-range low address high address
```

Or

```
user@host# set address address-prefix
```

For example:

```
[edit services nat pool pba-pool1]
user@host# set address-range low 32.32.32.1 high 32.32.32.253
```

3. Specify automatic port assignment by the Junos OS.

```
[edit services nat pool pba-pool1]
user@host# set port automatic
```

4. Configure deterministic port block allocation. Specify **block-size** or accept the default value of 512.

. You can also specify **include-boundary-addresses** if you want the lowest and highest addresses in the source address range of a NAT rule to be translated when the NAT pool is used.

```
[edit services nat pool pba-pool1]
user@host# set port deterministic-port-block-allocation block-size block-size
include-boundary-addresses
```

For example:

```
[edit services nat pool pba-pool1]
user@host# set port deterministic-port-block-allocation block-size 256
```



NOTE: In order for deterministic-port-block-allocation configuration changes to take effect, you must reboot the services PIC whenever you change any of the following nat pool options:

- address or address-range
- port range
- port deterministic-port-block-allocation block-size

Related Documentation • [Network Address Translation Configuration Overview on page 23](#)

Configuring Stateful NAT64

Stateful NAT64 is a mechanism used to move to an IPv6 network and at the same time deal with IPv4 address depletion. By allowing IPv6-only clients to contact IPv4 servers using unicast UDP, TCP, or ICMP, several IPv6-only clients can share the same public IPv4 server address. To allow sharing of the IPv4 server address, stateful NAT64 translates incoming IPv6 packets into IPv4, and vice versa.

To configure stateful NAT64, you must configure a rule at the **[edit services nat]** hierarchy level for translating the source address dynamically and the destination address statically.



BEST PRACTICE: When you configure the service set that includes your NAT rule, include the **set stateful-nat64 clear-dont-fragment-bit** at the **[edit services service-set service-set-name]** hierarchy level. This clears the DF (don't fragment) bit in order to prevent unnecessary creation of an IPv6 fragmentation header when translating IPv4 packets that are less than 1280 bytes. RFC 6145, *IP/ICMP Translation Algorithm*, provides a full discussion of the use of the DF flag to control generation of fragmentation headers. For more information on service sets for NAT, see *Configuring Service Sets for Network Address Translation*.

To configure stateful NAT64:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

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

2. Define the pool of source addresses to be used for dynamic translation.

```
[edit services nat]
user@host# set pool pool name address source addresses
user@host# set pool pool name port source ports
```

For example:

```
[edit services nat]
user@host# set pool src-pool-nat64 address 203.0.113.0/24
user@host# set pool src-pool-nat64 port automatic
```

3. Define a NAT rule for translating the source addresses. Set the **match-direction** statement of the rule as **input**. Then define a term that uses **stateful-nat64** as the translation type for translating the addresses of the pool defined in the previous step.

```
[edit services nat]
user@host# set rule rule name match-direction input
user@host# set rule rule name term term name from source-address source address
user@host# set rule rule name term term name from destination-address destination address
user@host# set rule rule name term term name then translated source-pool pool name
```

```

user@host# set rule rule name term term name then translated destination-prefix
destination prefix
user@host# set rule rule name term term name then translated translation-type
stateful-nat64

```

For example:

```

[edit services nat]
user@host# set rule stateful-nat64 match-direction input
user@host# set rule stateful-nat64 term t1 from source-address 2001:DB8::0/96
user@host# set rule stateful-nat64 term t1 from destination-address 64:FF9B::/96
user@host# set rule stateful-nat64 term t1 then translated source-pool src-pool-nat64
user@host# set rule stateful-nat64 term t1 then translated destination-prefix
64:FF9B::/96
user@host# set rule stateful-nat64 term t1 then translated translation-type
stateful-nat64

```

The following example configures dynamic source address (IPv6-to-IPv4) and static destination address (IPv6-to-IPv4) translation.

```

[edit services]
user@host# show
nat {
    pool src-pool-nat64 {
        address 203.0.113.0/24;
        port {
            automatic;
        }
    }
    rule stateful-nat64 {
        match-direction input;
        term t1 {
            from {
                source-address {
                    2001:db8::0/96;
                }
                destination-address {
                    64:ff9b::/96;
                }
            }
            then {
                translated {
                    source-pool src-pool-nat64;
                    destination-prefix 64:ff9b::/96;
                    translation-type {
                        stateful-nat64;
                    }
                }
            }
        }
    }
}
service-set sset-nat64 {
    nat-options {
        stateful-nat64 {
            clear-dont-fragment-bit;
        }
    }
    service-set-options;
    nat-rules stateful-nat64;
}

```

```
interface-service {  
    service-interface ms-0/1/0;  
}  
}
```

Configuring Protocol Translation Between IPv6 and IPv4 Networks - NAT-PT

To configure the translation type as **basic-nat-pt**, you must configure the DNS ALG application, the NAT pools and rules, a service set with a service interface, and trace options. This topic includes the following tasks:

- [Configuring the DNS ALG Application on page 54](#)
- [Configuring the NAT Pool and NAT Rule on page 54](#)
- [Configuring the Service Set for NAT on page 58](#)
- [Configuring Trace Options on page 58](#)

Configuring the DNS ALG Application

To configure the DNS ALG application:

1. In configuration mode, go to the **[edit applications]** hierarchy level.

```
[edit]  
user@host# edit applications
```

2. Configure the ALG to which the DNS traffic is destined at the **[edit applications]** hierarchy level. Define the application name and specify the application protocol to use in match conditions in the first NAT rule or term.

```
[edit applications]  
user@host# set application application-name application-protocol application-protocol
```

In the following example, the application name is **dns-alg** and application protocol is **dns**.

```
[edit applications]  
user@host# set application dns-alg application-protocol dns
```

3. Verify the configuration by using the **show** command at the **[edit applications]** hierarchy level.

```
[edit applications]  
user@host# show  
application dns-alg {  
    application-protocol dns;  
}
```

Configuring the NAT Pool and NAT Rule

To configure the NAT pool and NAT rule:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

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


2. Configure the NAT pool and its address.

```
[edit services nat]
user@host# set pool pool-name address address
```

In the following example, the name of the NAT pool is **p1** and the address is **10.10.10.2/32**.

```
[edit services nat]
user@host# set pool p1 address 10.10.10.2/32
```

3. Configure the source pool and its address.

```
[edit services nat]
user@host# set pool source-pool-name address address
```

In the following example, the name of the source pool is **src_pool0** and the source pool address is **20.1.1.1/32**.

```
[edit services nat]
user@host# set pool src_pool0 address 20.1.1.1/32
```

4. Configure the destination pool and its address.

```
[edit services nat]
user@host# set pool destination-pool-name address address
```

In the following example, the name of the destination pool is **dst_pool0** and the destination pool address is **50.1.1.2/32**.

```
[edit services nat]
user@host# set pool dst_pool0 address 50.1.1.2/32
```

5. Configure the rule and the match direction.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction
```

In the following example, the rule name is **rule-basic-nat-pt** and the match direction is **input**.

```
[edit services nat]
user@host# set rule basic-nat-pt match-direction input
```

6. Configure the term and the input conditions for the NAT term.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term term from from
```

In the following example, the term is **t1** and the input conditions are **source-address 2000::2/128**, **destination-address 4000::2/128**, and **applications dns_alg**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 from source-address 2000::2/128
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 from destination-address 4000::2/128
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 from applications dns_alg
```

7. Configure the NAT term action and the properties of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then term-action translated-property
```

In the following example, the term action is **translated** and the properties of the translated traffic are **source-pool** `src_pool0`, **destination-pool** `dst_pool0`, and **dns-alg-prefix** `10:10:10::0/96`.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated source-pool src_pool0
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated destination-pool
dst_pool0
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated dns-alg-prefix
10:10:10::0/96
```

8. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat-pt**.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t1 then translated translation-type
basic-nat-pt
```

9. Configure another term and the input conditions for the NAT term.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term term-name from from
```

In the following example, the term name is `t2` and the input conditions are **source-address** `2000::2/128` and **destination-address** `10:10:10::0/96`.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 from source-address 2000::2/128
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 from destination-address 10:10:10::0/96
```

10. Configure the NAT term action and the property of the translated traffic.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 then term-action translated-property
```

In the following example, the term action is **translated** and the property of the translated traffic is **source-prefix** `19.19.19.1/32`.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 then translated source-prefix
19.19.19.1/32
```

11. Configure the translation type.

```
[edit services nat]
user@host# set rule rule-basic-nat-pt term t2 then translated translation-type
translation-type
```

In the following example, the translation type is **basic-nat-pt**.

```
[edit services nat]
```

```
user@host# set rule rule-basic-nat-pt term t2 then translated translation-type
basic-nat-pt
```

12. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services nat]
user@host# show
pool p1 {
    address 10.10.10.2/32;
}
pool src_pool0 {
    address 20.1.1.1/32;
}
pool dst_pool0 {
    address 50.1.1.2/32;
}
rule rule-basic-nat-pt {
    match-direction input;
    term t1 {
        from {
            source-address {
                2000::2/128;
            }
            destination-address {
                4000::2/128;
            }
            applications dns_alg;
        }
        then {
            translated {
                source-pool src_pool0;
                destination-pool dst_pool0;
                dns_alg-prefix 10:10:10::0/96;
                translation-type {
                    basic-nat-pt;
                }
            }
        }
    }
}
term t2 {
    from {
        source-address {
            2000::2/128;
        }
        destination-address {
            10:10:10::0/96;
        }
    }
    then {
        translated {
            source-prefix 19.19.19.1/32;
            translation-type {
                basic-nat-pt;
            }
        }
    }
}
}
```

Configuring the Service Set for NAT

To configure the service set for NAT:

1. In configuration mode, go to the **[edit services]** hierarchy level.

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

2. Configure the service set.

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

In the following example, the name of the service set is **ss_dns**.

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

3. Configure the service set with NAT rules.

```
[edit services service-set ss_dns]
user@host# set nat-rules rule-name
```

In the following example, the rule name is **rule-basic-nat-pt**.

```
[edit services service-set ss_dns]
user@host# set nat-rules rule-basic-nat-pt
```

4. Configure the service interface.

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

In the following example, the name of service interface is **sp-1/2/0**.

```
[edit services service-set ss_dns]
user@host# set interface-service service-interface sp-1/2/0
```

5. Verify the configuration by using the **show services** command from the **[edit]** hierarchy level.

```
[edit]
user@host# show services
  service-set ss_dns {
    nat-rules rule-basic-nat-pt;
    interface-service {
      service-interface sp-1/2/0;
    }
  }
```

Configuring Trace Options

To configure the trace options:

1. In configuration mode, go to the **[edit services adaptive-services-pics]** hierarchy level.

```
[edit]
user@host# edit services adaptive-services-pics
```

2. Configure the trace options.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag tracing parameter
```

In the following example, the tracing parameter is **all**.

```
[edit services adaptive-services-pics]
user@host# set traceoptions flag all
```

3. Verify the configuration by using the **show** command at the **[edit services]** hierarchy level.

```
[edit services]
user@host# show
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

The following example configures the translation type as **basic-nat-pt**.

```
[edit]
user@host# show services
service-set ss_dns {
  nat-rules rule-basic-nat-pt;
  interface-service {
    service-interface sp-1/2/0;
  }
}
nat {
  pool p1 {
    address 10.10.10.2/32;
  }
  pool src_pool0 {
    address 20.1.1.1/32;
  }
  pool dst_pool0 {
    address 50.1.1.2/32;
  }
  rule rule-basic-nat-pt {
    match-direction input;
    term t1 {
      from {
        source-address {
          2000::2/128;
        }
        destination-address {
          4000::2/128;
        }
        applications dns_alg;
      }
      then {
        translated {
          source-pool src_pool0;
          destination-pool dst_pool0;
          dns_alg-prefix 10:10:10::0/96;
          translation-type {
            basic-nat-pt;
          }
        }
      }
    }
  }
}
```

```
    }
    term t2 {
      from {
        source-address {
          2000::2/128;
        }
        destination-address {
          10:10:10::0/96;
        }
      }
      then {
        translated {
          source-prefix 19.19.19.1/32;
          translation-type {
            basic-nat-pt;
          }
        }
      }
    }
  }
}
adaptive-services-pics {
  traceoptions {
    flag all;
  }
}
```

Example: Assigning Addresses from a Dynamic Pool for Static Use

The following configuration statically assigns a subset of addresses that are configured as part of a dynamic pool (**dynamic-pool**) to two separate static pools (**static-pool** and **static-pool2**).

```
[edit services nat]
pool dynamic-pool {
  address 20.20.10.0/24;
}
pool static-pool {
  address-range low 20.20.10.10 high 10.20.10.12;
}
pool static-pool2 {
  address 20.20.10.15/32;
}
rule src-nat {
  match-direction input;
  term t1 {
    from {
      source-address 30.30.30.0/24;
    }
    then {
      translation-type dynamic-nat44;
      source-pool dynamic-pool;
    }
  }
}
term t2 {
  from {
    source-address 10.10.10.2;
```

```


    }
    then {
        translation-type basic-nat44;
        source-pool static-pool;
    }
}
term t3 {
    from {
        source-address 10.10.10.10;
    }
    then {
        translation-type basic-nat44;
        source-pool static-pool2;
    }
}
}
}

```

Example: Configuring NAT for Multicast Traffic

Figure 7 on page 61 illustrates the network setup for the following configuration, which allows IP multicast traffic to be sent to the Multiservices PIC.

Figure 7: Configuring NAT for Multicast Traffic



- [Rendezvous Point Configuration on page 61](#)
- [Router 1 Configuration on page 64](#)

Rendezvous Point Configuration

On the rendezvous point (RP), all incoming traffic from the multicast source at **192.168.254.0/27** is sent to the static NAT pool **mcast_pool**, where its source is translated to **20.20.20.0/27**. The service set **nat_ss** is a next-hop service set that allows IP multicast traffic to be sent to the Multiservices DPC or Multiservices PIC. The inside interface on the PIC is **ms-1/1/0.1** and the outside interface is **ms-1/1/0.2**.

```

[edit services]
nat {
    pool mcast_pool {
        address 20.20.20.0/27;
    }
    rule nat_rule_1 {
        match-direction input;
        term 1 {
            from {
                source-address 192.168.254.0/27;
            }
        }
        then {

```

```
        translated {
            source-pool mcast_pool;
            translation-type basic-nat44;
        }
        syslog;
    }
}
service-set nat_ss {
    allow-multicast;
    nat-rules nat_rule_1;
    next-hop-service {
        inside-service-interface ms-1/1/0.1;
        outside-service-interface ms-1/1/0.2;
    }
}
```

The Gigabit Ethernet interface **ge-0/3/0** carries traffic out of the RP to Router 1. The multiservices interface **ms-1/1/0** has two logical interfaces: **unit 1** is the inside interface for next-hop services and **unit 2** is the outside interface for next-hop services. Multicast source traffic comes in on the Fast Ethernet interface **fe-1/2/1**, which has the firewall filter **fbf** applied to incoming traffic.

```
[edit interfaces]
ge-0/3/0 {
    unit 0 {
        family inet {
            address 10.10.1.1/30;
        }
    }
}
ms-1/1/0 {
    unit 0 {
        family inet;
    }
    unit 1 {
        family inet;
        service-domain inside;
    }
    unit 2 {
        family inet;
        service-domain outside;
    }
}
fe-1/2/1 {
    unit 0 {
        family inet {
            filter {
                input fbf;
            }
            address 192.168.254.27/27;
        }
    }
}
```


Multicast packets can only be directed to the Multiservices DPC or the Multiservices PIC using a next-hop service set. In the case of NAT, you must also configure a VPN routing and forwarding instance (VRF). Therefore, the routing instance **stage** is created as a “dummy” forwarding instance. To direct incoming packets to **stage**, you configure filter-based forwarding through a firewall filter called **fbf**, which is applied to the incoming interface **fe-1/2/1**. A lookup is performed in **stage.inet.0**, which has a multicast static route that is installed with the next hop pointing to the PIC’s inside interface. All multicast traffic matching this route is sent to the PIC.

```
[edit firewall]
filter fbf {
  term 1 {
    then {
      routing-instance stage;
    }
  }
}
```

The routing instance **stage** forwards IP multicast traffic to the inside interface **ms-1/1/0.1** on the Multiservices DPC or Multiservices PIC:

```
[edit]
routing-instances stage {
  instance-type forwarding;
  routing-options {
    static {
      route 224.0.0.0/4 next-hop ms-1/1/0.1;
    }
  }
}
```

You enable OSPF and Protocol Independent Multicast (PIM) on the Fast Ethernet and Gigabit Ethernet logical interfaces over which IP multicast traffic enters and leaves the RP. You also enable PIM on the outside interface (**ms-1/1/0.2**) of the next-hop service set.

```
[edit protocols]
ospf {
  area 0.0.0.0 {
    interface fe-1/2/1.0 {
      passive;
    }
    interface lo0.0;
    interface ge-0/3/0.0;
  }
}
pim {
  rp {
    local {
      address 10.255.14.160;
    }
  }
  interface fe-1/2/1.0;
  interface lo0.0;
  interface ge-0/3/0.0;
```

```
interface ms-1/1/0.2;  
}
```

As with any filter-based forwarding configuration, in order for the static route in the forwarding instance **stage** to have a reachable next hop, you must configure routing table groups so that all interface routes are copied from **inet.0** to the routing table in the forwarding instance. You configure routing tables **inet.0** and **stage.inet.0** as members of **fbf_rib_group**, so that all interface routes are imported into both tables.

```
[edit routing-options]  
interface-routes {  
  rib-group inet fbf_rib_group;  
}  
rib-groups fbf_rib_group {  
  import-rib [ inet.0 stage.inet.0 ];  
}  
multicast {  
  rpf-check-policy no_rpf;  
}
```

Reverse path forwarding (RPF) checking must be disabled for the multicast group on which source NAT is applied. You can disable RPF checking for specific multicast groups by configuring a policy similar to the one in the example that follows. In this case, the **no_rpf** policy disables RPF check for multicast groups belonging to **224.0.0.0/4**.

```
[edit policy-options]  
policy-statement no_rpf {  
  term 1 {  
    from {  
      route-filter 224.0.0.0/4 orlonger;  
    }  
    then reject;  
  }  
}
```

Router 1 Configuration

The Internet Group Management Protocol (IGMP), OSPF, and PIM configuration on Router 1 is as follows. Because of IGMP static group configuration, traffic is forwarded out **fe-3/0/0.0** to the multicast receiver without receiving membership reports from host members.

```
[edit protocols]  
igmp {  
  interface fe-3/0/0.0 {  
  }  
}  
ospf {  
  area 0.0.0.0 {  
    interface fe-3/0/0.0 {  
      passive;  
    }  
    interface lo0.0;  
    interface ge-7/2/0.0;  
  }  
}
```

```

pim {
  rp {
    static {
      address 10.255.14.160;
    }
  }
  interface fe-3/0/0.0;
  interface lo0.0;
  interface ge-7/2/0.0;
}
}

```

The routing option creates a static route to the NAT pool, **mcast_pool**, on the RP.

```

[edit routing-options]
static {
  route 20.20.20.0/27 next-hop 10.10.1.1;
}

```

Configuring Port Forwarding for Static Destination Address Translation

Starting with Junos OS Release 11.4, you can map an external IP address and port with an IP address and port in a private network. This allows the destination address and port of a packet to be changed to reach the correct host in a Network Address Translation (NAT) gateway. The translation facilitates reaching a host within a masqueraded, typically private, network, based on the port number on which the packet was received from the originating host. Port forwarding allows remote computers, such as public machines on the Internet, to connect to a non-standard port (port other than 80) of a specific computer within a private network. An example of this type of destination is the host of a public HTTP server within a private network. Port forwarding is supported only with **dnat-44** and **twice-napt-44** on IPv4 networks. Port forwarding works only with the FTP application-level gateway (ALG). Port forwarding also supports endpoint-independent mapping (EIM), endpoint-independent filtering (EIF), and address pooling paired (APP). Port forwarding has no support for technologies such as IPv6 rapid deployment (6rd) and dual-stack lite (DS-Lite) that offer IPv6 services over IPv4 infrastructure.

To configure destination address translation with port forwarding in IPv4 networks:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```

[edit]
user@host# edit services nat

```

2. Configure the NAT pool with an address.

```

[edit services nat]
user@host# set pool pool-name address address

```

In the following example, **dest-pool** is used as the pool name and **4.1.1.2** as the address.

```

user@host# set pool dest-pool address 4.1.1.2

```

3. Configure the rule, match direction, term, and destination address.

```

[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name from
destination-address address

```

In the following example, the name of the rule is **rule-dnat44**, the match direction is **input**, the name of the term is **t1**, and the address is **20.20.20.20**.

```
[edit services nat]
user@host# set rule rule-dnat44 match-direction input term t1 from
destination-address 20.20.20.20
```

4. Configure the destination port range.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name from
destination-port range range high | low
```

In the following example, the upper port range is **50** and the lower port range is **20**.

```
[edit services nat]
user@host# set rule rule-dnat44 match-direction input term t1 from destination-port
range range high 50 low 20
```

5. Go to the **[edit services nat rule rule-dnat44 term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-dnat44 term t1
```

6. Configure the destination pool.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool-name
```

In the following example, the destination pool name is **dest-pool**, and the translation type is **dnat-44**.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then translated destination-pool dest-pool
```

7. Configure the mapping for port forwarding and the translation type.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then port-forwarding-mappings map-name translation-type
translation-type
```

In the following example, the port forwarding map name is **map1**, and the translation type is **dnat-44**.

```
[edit services nat rule rule-dnat44 term t1]
user@host# set then port-forwarding-mappings map1 translation-type dnat-44
```

8. Go to the **[edit services nat port-forwarding map1]** hierarchy level.

```
[edit services nat]
user@host# edit port-forwarding map1
```

9. Configure the mapping for port forwarding.

```
[edit port-forwarding map1]
user@host# set destined-port port-id
user@host# set translated-port port-id
```

In the following example, the destination port is **45** and the translated port is **23**.

```
[edit port-forwarding map1]
user@host# set destined-port 23
user@host# set translated-port 45
```

**NOTE:**

- Multiple port mappings are supported with port forwarding. Up to 32 port maps can be configured for port forwarding.
- The destination port should not overlap the port range configured for NAT.

10. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services]
user@host# show
nat {
  pool dest-pool {
    address 4.1.1.2/32;
  }
  rule rule-dnat44 {
    match-direction input;
    term t1 {
      from {
        destination-address {
          20.20.20.20/32;
        }
        destination-port {
          range low 20 high 50;
        }
      }
      then {
        port-forwarding-mappings map1;
        translated {
          destination-pool dest-pool;
          translation-type {
            dnat-44;
          }
        }
      }
    }
  }
}
port-forwarding map1 {
  destined-port 45;
  translated-port 23;
}
```

**NOTE:**

- A similar configuration is possible with twice NAT for IPv4. See [“Example: Configuring Port Forwarding with Twice NAT” on page 69](#).
- Port forwarding and stateful firewall can be configured together. Stateful firewall has precedence over port forwarding.

**Related
Documentation**

- [Configuring Static Destination Address Translation in IPv4 Networks on page 35](#)

Configuring Port Forwarding Without Destination Address Translation

Starting with Junos OS Release 12.1, you can configure port forwarding without translating a destination address.

To configure port forwarding without destination address translation in IPv4 networks:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

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

2. Configure the rule, match direction, term, and destination address.

```
[edit services nat]
user@host# set rule rule-name match-direction match-direction term term-name
```

In the following example, the name of the rule is **rule-port-forwarding**, the match direction is **input**, and the name of the term is **t1**.

```
[edit services nat]
user@host# set rule rule-port-forwarding match-direction input term t1
```

3. Go to the **[edit services nat rule rule-port-forwarding term t1]** hierarchy level.

```
[edit services nat]
user@host# edit rule rule-port-forwarding term t1
```

4. Specify that there is no address translation for this rule.

```
[edit services nat rule rule-port-forwarding term t1]
user@host# set then no-translation
```

5. Configure the mapping for port forwarding and the translation type.

```
[edit services nat rule rule-port-forwarding term t1]
user@host# set then port-forwarding-mappings map-name
```

In the following example, the port forwarding map name is **map1**.

```
[edit services nat rule rule-port-forwarding term t1]
user@host# set then port-forwarding-mappings map1
```

6. Go to the **[edit services nat port-forwarding map1]** hierarchy level.

```
[edit services nat]
user@host# edit port-forwarding map1
```

7. Configure the mapping for port forwarding.

```
[edit port-forwarding map1]
user@host# set destined-port port-id
user@host# set translated-port port-id
```

In the following example, the destination port is **45** and the translated port is **23**.

```
[edit port-forwarding map1]
user@host# set destined-port 23
user@host# set translated-port 45
```

**NOTE:**

- Multiple port mappings are supported with port forwarding. Up to 32 port maps can be configured for port forwarding.
- The destination port should not overlap the port range configured for NAT.

8. Verify the configuration by using the **show** command at the **[edit services nat]** hierarchy level.

```
[edit services]
user@host# show
nat {
  rule rule-port-forwarding {
    match-direction input;
    term t1 {
      then {
        port-forwarding-mappings map1;
        no-translation      }
      }
    }
  }
  port-forwarding map1 {
    destined-port 45;
    translated-port 23;
  }
}
```



NOTE: Port forwarding and stateful firewall can be configured together. Stateful firewall has precedence over port forwarding.

Example: Configuring Port Forwarding with Twice NAT

The following example configures port forwarding with **twice-napt-44** as the translation type. The example also has stateful firewall and multiple port maps configured.

```
[edit services]
user@host# show
service-set in {
  syslog {
    host local {
      services any;
    }
  }
  stateful-firewall-rules r;
  nat-rules r;
  interface-service {
    service-interface sp-10/0/0.0;
  }
}
stateful-firewall {
  rule r {
    match-direction input;
    term t {
```

```
        from {
            destination-port {
                range low 20 high 5000;
            }
        }
        then {
            reject;
        }
    }
}
nat {
    pool x {
        address 12.0.0.2/32;
    }
    rule r {
        match-direction input;
        term t {
            from {
                destination-address {
                    14.0.0.2/32;
                }
                destination-port {
                    range low 10 high 20000;
                }
            }
            then {
                port-forwarding-mappings y;
                translated {
                    destination-pool x;
                    translation-type {
                        twice-napt-44;
                    }
                }
            }
        }
    }
}
port-forwarding y {
    destined-port 45;
    translated-port 23;
    destined-port 55;
    translated-port 33;
    destined-port 65;
    translated-port 43;
}
}
adaptive-services-pics {
    traceoptions {
        file sp-trace;
        flag all;
    }
}
```


**NOTE:**

- Stateful firewall has precedence over port forwarding. In this example, for instance, no traffic destined to any port between 20 and 5000 will be translated.
- Up to 32 port maps can be configured.

Related Documentation

- [Configuring Port Forwarding for Static Destination Address Translation on page 65](#)

Configuring Port Control Protocol

This topic describes the following configuration tasks:

- [Configuring PCP Server Options on page 71](#)
- [Configuring a PCP Rule on page 72](#)
- [Configuring a Service Set to Apply PCP on page 73](#)
- [SYSLOG Message Configuration on page 73](#)

Configuring PCP Server Options

1. Go to the `[edit services pcp pcp-server server-name]` hierarchy level and specify a PCP server name.

```
user @host# edit services pcp 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).

```
[edit services pcp pcp-server s1]
user @host# set ipv6-address ipv6-address
```

or

```
[edit services pcp pcp-server s1]
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 pcp-server s1]
user @host# set software-concentrator software-concentrator-name
```

4. Specify the minimum and maximum mapping lifetimes for the server.

```
[edit services pcp pcp-server s1]
user @host# set mapping-lifetime-minimum mapping-lifetime-minimum
user @host# set mapping-lifetime-maximum mapping-lifetime-maximum
```

5. Specify the time limits for generating short lifetime or long lifetime errors.

```
[edit services pcp pcp-server s1]
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 pcp-server s1]
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 pcp-server s1]
user @host# set nat-options pcp-nat-pool 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; 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.

8. (Optional)—Configure the maximum number of mappings per client. The default is 32 and maximum is 128.

```
[edit services pcp pcp-server s1]
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 **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 pcp server that handles selected traffic

1. Go to the **[edit services pcp rule *rulename*]** hierarchy level and specify **match-direction** input.

```
user @host# edit services pcp rule rulename
user @host# set match-direction input
```

2. Go to the **[edit services pcp rule *rulename* term *termname*]** hierarchy level and provide a termname.

```
user @host# edit term termname
```

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.

4. Set the **then** option to identify the target pcp server.

```
[edit services pcp rule rulename term termname]
user @host# set then pcp-server server-name
```

Configuring a Service Set to Apply PCP

To use PCP, you must provide the rule-name (or name of a list of rulenames) in the **pcp-rule *rulename*** 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 **software-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.

CHAPTER 6

Carrier-Grade NAT Complete Configuration Examples

- [Example: Configuring Basic NAT44 on page 75](#)
- [Example: NAPT Configuration for the MS-MPC on page 77](#)
- [Example: Configuring NAT-PT on page 82](#)
- [Example: Configuring Inline Network Address Translation - Interface-Service Service Set on page 96](#)
- [Port Control Protocol Configuration Examples on page 104](#)

Example: Configuring Basic NAT44

This example describes how to implement a basic NAT44 configuration.

- [Requirements on page 75](#)
- [Overview on page 75](#)
- [Configuring Basic NAT44 on page 76](#)

Requirements

This example uses the following hardware and software components:

- An MX Series 3D Universal Edge router with a Services DPC or an M Series Multiservice Edge router with a services PIC
- A domain name server (DNS)
- Junos OS Release 11.4 or higher

Overview

This example shows a complete CGN NAT44 configuration and advanced options.

Configuring Basic NAT44

Chassis Configuration

Step-by-Step Procedure

To configure the service PIC (FPC 5 Slot 0) with the Layer 3 service package:

1. Go to the **[edit chassis]** hierarchy level.

```
user@host# edit chassis
```
2. Configure the Layer 3 service package.

```
[edit chassis]  
user@host# set fpc 5 pic 0 adaptive-services service-package layer-3
```

Interfaces Configuration

Step-by-Step Procedure

To configure interfaces to the private network and the public Internet:

1. Define the interface to the private network.

```
user@host# edit interfaces ge-1/3/5  
[edit interfaces ge-1/3/5]  
user@host# set description "Private"  
user@host# edit unit 0 family inet  
[edit interfaces ge-1/3/5 unit 0 family inet]  
user@host# set service input service-set ss2  
user@host# set service output service-set ss2  
user@host# set address 9.0.0.1/24
```
2. Define the interface to the public Internet.

```
user@host# edit interfaces ge-1/3/6  
[edit interfaces ge-1/3/6]  
user@host# set description "Public"  
user@host# set unit 0 family inet address 128.0.0.1/24
```
3. Define the service interface for NAT processing.

```
user@host# edit interfaces sp-5/0/0  
[edit interfaces sp-5/0/0]  
user@host# set unit 0 family inet
```

```

Results user@host# show interfaces ge-1/3/5
description Private;
unit 0 {
    family inet {
        service {
            input {
                service-set sset2;
            }
            output {
                service-set sset2;
            }
        }
        address 9.0.0.1/24;
    }
}

user@host# show interfaces ge-1/3/6
description Public;;
unit 0 {
    family inet {
        address 128.0.0.1/24;
    }
}

user@host# show interfaces sp-5/0/0
unit 0 {
    family inet;
}

```

Example: NAT Configuration for the MS-MPC

This example shows how to configure network address translation with port translation (NAPT) on an MX series router using a MultiServices Modular Port Concentrator (MS-MPC) as a services interface card.

- [Requirements on page 77](#)
- [Overview on page 77](#)
- [Configuration on page 78](#)

Requirements

This example uses the following hardware and software components:

- MX-series router
- MultiServices Modular Port Concentrator (MS-MPC)
- Junos OS Release 13.2R1 or higher

Overview

A service provider has chosen an MS-MPC as a platform to provide NAT services to accommodate new subscribers.

Configuration

To configure NAPT44 using the MS-MPC as a services interface card, perform these tasks:

- [Configuring Interfaces on page 79](#)
- [Configure an Application Set of Acceptable ALG traffic on page 79](#)
- [Configuring a Stateful Firewall Rule on page 80](#)
- [Configuring NAT Pool and Rule on page 80](#)
- [Configuring the Service Set on page 81](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces ge-0/2/0 unit 0 family inet address 10.255.248.2/24
set interfaces xe-1/1/0 unit 0 family inet address 10.255.247.2/24
set interfaces xe-1/1/0 unit 0 family inet service input service-set sset1
set interfaces xe-1/1/0 unit 0 family inet service output service-set sset1
set interfaces ms-3/0/0 unit 0 family inet
set applications application-set accept-algs application junos-http
set applications application-set accept-algs application junos-ftp
set applications application-set accept-algs application junos-tftp
set applications application-set accept-algs application junos-telnet
set applications application-set accept-algs application junos-sip
set applications application-set accept-algs application junos-rtcp
set services stateful-firewall rule sf-rule1 match-direction input-output
set services stateful-firewall rule sf-rule1 term sf-term1 from source-address
    10.255.247.0/24
set services stateful-firewall rule sf-rule1 term sf-term1 from application-sets accept-algs
set services stateful-firewall rule sf-rule1 term sf-term1 then accept
set services nat pool napt-pool address 1.1.1.0/24
set services nat pool napt-pool port automatic
* nat rule for napt
set services nat rule nat-rule1 match-direction input
set services nat rule nat-rule1 term nat-term1 from source-address 10.255.247.0/24
set services nat rule nat-rule1 term nat-term1 from application-sets accept-algs
set services nat rule nat-rule1 term nat-term1 then translated source-pool napt-pool
set services nat rule nat-rule1 term nat-term1 then translated translation-type napt-44
* nat rule for basic nat
set services service-set sset1 stateful-firewall-rules sf-rule1
set services service-set sset1 nat-rules nat-rule1
set services service-set sset1 interface-service service-interface ms-3/0/0
```


Configuring Interfaces

- Step-by-Step Procedure** Configure the interfaces required for NAT processing. You will need the following interfaces:
- A customer-facing interface that handles traffic from and to the customer.
 - An internet-facing interface.
 - A services interface that provides NAT and stateful firewall services to the customer-facing interface
1. Configure the interface for the customer-facing interface.


```
user@host# edit
[edit ]
user@host# set interfaces xe-1/1/0 unit 0 family inet address 10.255.247.2/24
user@host# set interfaces xe-1/1/0 unit 0 family inet service input service-set sset1
user@host# set interfaces xe-1/1/0 unit 0 family inet service output service-set sset1
```
 2. Configure the interface for the Internet-facing interface.


```
[edit ]
set interfaces ge-0/2/0 unit 0 family inet address 10.255.248.2/24
```
 3. Configure the interface for the service set that will connect services to the customer-facing interface. In our example, the interface resides on an MS-MPC.


```
[edit ]
user@host# set interfaces ms-3/0/0 unit 0 family inet
```

Configure an Application Set of Acceptable ALG traffic

- Step-by-Step Procedure** Identify the acceptable ALGs for incoming traffic.
1. Specify an application set that contains acceptable incoming ALG traffic.


```
user@host# set applications application-set accept-algs application junos-http
user@host# set applications application-set accept-algs application junos-ftp
user@host# set applications application-set accept-algs application junos-tftp
user@host# set applications application-set accept-algs application junos-telnet
user@host# set applications application-set accept-algs application junos-sip
user@host# set applications application-set accept-algs application junos-rtcp
```

Results user@host#edit services applications application-set accept-algs
user@host#show
application junos-http;
application junos-ftp;
application junos-tftp;
application junos-telnet;
application junos-sip;
application junos-

Configuring a Stateful Firewall Rule

Step-by-Step Procedure Configure a stateful firewall rule that will accept all incoming traffic.

1. Specify firewall matching for all input and output
user@host# set services stateful-firewall rule sf-rule1 match-direction input-output
2. Identify source-address and acceptable ALG traffic from the customer-facing interface.
user@host# set services stateful-firewall rule sf-rule1 term sf-term1 from source-address 10.255.247.0/24
user@host# set services stateful-firewall rule sf-rule1 term sf-term1 from application-sets accept-algs
user@host# set services stateful-firewall rule sf-rule1 term sf-term1 then accept

Results user@host# edit services stateful-firewall
user@host# show
rule sf-rule1 {
 match-direction input-output;
 term sf-term1 {
 from {
 source-address {
 10.255.247.0/24;
 }
 application-sets accept-algs;
 }
 then {
 accept;
 }
 }
}

Configuring NAT Pool and Rule

Step-by-Step Procedure Configure a NAT pool and rule for address translation with automatic port assignment.

1. Configure the NAT pool with automatic port assignment.
user@host# set services nat pool napt-pool address 1.1.1.0/24
user@host# set services nat pool napt-pool port automatic
2. Configure a NAT rule that applies translation type **napt-44** using the defined NAT pool.
user@host# set services nat rule nat-rule1 term nat-term1 from application-sets accept-algs

```

user@host# set services nat rule nat-rule1 term nat-term1 then translated source-pool
napt-pool
user@host# set services nat rule nat-rule1 term nat-term1 then translated
translation-type napt-44

```

Results

```

user@host# edit services nat
user@host# show

pool napt-pool {
    address 1.1.1.0/24;
    port {
        automatic;
    }
}
rule nat-rule1 {
    match-direction input;
    term nat-term1 {
        from {
            source-address {
                10.255.247.0/24;
            }
            application-sets accept-algs;
        }
        then {
            translated {
                source-pool napt-pool;
                translation-type {
                    napt-44;
                }
            }
        }
    }
}

```

Configuring the Service Set

Step-by-Step Procedure Configure an interface type service set.

1. Specify the NAT and stateful firewall rules that apply to customer traffic.


```

user@host set services service-set sset1 stateful-firewall-rules sf-rule1
user@host set services service-set sset1 nat-rules bat-rule1

```
2. Specify the services interface that applies the rules to customer traffic.


```

set services service-set sset1 interface-service service-interface ms-3/0/0

```

Results

```

user@host# edit services service-set sset1
user@host# show
set services service-set sset1 stateful-firewall-rules sf-rule1
set services service-set sset1 nat-rules nat-rule1
set services service-set sset1 interface-service service-interface ms-3/0/0

```

Related Documentation

- *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*

Example: Configuring NAT-PT

A Domain Name System application-level gateway (DNS ALG) is used with Network Address Translation-Protocol Translation (NAT-PT) to facilitate name-to-address mapping. You can configure the DNS ALG to map addresses returned in the DNS response to an IPv6 address.

When you configure NAT-PT with DNS ALG support, you must configure two NAT rules or one rule with two terms. In this example, you configure two rules. The first NAT rule ensures that the DNS query and response packets are translated correctly. For this rule to work, you must configure a DNS ALG application and reference it in the rule. The second rule is required to ensure that NAT sessions are destined to the address mapped by the DNS ALG.

Then, you must configure a service set, and then apply the service set to the interfaces.

This example describes how to configure NAT-PAT with DNS ALG:

- [Requirements on page 82](#)
- [Overview and Topology on page 82](#)
- [Configuration of NAT-PT with DNS ALGs on page 84](#)

Requirements


This example uses the following hardware and software components:

- Junos OS Release 11.2
- A multiservices interface (**ms-**)

Overview and Topology

The following scenario shows the process of NAT-PT with DNS ALG when a laptop in an IPv6-only domain requests access to a server in an IPv4-only domain.

Figure 8: Configuring DNS ALGs with NAT-PT Network Topology



The Juniper Networks router in the center of the illustration performs address translation in two steps. When the laptop requests a session with the **www.example.com** server that is in an IPv4-only domain, the Juniper Networks router performs the following:

- Translates the IPv6 laptop and DNS server addresses into IPv4 addresses.
- Translates the AAAA request from the laptop into an A request so that the DNS server can provide the IPv4 address.

When the DNS server responds with the A request, the Juniper Networks router performs the following:

- Translates the IPv4 DNS server address back into an IPv6 address.
- Translates the A request back into a AAAA request so that the laptop now has the 96-bit IPv6 address of the **www.example.com** server.

After the laptop receives the IPv6 version of the **www.example.com** server address, the laptop initiates a second session using the 96-bit IPv6 address to access that server. The Juniper Networks router performs the following:

- Translates the laptop IPv4 address directly into its IPv4 address.
- Translates the 96-bit IPv6 **www.example.com** server address into its IPv4 address.

Configuration of NAT-PT with DNS ALGs

To configure NAT-PT with DNS ALG, perform the following tasks:

- [Configuring the Application-Level Gateway on page 84](#)
- [Configuring the NAT Pools on page 85](#)
- [Configuring the DNS Server Session: First NAT Rule on page 86](#)
- [Configuring the HTTP Session: Second NAT Rule on page 89](#)
- [Configuring the Service Set on page 91](#)
- [Configuring the Stateful Firewall Rule on page 93](#)
- [Configuring Interfaces on page 94](#)

Configuring the Application-Level Gateway

Step-by-Step Procedure

Configure the DNS application as the ALG to which the DNS traffic is destined. The DNS application protocol closes the DNS flow as soon as the DNS response is received. When you configure the DNS application protocol, you must specify the UDP protocol as the network protocol to match in the application definition.

To configure the DNS application:

1. In configuration mode, go to the **[edit applications]** hierarchy level.

```
user@host# edit applications
```

2. Define the application name and specify the application protocol to use in match conditions in the first NAT rule.

```
[edit applications]
user@host# set application application-name application-protocol protocol-name
```

For example:

```
[edit applications]
user@host# set application dns_alg application-protocol dns
```

3. Specify the protocol to match, in this case UDP.

```
[edit applications]
user@host# set application application-name protocol type
```

For example:

```
[edit applications]
user@host# set application dns_alg protocol udp
```

4. Define the UDP destination port for additional packet matching, in this case the domain port.

```
[edit applications]
user@host# set application application-name destination-port value
```

For example:

```
[edit applications]
user@host# set application dns_alg destination-port 53
```

Results [edit applications]
 user@host# show
 application dns_alg {
 application-protocol dns;
 protocol udp;
 destination-port 53;
 }

Configuring the NAT Pools

Step-by-Step Procedure In this configuration, you configure two pools that define the addresses (or prefixes) used for NAT. These pools define the IPv4 addresses that are translated into IPv6 addresses. The first pool includes the IPv4 address of the source. The second pool defines the IPv4 address of the DNS server. To configure NAT pools:

1. In configuration mode, go to the [edit services nat] hierarchy level.
 user@host# edit services nat
2. Specify the name of the first pool and the IPv4 source address (laptop).

```
[edit services nat]
user@host# set pool nat-pool-name address ip-prefix
```

For example:

```
[edit services nat]
user@host# set pool pool1 address 40.1.1.1/32
```

3. Specify the name of the second pool and the IPv4 address of the DNS server.

```
[edit services nat]
user@host# set pool nat-pool-name address ip-prefix
```

For example:

```
[edit services nat]
user@host# set pool pool2 address 50.1.1.1/32
```

Results The following sample output shows the configuration of NAT pools.

```
[edit services nat]
user@host# show
pool pool1 {
  address 40.1.1.1/32;
}
pool pool2 {
  address 50.1.1.1/32;
}
```

Configuring the DNS Server Session: First NAT Rule

Step-by-Step Procedure The first NAT rule is applied to DNS traffic going to the DNS server. This rule ensures that the DNS query and response packets are translated correctly. For this rule to work, you must configure a DNS ALG application and reference it in the rule. The DNS application was configured in [“Configuring the DNS ALG Application” on page 54](#). In addition, you must specify the direction in which traffic is matched, the source address of the laptop, the destination address of the DNS server, and the actions to take when the match conditions are met.

To configure the first NAT rule:

1. In configuration mode, go to the **[edit services nat]** hierarchy level.

```
user@host# edit services nat
```

2. Specify the name of the NAT rule.

```
[edit services nat]
user@host# edit rule rule-name
```

For example:

```
[edit services nat]
user@host# edit rule rule1
```

3. Specify the name of the NAT term.

```
[edit services nat rule rule-name]
user@host# edit term term-name
```

For example:

```
[edit services nat rule rule1]
user@host# edit term term1
```

4. Define the match conditions for this rule.

- a. Specify the IPv6 source address of the device (laptop) attempting to access an IPv4 address.

```
[edit services nat rule rule-name term term-name]
user@host# set from source-address source-address
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set from source-address 2000::2/128
```

- b. Specify the IPv6 destination address of the DNS server.

```
[edit services nat rule rule-name term term-name]
user@host# set from destination-address prefix
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set from destination-address 4000::2/128
```

- c. Reference the DNS application to which the DNS traffic destined for port 53 is applied.


```
[edit services nat rule rule1 term term1]
user@host# set from applications application-name
```

In this example, the application name configured in the *Configuring the DNS Application* step is `dns_alg`:

```
[edit services nat rule rule1 term term1]
user@host# set from applications dns_alg
```

5. Define the actions to take when the match conditions are met. The source and destination pools you configured in “[Configuring the NAT Pools](#)” on page 85 are applied here.

- a. Apply the NAT pool configured for source translation.

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

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated source-pool pool1
```

- b. Apply the NAT pool configured for destination translation.

```
[edit services nat rule rule-name term term-name]
user@host# set then translated destination-pool nat-pool-name
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated source-pool pool2
```

6. Define the DNS ALG 96-bit prefix for IPv4-to-IPv6 address mapping.

```
[edit services nat rule rule-name term term-name]
user@host# set then translated dns-alg-prefix dns-alg-prefix
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated dns-alg-prefix 10:10:10::0/96
```

7. Specify the type of NAT used for source and destination traffic.

```
[edit services nat rule rule-name term term-name]
user@host# set then translated translation-type basic-nat-pt
```

For example:

```
[edit services nat rule rule1 term term1]
user@host# set then translated translation-type basic-nat-pt
```



NOTE: In this example, since NAT is achieved using address-only translation, the `basic-nat-pt` translation type is used. To achieve NAT using address and port translation (NAPT), use the `napt-pt` translation type.

8. Specify the direction in which to match traffic that meets the rule conditions.

```
[edit services nat rule rule-name]  
user@host# set match-direction (input | output)
```

For example:

```
[edit services nat rule rule1]  
user@host# set match-direction input
```

9. Configure system logging to record information from the services interface to the /var/log directory.

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

For example:

```
[edit services nat rule rule1 term term1]  
user@host# set then syslog
```

Results The following sample output shows the configuration of the first NAT rule that goes to the DNS server.

```
[edit services nat]
user@host# show
rule rule1 {
  match-direction input;
  term term1 {
    from {
      source-address {
        2000::2/128;
      }
      destination-address {
        4000::2/128;
      }
      applications dns_alg;
    }
    then {
      translated {
        source-pool pool1;
        destination-pool pool2;
        dns-alg-prefix 10:10:10::0/96;
        translation-type {
          basic-nat-pt;
        }
      }
    }
    syslog;
  }
}
```

Configuring the HTTP Session: Second NAT Rule

Step-by-Step Procedure

The second NAT rule is applied to destination traffic going to the IPv4 server (www.example.com). This rule ensures that NAT sessions are destined to the address mapped by the DNS ALG. For this rule to work, you must configure the DNS ALG address map that correlates the DNS query or response processing done by the first rule with the actual data sessions processed by the second rule. In addition, you must specify the direction in which traffic is matched: the IPv4 address for the IPv6 source address (laptop), the 96-bit prefix to prepend to the IPv4 destination address (www.example.com), and the translation type.

To configure the second NAT rule:

1. In configuration mode, go to the following hierarchy level.

```
user@host# edit services nat
```

2. Specify the name of the NAT rule and term.

```
[edit services nat]
user@host# edit rule rule-name term term-name
```

For example:

```
[edit services nat]
user@host# edit rule rule2 term term1
```

3. Define the match conditions for this rule:
 - a. Specify the IPv6 address of the device attempting to access the IPv4 server.

```
[edit services nat rule rule-name term term-name]  
user@host# set from source-address source-address
```

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set from source-address 2000::2/128
```

- b. Specify the 96-bit IPv6 prefix to prepend to the IPv4 server address.

```
[edit services nat rule rule-name term term-name]  
user@host# set from destination-address prefix
```

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set from destination-address 10:10:10::c0a8:108/128
```

4. Define the actions to take when the match conditions are met.
 - Specify the prefix for the translation of the IPv6 source address.

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

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set then translated source-prefix 19.19.19.1/32
```

5. Specify the type of NAT used for source and destination traffic.

```
[edit services nat rule rule-name term term-name]  
user@host# set then translated translation-type basic-nat-pt
```

For example:

```
[edit services nat rule rule2 term term1]  
user@host# set then translated translation-type basic-nat-pt
```



NOTE: In this example, since NAT is achieved using address-only translation, the *basic-nat-pt* translation type is used. To achieve NAT using address and port translation (NAPT), you must use the *napt-pt* translation type.

6. Specify the direction in which to match traffic that meets the conditions in the rule.

```
[edit services nat rule rule-name]  
user@host# set match-direction (input | output)
```

For example:

```
[edit services nat rule rule2]  
user@host# set match-direction input
```

Results The following sample output shows the configuration of the second NAT rule.

```
[edit services nat]
user@host# show
rule rule2 {
  match-direction input;
  term term1 {
    from {
      source-address {
        2000::2/128;
      }
      destination-address {
        10:10:10::c0a8:108/128;
      }
    }
    then {
      translated {
        source-prefix 19.19.19.1/32;
        translation-type {
          basic-nat-pt;
        }
      }
    }
  }
}
```

Configuring the Service Set

Step-by-Step Procedure This service set is an interface service set used as an action modifier across the entire services (**ms-**) interface. Stateful firewall and NAT rule sets are applied to traffic processed by the services interface.

To configure the service set:

1. In configuration mode, go to the **[edit services]** hierarchy level.

```
user@host# edit services
```

2. Define a service set.

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

For example:

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

3. Specify properties that control how system log messages are generated for the service set.

```
[edit services service-set ss]
user@host# set syslog host local services severity-level
```

The example below includes all severity levels.

```
[edit services service-set ss]
user@host# set syslog host local services any
```

4. Specify the stateful firewall rule included in this service set.

```
[edit services service-set ss]
user@host# set stateful-firewall-rules rule1 severity-level
```

The example below references the stateful firewall rule defined in [“Configuring the Stateful Firewall Rule” on page 93](#).

```
[edit services service-set ss]
user@host# set stateful-firewall-rules rule1
```

5. Define the NAT rules included in this service set.

```
[edit services service-set ss]
user@host# set nat-rules rule-name
```

The example below references the two rules defined in this configuration example.

```
[edit services service-set ss]
user@host# set nat-rules rule1
user@host# set nat-rules rule2
```

6. Configure an adaptive services interface on which the service is to be performed.

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

For example:

```
[edit services service-set ss]
user@host# interface-service service-interface ms-2/0/0
```

Only the device name is needed, because the router software manages logical unit numbers automatically. The services interface must be an adaptive services interface for which you have configured **unit 0 family inet** at the `[edit interfaces interface-name]` hierarchy level in [“Configuring Interfaces” on page 94](#).

Results The following sample output shows the configuration of the service set.

```
[edit services]
user@host# show
service-set ss {
  syslog {
    host local {
      services any;
    }
  }
  stateful-firewall-rules rule1;
  nat-rules rule1;
  nat-rules rule2;
  interface-service {
    service-interface ms-2/0/0;
  }
}
```

Configuring the Stateful Firewall Rule

Step-by-Step Procedure This example uses a stateful firewall to inspect packets for state information derived from past communications and other applications. The NAT-PT router checks the traffic flow matching the direction specified by the rule, in this case both input and output. When a packet is sent to the services (**ms-**) interface, direction information is carried along with it.

To configure the stateful firewall rule:

1. In configuration mode, go to the **[edit services stateful firewall]** hierarchy level.

```
user@host# edit services stateful firewall
```

2. Specify the name of the stateful firewall rule.

```
[edit services stateful-firewall]
user@host# edit rule rule-name
```

For example:

```
[edit services stateful-firewall]
user@host# edit rule rule1
```

3. Specify the direction in which traffic is to be matched.

```
[edit services stateful-firewall rule rule-name]
user@host# set match-direction (input | input-output | output)
```

For example:

```
[edit services stateful-firewall rule rule1]
user@host# set match-direction input-output
```

4. Specify the name of the stateful firewall term.

```
[edit services stateful-firewall rule rule-name]
user@host# edit term term-name
```

For example:

```
[edit services stateful-firewall rule rule1]
user@host# edit term term1
```

5. Define the terms that make up this rule.

```
[edit services stateful-firewall rule rule-name term term-name]  
user@host# set then accept
```

For example:

```
[edit services stateful-firewall rule rule1 term term1]  
user@host# set then accept
```

Results The following sample output shows the configuration of the services stateful firewall.

```
[edit services]  
user@host# show  
stateful-firewall {  
  rule rule1 {  
    match-direction input-output;  
    term term1 {  
      then {  
        accept;  
      }  
    }  
  }  
}
```

Configuring Interfaces

Step-by-Step Procedure After you have defined the service set, you must apply services to one or more interfaces installed on the router. In this example, you configure one interface on which you apply the service set for input and output traffic. When you apply the service set to an interface, it automatically ensures that packets are directed to the services (**ms-**) interface.

To configure the interfaces:

1. In configuration mode, go to the **[edit interfaces]** hierarchy level.

```
user@host# edit interfaces
```

2. Configure the interface on which the service set is applied to automatically ensure that packets are directed to the services (**ms-**) interface.

- a. For IPv4 traffic, specify the IPv4 address.

```
[edit interfaces]  
user@host# set ge-1/0/9 unit 0 family inet address 30.1.1.1/24
```

- b. Apply the service set defined in [“Configuring Interfaces” on page 94](#).

```
[edit interfaces]  
user@host# set ge-1/0/9 unit 0 family inet6 service input service-set ss  
user@host# set ge-1/0/9 unit 0 family inet6 service output service-set ss
```

- c. For IPv6 traffic, specify the IPv6 address.

```
[edit interfaces]  
user@host# set ge-1/0/9 unit 0 family inet6 address 2000::1/64
```

3. Specify the interface properties for the services interface that performs the service.


```
[edit interfaces]
user@host# set ms-2/0/0 services-options syslog host local services any
user@host# set ms-2/0/0 unit 0 family inet
user@host# set ms-2/0/0 unit 0 family inet6
```

Results The following sample output shows the configuration of the interfaces for this example.

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

ge-1/0/9 {
  unit 0 {
    family inet {
      address 30.1.1.1/24;
    }
    family inet6 {
      service {
        input {
          service-set ss;
        }
        output {
          service-set ss;
        }
      }
      address 2000::1/64;
    }
  }
}

ms-2/0/0 {
  services-options {
    syslog {
      host local {
        services any;
      }
    }
  }
  unit 0 {
    family inet;
    family inet6;
  }
}
```

- Related Documentation**
- *Network Address Translation Overview for MS-DPC, MS-MPC, and MS-MIC Line Cards*
 - *Configuring Service Sets to be Applied to Services Interfaces*
 - *Example: Configuring Layer 3 Services and the Services SDK on Two PICs*
 - [dns-alg-prefix on page 131](#)
 - [dns-alg-pool on page 131](#)

Example: Configuring Inline Network Address Translation - Interface-Service Service Set

- [Requirements on page 96](#)
- [Overview on page 96](#)
- [Configuration on page 98](#)

Requirements

This example uses the following hardware and software components:

- MX-series router
- Modular Port Concentrator (MPC) with Trio chipset
- Junos OS Release 11.4R1 or higher

Overview

This example is configured for the network of a large financial services firm. This Application Service Provider (ASP) has an IP/MPLS-based backbone and provides L3VPN connectivity. In our example, the ASP acts like an Internet Service Provider (ISP) and its servers have public IPv4 addresses.

A large subscriber base relies heavily on the market data feeds that the ASP provides. Like many of the enterprise networks today, a private addressing scheme has been in place for majority of ASP's customers. NAT is required to maintain access to ASP's shared services.

Requirements for the solution include:

- Ease subscriber addressing challenges of their by providing NAT services in ASP's network.
- Support access to common services by a large number of customers, even when these are hosted across in different VRFs and use overlapping addresses.
- Provide high throughput, low latency packet forwarding with NAT enabled.
- Provide operational simplicity and efficiency.
- Reduce cost of operations.

By deploying Juniper's MX's inline NAT service, the ASP can offer scalable solutions with uncompromised performance that fit the requirements of financial markets customers. Operational cost can be dramatically reduced by eliminating the need for a dedicated services PIC. Enabling subscribers to keep their existing addressing scheme by outsourcing the address translation function to the ASP greatly simplifies their network operations.

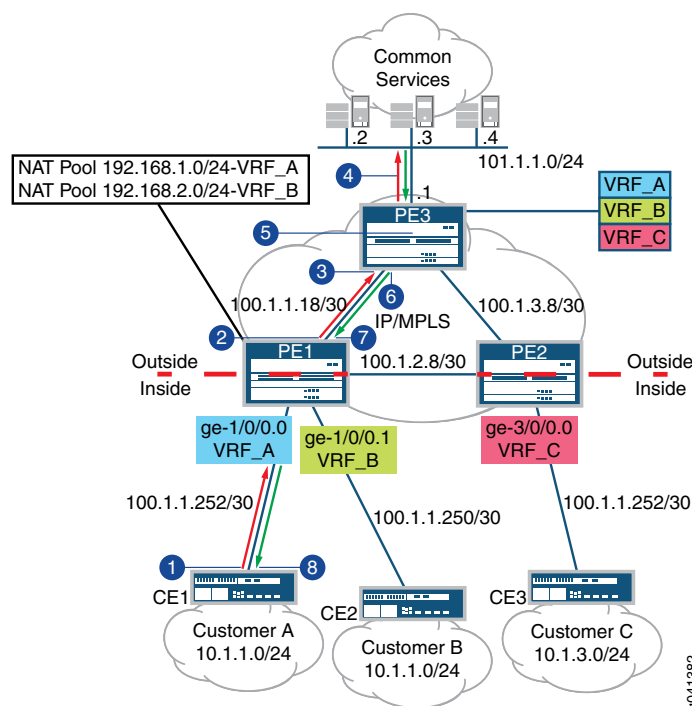
Topology

The topology for this application is show in [Figure 9 on page 97](#)

The ASP's shared services are located on LAN segment 101.1.1.0/24 behind PE3. PE1 and PE2 are used to connect subscribers. Traditional MPLS-VPN is deployed between provider edge routers. In the case of PE1, subscriber A and B have overlapping addressing schemes of 10.1.1.0/24; NAT is needed so the subscribers can access the same server. NAT pools 192.168.1.0/24 and 192.168.2.0/24 have been allocated to customer A and B respectively.

We will use host 10.1.1.2 from customer A to illustrate packet flow at a high level, as shown in Figure 9 on page 97

Figure 9: Deploy Inline NAT within L3VPN



1. CE1 forwards request from host 10.1.1.2 with a server destination of 101.1.1.2
2. With configured service set on PE1 for VRF_A, source address of 10.1.1.2 will be translated into 192.168.1.2. VPN label and IGP label will be imposed after the translation.
3. Packets will then be forwarded to PE3 using IGP label
4. PE3 receives the packet and performs a lookup in its VPN routing table. It then forwards the packets to server 101.1.1.2 after label disposition.
5. The server returns the packet with destination address of 192.168.1.2.
6. PE3 imposes VPN and IGP labels for the above destination and label switched the packets to PE1.
7. PE1 sends the packet to VRF_A after a FIB lookup. Destination address 192.168.1.2 will be translated 10.1.1.2.
8. CE1 receives the packets for host 10.1.1.2 and forwards them on.

Configuration

By using a **si-** (service-inline) interface, the operator can configure both **interface-service** and **next-hop** service-sets to perform inline NAT. This example uses the **interface-service** service set.

To configure inline NAT, perform these tasks:

- [Configure Interfaces on page 98](#)
- [Configuring Bandwidth for the Service Inline \(si-\) Interface on page 100](#)
- [Configuring NAT Pool and Rule on page 101](#)
- [Configuring the Service Set on page 103](#)

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```
set interfaces si-0/0/0 unit 0 family inet
set interfaces ge-1/0/0 unit 0 family inet service input service-set nat1
set interfaces ge-1/0/0 unit 0 family inet service output service-set nat1
set interfaces ge-1/0/0 unit 0 family inet address 10.1.1.1/24
set interfaces ge-1/0/1 unit 0 family inet service input service-set nat2
set interfaces ge-1/0/1 unit 0 family inet service output service-set nat2
set interfaces ge-1/0/1 unit 0 family inet address 192.168.1.1/24
set interfaces ge-1/0/2 unit 0 family inet service input service-set nat3
set interfaces ge-1/0/2 unit 0 family inet service output service-set nat3
set interfaces ge-1/0/2 unit 0 family inet address 10.1.1.100/24
set chassis fpc 0 pic 0 inline-services bandwidth 10g
set services nat pool p1 address 20.1.1.0/24
set services nat pool p2 address 21.1.1.2/32
set services nat pool p3 address 120.1.1.1/32
set services nat rule r1 match-direction input
set services nat rule r1 term t1 from source-address 10.1.1.0/24
set services nat rule r1 term t1 then translated source-pool p1 translation-type basic-nat44
set services nat rule r2 match-direction input
set services nat rule r2 term t1 from source-address 192.168.1.2/32
set services nat rule r2 term t1 then translated source-pool p2 translation-type basic-nat44
set services nat rule r3 match-direction input
set services nat rule r3 term t1 from source-address 10.1.1.8/32
set services nat rule r3 term t1 then translated source-pool p3 translation-type basic-nat44
set services service-set nat1 nat-rules r1
set services service-set nat1 interface-service service-interface si-0/0/0.0
set services service-set nat2 nat-rules r2
set services service-set nat2 interface-service service-interface si-0/0/0.0
set services service-set nat3 nat-rules r3
```

Configure Interfaces

Step-by-Step Procedure

To configure interfaces required for inline NAT:

1. Configure the inline interface for NAT services.

```
user@host# edit interfaces si-0/0/0
[edit interfaces si-0/0/0]
user@host# set unit 0 family inet
```

2. Configure the interface for traffic to be handled by service set nat1

```
user@host# edit interfaces ge-1/0/0
[edit interfaces ge-1/0/0]
user@host# edit unit 0 family inet service
[edit unit 0 family inet service]
user@host# set input service-set nat1 output service-set nat1
user@host# set address 10.1.1/24
```

3. Configure the interface for traffic to be handled by service set nat2

```
user@host# edit interfaces ge-1/0/1
[edit interfaces ge-1/0/0]
user@host# edit unit 0 family inet service input service
[edit unit 0 family inet service]
user@host# set input service-set nat2 output service-set nat2
user@host# set address 192.168.1/24
```

4. Configure the interface for traffic to be handled by service set nat3

```
user@host# edit interfaces ge-1/0/2
[edit interfaces ge-1/0/0]
user@host# edit unit 0 family inet service input service
[edit unit 0 family inet service]
user@host# set input service-set nat3 output service-set nat3
user@host# set address 10.1.1.100/24
```

```

Results si-0/0/0 {
        unit 0 {
            family inet;
        }
    }
ge-1/0/0 {
    unit 0 {
        family inet {
            service {
                input {
                    service-set nat1;
                }
                output {
                    service-set nat1;
                }
            }
            address 10.1.1.1/24;
        }
    }
}
ge-1/0/1 {
    unit 0 {
        family inet {
            service {
                input {
                    service-set nat2;
                }
                output {
                    service-set nat2;
                }
            }
            address 192.168.1.1/24;
        }
    }
}
ge-1/0/2 {
    unit 0 {
        family inet {
            service {
                input {
                    service-set nat3;
                }
                output {
                    service-set nat3;
                }
            }
            address 10.1.1.1/24;
        }
    }
}

```

Configuring Bandwidth for the Service Inline (si-) Interface

Step-by-Step Procedure

1. Go to the configuration hierarchy for the fpc and pic used for inline NAT services.

```

user@host# edit chassis fpc 0 pic 0
[edit chassis fpc - pic 0]

```
2. Set the bandwidth for inline services.

```

[edit chassis fpc 0 pic 0]

```

```
user@host# set inline-services bandwidth 10g
```

Configuring NAT Pool and Rule

Step-by-Step Procedure

1. Go to the services NAT hierarchy.

```
user@host# edit services nat
```
2. Configure three NAT pools.

```
[edit services nat]
user@host# set nat pool p1 address 20.1.1.0/24
user@host# set nat pool p2 address 21.1.1.2/32
user@host# set nat pool p3 address 120.1.1.1/32
```
3. Configure NAT rule for source pool p1.

```
[edit services nat]
user@host# set nat rule r1 match-direction input
user@host# set nat rule r1 term t1 from source-address 10.1.1.0/24 then
[nat pool r1 term t1 from source-address 10.1.1.0/24 then]
user@host# set translated source-pool p1 translation-type basic-nat44
```
4. Configure NAT rule for source pool p2.

```
[edit services nat]
user@host# set nat rule r2 match-direction input
user@host# edit nat rule r2 term t1 from source-address 192.168.1.2/32 then
[nat pool r2 term t1 from source-address 192.168.1.2/32 then]
user@host# set translated source-pool p2 translation-type basic-nat44.
```
5. Configure NAT rule for source pool p3.

```
[edit services nat]
user@host# set nat rule r3 match-direction input
user@host# edit nat rule r3 term t1 from source-address 10.1.1.8/32 then
[nat pool r1 term t1 from source-address 10.1.1.8/32 then]
user@host# set translated source-pool p1 translation-type basic-nat44
```

```
Results user@host# edit services nat
user@host# show

pool p1 {
    address 20.1.1.0/24;
}
pool p2 {
    address 21.1.1.2/32;
}
pool p3 {
    address 120.1.1.1/32;
}
rule r1 {
    match-direction input;
    term t1 {
        from {
            source-address {
                10.1.1.0/24;
            }
        }
        then {
            translated {
                source-pool p1;
                translation-type {
                    basic-nat44;
                }
            }
        }
    }
}
rule r2 {
    match-direction input;
    term t1 {
        from {
            source-address {
                192.168.1.2/32;
            }
        }
        then {
            translated {
                source-pool p2;
                translation-type {
                    basic-nat44;
                }
            }
        }
    }
}
rule r3 {
    match-direction input;
    term t1 {
        from {
            source-address {
                10.1.1.8/32;
            }
        }
        then {
            translated {
                source-pool p3;
                translation-type {
```



```

        basic-nat44;
    }
}
}
}
}

```

Configuring the Service Set

Step-by-Step Procedure

1. Configure a service set using NAT rule r1, associated with NAT pool p1.

```

user@host# edit services service-set nat1
[edit services service-set nat1]
user@host# set nat rules r1
user@host# set interface-service service-interface si-0/0/0.0

```
2. Configure a service set using NAT rule r2, associated with NAT pool p2.

```

user@host# edit services service-set nat2
[edit services service-set nat1]
user@host# set nat rules r2
user@host# set interface-service service-interface si-0/0/0.0

```
3. Configure a service set using NAT rule r3, associated with NAT pool p3.

```

user@host# edit services service-set nat3
[edit services service-set nat1]
user@host# set nat rules r3
user@host# set interface-service service-interface si-0/0/0.0

```

Results

```

user@host# edit services service-set nat1
user@host# show
nat-rules r1;
interface-service {
    service-interface si-0/0/0.0;
}

user@host# edit services service-set nat2
user@host# show
nat-rules r2;
interface-service {
    service-interface si-0/0/0.0;
}

user@host# edit services service-set nat3
user@host# show
nat-rules r3;
interface-service {
    service-interface si-0/0/0.0;
}

```

Related Documentation

- *Inline Network Address Translation Overview for MPC Types 1, 2, and 3*

Port Control Protocol Configuration Examples

This topic contains the following Port Control Protocol (PCP) configuration examples.

- [Example: Configuring Port Control Protocol with NAPT44 on page 104](#)

Example: Configuring Port Control Protocol with NAPT44

- [Requirements on page 104](#)
- [Overview on page 104](#)
- [PCP Configuration on page 104](#)

Requirements

Hardware Requirements

- UEs with PCP clients.
- An MX 3D Router with an MS-DPC services PIC.

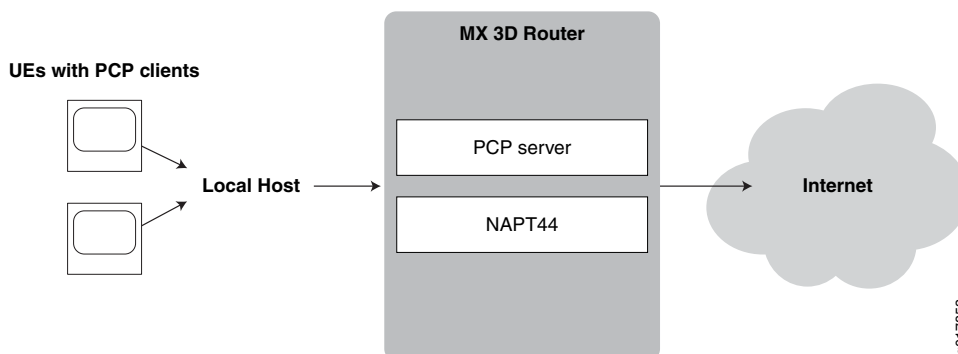
Software Requirements

- Junos OS 13.2
- Layer-3 Services Package

Overview

An ISP wants to enable UEs with PCP clients to maintain connections to servers without timing out. The PCP clients generate PCP requests for the type and duration of the connection they require. Connections may be of a long duration, such as applications using a webcam, or a shorter duration, such as online games. An MX 3D router provides a PCP server to interpret PCP client requests, and NAPT44. [Figure 10 on page 104](#) shows the basic topology for this example.

Figure 10: PCP with NAPT44



PCP Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network

configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level.

```

set chassis fpc 2 pic 0 adaptive-services service-package layer-3
set interfaces sp-2/0/0 services-options inactivity-timeout 180 cgn-pic
set interfaces sp-2/0/0 unit 0 family inet
set interfaces xe-3/2/0 unit 0 family inet service input service-set sset_0
set interfaces xe-3/2/0 unit 0 family inet service output service-set sset_0
set interfaces xe-3/2/0 unit 0 family inet address 30.0.0.1/24
set interfaces xe-5/0/0 unit 0 family inet address 25.0.0.1/24
set services nat pool pcp-pool address 44.0.0.0/16
set services nat pool pcp-pool port automatic random-allocation address-allocation
    round-robin
set services nat pool pcp-pool address-allocation round-robin
set services nat rule pcp-rule match-direction input
set services nat rule pcp-rule term t0 then translated source-pool pcp-pool
    translation-type napt-44
set services nat rule pcp-rule term t0 then translated mapping-type endpoint-independent
    filtering-type endpoint-independent
set services nat rule pcp-rule term t0 then translated mapping-type endpoint-independent
    filtering-type endpoint-independent
set services pcp server pcp-s1 ipv4-address 124.124.124.122 mapping-lifetime-minimum
    600 mapping-lifetime-minimum 600
set services pcp server pcp-s1 mapping-lifetime-minimum 600
    mapping-lifetime-maximum 86500
set services pcp server pcp-s1 short-lifetime-error 120 long-lifetime-error 1200
set services pcp server pcp-s1 max-mappings-per-client 128 pcp-options third-party
    prefer-failure
set services service-set sset_0 pcp-rules r1
set services service-set sset_0 nat-rules pcp-rule
set services service-set sset_0 interface-service service-interface sp-2/0/0.0

```

Chassis Configuration

Step-by-Step Procedure To configure the service PIC (FPC 2 Slot 0) with the Layer 3 service package:

1. Go to the [edit chassis] hierarchy level.

```
user@host# edit chassis
```
2. Configure the Layer 3 service package.

```
[edit chassis]
user@host# set fpc 2 pic 0 adaptive-services service-package layer-3
```

Results user@host# show chassis fpc 2 pic 0

```
pcp-rules pcp-napt44-rule;
nat-rules pcp-rule;
interface-service {
    service-interface sp-2/0/0.0;
}
```

Interface Configuration

**Step-by-Step
Procedure**

1. Configure the services MS-DPC.

 user@host# set interfaces sp-2/0/0 services-options inactivity-timeout 180 cgn-pic
 user@host# set interfaces sp-2/0/0 unit 0 family inet
2. Configure the customer-facing interface used for NAT and PCP services.

 user@host# set interfaces xe-3/2/0 unit 0 family inet service input service-set
 sset_0
 user@host# set interfaces xe-3/2/0 unit 0 family inet service output service-set
 sset_0
 user@host# set interfaces xe-3/2/0 unit 0 family inet address 30.0.0.1/24
3. Configure the Internet-facing interface.

 user@host# set interfaces xe-5/0/0 unit 0 family inet address 25.0.0.1/24

```

Results user@host#
sp-2/0/0 {
  services-options {
    inactivity-timeout 180;
    cgn-pic;
  }
  unit 0 {
    family inet;
  }
}
xe-3/2/0 {
  unit 0 {
    family inet {
      service {
        input {
          service-set sset_0;
        }
        output {
          service-set sset_0;
        }
      }
      address 30.0.0.1/24;
    }
  }
}
xe-5/0/0 {
  unit 0 {
    family inet {
      address 25.0.0.1/24;
    }
  }
}

```

NAT Configuration

Step-by-Step Procedure

1. Go the [edit services nat] hierarchy.

```

user@host# edit services nat

```
2. Configure a NAT pool called **pcp-pool**.

```

[edit services nat]
user@host# set pool pcp-pool address 44.0.0.0/16
user@host# set pool pcp-pool port automatic random-allocation
user@host# set pool pcp-pool address-allocation round-robin

```
3. Configure a NAT rule called **pcp-rule**.

```

[edit services nat]
user@host# set rule pcp-rule term t0 then translated source-pool pcp-pool
translation-type napt-44
user@host# set rule pcp-rule term t0 then translated mapping-type
endpoint-independent filtering-type endpoint-independent

```

```

Results user@host# show services nat
pool pcp-pool {
    address 44.0.0.0/16;
    port {
        automatic {
            random-allocation;
        }
    }
    address-allocation round-robin;
}
rule pcp-rule {
    match-direction input;
    term t0 {
        then {
            translated {
                source-pool pcp-pool;
                translation-type {
                    napt-44;
                }
                mapping-type endpoint-independent;
                filtering-type {
                    endpoint-independent;
                }
            }
        }
    }
}

```

PCP Configuration

Step-by-Step Procedure To configure the PCP server and PCP rule options.

1. Go to the **edit services pcp** hierarchy level for server **pcp-s1**

```
user@host# edit services pcp server pcp-s1
```
2. Configure the PCP server options.

```
[edit services pcp server pcp-s1]
user@host# set ipv4-address 124.124.124.122
user@host# set mapping-lifetime-minimum 600
user@host# set mapping-lifetime-maximum 86500
user@host# set short-lifetime-error 120
user@host# set long-lifetime-error 1200
user@host# set max-mappings-per-client 128
user@host# set pcp-options third-party prefer-failure
```
3. Create the PCP rule.

```
[edit services pcp rule pcp-napt44-rule]
user@host# edit rule pcp-napt44-rule
```
4. Configure the PCP rule options.

```
[edit services pcp rule pcp-napt44-rule]
user@host# set match-direction input
user@host# set term t0 then pcp-server pcp-s1
```

Results regress@montag# show services pcp

```
server pcp-s1 {
    ipv4-address 124.124.124.122;
    mapping-lifetime-minimum 600;
    mapping-lifetime-maximum 86500;
    short-lifetime-error 120;
    long-lifetime-error 1200;
    max-mappings-per-client 128;
    pcp-options third-party prefer-failure;
}
rule pcp-napt44-rule {
    match-direction input;
    term t0 {
        then {
            pcp-server pcp-s1;
        }
    }
}
```

Service Set Configuration

Step-by-Step Procedure 1. Create a service set, **sset_0**, at the **edit services service-set** hierarchy level.

```
user@router# edit services service-set sset_0

service-set sset_0 {
    pcp-rules pcp-napt44-rule;
    nat-rules pcp-rule;
    interface-service {
        service-interface sp-2/0/0.0;
    }
}
```

2. Identify the NAT rule associated with the service set.

```
[edit services service-set sset_0]
user@router# set nat-rules pcp-rule
```

3. Identify the PCP rule associated with the service set.

```
[edit services service-set sset_0]
user@router# set pcp-rules r1
```

4. Identify the service interface associated with the service set.

```
[edit services service-set sset_0]
user@router# set interface-service service-interface sp-2/0/0.0
```

Results user@host# show

```
pcp-rules pcp-napt44-rule;
nat-rules pcp-rule;
interface-service {
    service-interface sp-2/0/0.0;
}
```


CHAPTER 7

Carrier-Grade NAT Implementation Best Practices

- [Carrier-Grade NAT Implementation: Best Practices on page 111](#)

Carrier-Grade NAT Implementation: Best Practices

The following topics present the best practices for carrier-grade NAT implementation on MS-DPCs using the Layer 3 services package:

- [Use APP and Round-Robin Address-Allocation on page 111](#)
- [Do Not Use EIM with SIP on page 112](#)
- [Do Not Use EIM with HTTP, DNS, or When Not Needed on page 112](#)
- [Define PBA Blocks Based on User Profiles on page 113](#)
- [Do Not Change the PBA Configuration on Running Systems on page 114](#)
- [Do Not Allocate Excessively Large NAT Pools on page 115](#)
- [Configure the System Log for PBA Only When Needed on page 115](#)
- [Use Redundant Service PIC \(RSP\) Interfaces for Failover on page 117](#)
- [Contain the Effects of Missing IP Fragments on page 118](#)
- [Do Not Use Configurations Prone to Routing Loops on page 118](#)

Use APP and Round-Robin Address-Allocation

Scenario:

- Address-pooling paired (APP) allows a private IP address to be mapped to the same public IP address from a NAT pool for all its sessions. The binding between private IP and public IP is triggered by the first packet seen from such private host.
- By default, an MS-DPC or MS-PIC allocates ports from a NAT pool in a sequential fashion from each consecutive IP address available in the pool.
- Sequential allocation, together with APP, can result in mapping multiple private hosts to the same public IP address, resulting in fast port exhaustion for the interested public IP address while other ports are still available from the remaining of NAT pool.



BEST PRACTICE: Configure round-robin address allocation for the NAT pool used by traffic served with APP. Round-robin allocation allocates ports from different IP addresses.

The following snippet provides an example of round-robin address allocation.

```
user@router# show services nat pool natpool-1
address-range low 9.9.9.1 high 9.9.9.10;
port {
    automatic;
}
address-allocation round-robin;
mapping-timeout 120;
```

Do Not Use EIM with SIP

Scenario:

- Session Initiation Protocol (SIP) traffic requires an Application Level Gateway (ALG) to allow SIP servers and clients on the public side of the CGNAT to communicate with the SIP hosts on the private side.
- The SIP ALG opens the pinholes in the CGNAT router to permit the forwarding of outbound traffic based on any supported SIP feature.
- Endpoint-independent mapping (EIM) is not needed by SIP to function, nor by the SIP ALG to create the flows for forwarding the SIP traffic



BEST PRACTICE: Do *not* configure EIM together with the SIP ALG; doing so adds processing overhead with no benefit.

```
user@router# show services nat rule natrule-1
match-direction input;
term 1 {
    from {
        applications junos-sip;
    }
    then {
        translated {
            source-pool natpool-3;
            translation-type {
                napt-44;
            }
            address-pooling paired;
        }
    }
}
```

Do Not Use EIM with HTTP, DNS, or When Not Needed

Scenario:

- Most Internet traffic uses HTTP, and there is no browser on any OS that reuses the same source port for sending traffic to different destinations. EIM provides no benefit for HTTP traffic.

- Because none of the junos-algs require EIM to work, avoid using EIM with the ALGs.
- EIM allocates memory for each mapping; this is in addition to the memory used for flow allocation. This reduces the maximum number of flows that can be established through the services PIC, and causes processing overhead for the creation and deletion of flows and mappings.



BEST PRACTICE:

- Don't enable EIM for applications that are defined ALGs or are known not to use Session Traversal Utilities for NAT (STUN) servers to discover the presence of a NAT router.
- Enable EIM for applications that do reuse the source ports and rely on a CGNAT device to maintain the same address:port mapping for all traffic sent to different destinations, such as on-line gaming applications like Xbox and PS3, or applications that use unilateral self-address fixing methods (UNSAF). see (*IETF RFC 3424 IAB Considerations for UNilateral Self-Address Fixing (UNSAF) Across Network Address Translation*).

Define PBA Blocks Based on User Profiles

Scenario:

- When a user connects to a website that requires the establishment of a significant number of sockets for a single HTML page, a corresponding number of new ports must be allocated. Port blocks should be large enough to prevent continual allocation of new blocks.
- If the number of concurrent sessions exceeds the number of ports available in the active port block, the other allocated port-blocks will be scanned for available ports to use or a new block will be allocated from the free block pool.
- The process of continually scanning the allocated port-blocks and/or allocating additional blocks from the free block pool could result in experienced latency for setting up new sessions and delay loading of web pages.
- Having a user continuously allocating or de-allocating from different PBA blocks impacts performance.



BEST PRACTICE: Define PBA blocks with a size that is a power of 2 or 4 related to the average number of sessions a user is expected to have active. For example, if a user is expected to have an average of approximately 200 to 250 sessions active, configuring the PBA block size to 512 or 1024 will provide a liberal allocation.

```
user@router# show services nat pool natpool-1
address-range low 9.9.9.1 high 9.9.9.10;
port {
    automatic;
```

```
secure-port-block-allocation {  
    block-size 1024;  
    max-blocks-per-user 8; /* Max 2048, default 8 */  
    active-block-timeout 300;  
}  
}  
mapping-timeout 300;
```

Do Not Change the PBA Configuration on Running Systems

Scenario:

- PBA settings in NAT pools are mapped to memory at the time of the Service PIC boot up and cannot be changed while processing traffic.
- Do not change the following settings:
 - Update any NAT pool PBA configuration.
 - Change a PBA NAT pool to a non-PBA NAT pool.
 - Change a non-PBA NAT pool to a PBA NAT pool.

Any of these changes result in the logging of the following message:

PBA_CATASTROPIC_CHANGE: The recent PBA configuration changes will reflect in the Service-PIC only after deactivate and activate of the service-set again



NOTE: Starting with Junos OS Release 14.1, when you deactivate a service-set that contains address pooling paired (APP) or 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 APP or 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.



BEST PRACTICE: When changing PBA configurations, restart the services PIC if possible. Minimally, you must deactivate and reactivate the affected service set.

Do Not Allocate Excessively Large NAT Pools

Scenario:

- The maximum number of flows supported by the MS-DPC and each PIC on an MS-DPC is 8 million.
- Assuming that the 8 million flow maximum consists of 4 million sessions (1 reverse flow for each forward flow), these sessions would require a maximum of 4 million ports that are available from 64 IP addresses within the 1024 to 65,535 ports range (64K ports per IP address).
- Do not configure ports to support more than 8 million flows; they will never be needed.
- This scenario assumes that APP, EIM, and EIF are not enabled. When they are enabled, the total number of flows is lower, which means that you should configure the number of IP addresses in the NAT pool based on the maximum supported flows.



BEST PRACTICE: Do not configure NAT pools with more than 64 addresses (that is, a /26 network) and round-robin configured and 64K ports from each address.

Configure the System Log for PBA Only When Needed

Scenario:

- Session logging can negatively affect performance depending on the frequency of creation and deletion of flows.
- PBA is meant to reduce the need for logging.
- Deterministic NAT is designed to eliminate the need for logging.
- All system log messages created by the services PIC constitutes traffic that will be sent to the Packet Forwarding Engine, competing with user traffic to reach the external destination.



BEST PRACTICE:

- Use logging to the system log at the service-set level rather than at the services PIC interface level when possible.
- Do not enable logging for redundant information. When using PBA, you don't need to configure logs per session because knowing the PBA block and the block size enables you to derive the ports allocated to each user. In this case, a log that reports all sessions created by that user with ports belonging to a block is redundant. If you have configured deterministic NAT (DetNat) a log is completely unnecessary because all information on port allocation can be deduced mathematically.

- Rate-limit the number of logs generated from an sp- interface. When not set, the default limits apply: 10K for the local host system log server (RE) and 200K for the external system log server.

```
user@router# show interfaces sp-1/1/0 services-options
system log {
  host 1.2.3.4 {
    services info;
  }
  message-rate-limit 1000;
}
```

- Always system log to an external server to avoid loading the Routing Engine and specify system log class to restrict logging.
 - If you do not specify system log class, all log messages are allowed (subject to priority check and rate limiting).
 - When you specify system log class, only messages meeting the class criteria are retained.
 - Use the `show services service-sets statistics system log detail` command to check what is being dropped by unconfigured classes.

```
user@router# show services service-set S-SET-1 system log
host 1.2.3.4 {
  services info;
  class {
    session-logs open close;
    packet-logs;
    stateful-firewall-logs;
    alg-logs;
    nat-logs;
    ids-logs;
  }
}
```



BEST PRACTICE: System log generation can be *rule-based* or *event-based*.

- Use rule-based system logging with care; it generates a log for every packet that enters the rule term, since rule-based logging is not subject to class or priority filtering.
- System log messages can be dropped only as a result of message rate limiting. Make sure you have set a realistic rate-limit that is unlikely to be exceeded.
- Use rule-based logging only for discarded traffic (a relatively small percentage of the traffic) or for troubleshooting. Since rule-based logging applies to all traffic that enters the PIC and creates a flow, logging can be excessive, resulting in reaching the configured induce rate limit with a consequent loss of needed logs.

```
cli# show services stateful-firewall
rule rule-sfw-accept {
  match-direction input-output;
  term term-sfw-accept {
```

```

        then {
            accept;
            system log;
        }
    }
}
rule rule-sfw-reject {
    match-direction input-output;
    term term-sfw-reject {
        then {
            reject;
            system log;
        }
    }
}
}

```

**BEST PRACTICE:**

All rule match logs are enabled by their respective rules:

- ASP_COS_RULE_MATCH (class-of-service rules)
- ASP_COS_RULE_MATCH (class-of-service rules)
- ASP_IDS_RULE_MATCH (ids rules)
- ASP_NAT_RULE_MATCH (nat rule)
- ASP_SFW_RULE_ACCEPT (stateful firewall rules)
- ASP_SFW_RULE_DISCARD
- ASP_SFW_RULE_REJECT

Use Redundant Service PIC (RSP) Interfaces for Failover

**BEST PRACTICE:**

- The usage of Redundant Service PIC (RSP) interfaces, allows the active services PIC to perform an immediated switchover to the secondary services PIC in case of major issues that require a services PIC reboot.
- This results in a minimal service impact for user traffic.
- There are two modes for redunancy: warm-standby (default) and hot-standby. Hot-standby provides 1:1 redundancy, while warm-standby provides 1:N redundancy. With both modes , there is no impact on the UDP forwarding.
- When the secondary services PIC is shared among multiple RSPs, only warm-standby is possible and the impact to traffic is limited to the time to load the appropriate configuration on the secondary PIC.

```
user@router# show interfaces rsp0
```

```
redundancy-options {  
  primary sp-0/1/0;  
  secondary sp-1/1/0;  
  hot-standby;  
}
```

Contain the Effects of Missing IP Fragments

Scenario:

- IP fragments are buffered as they arrive to facilitate the integrity check of the completely reassembled packet before being serviced by the services PIC.
- Missing fragments cause received fragments to be held until the internal buffer is full and are flushed out. This causes CPU usage overhead and reduced traffic forwarding.



BEST PRACTICE: Configure the `fragment-limit`, the maximum number of fragments for a packet, and `reassembly-timeout`, the maximum wait for a missing fragment, after which all other fragments for the same packet are flushed out.

```
user@router# show interfaces sp-0/0/0  
services-options {  
  open-timeout 5;  
  close-timeout 5;  
  inactivity-timeout 30;  
  tcp-tickles 4;  
  fragment-limit 10;  
  reassembly-timeout 3;  
  cgn-pic;  
}
```

Do Not Use Configurations Prone to Routing Loops

Scenario:

- Sudden and persistent high CPU usage is most likely an indication of packet looping between the Packet Forwarding Engine and the services PIC. Depending on whether the configuration uses interface-style or next-hop-style service sets, different network flaps can lead to routing loops.



BEST PRACTICE:

Ensure that only the intended traffic is allowed to reach the services PIC and is serviced based on service set rule.

- Configure a firewall filter that accepts only the traffic meant to go to the services PIC on the output direction of the sp- interface. That is, accept only traffic identified in the NAT rule from option as received from the

source-address that identifies the customer private network; discard and log all the rest.

- Allow only intended traffic to be serviced by the service set by configuring the stateful-firewall rules and NAT rules to translate only the traffic from the customer private source address ranges and intended applications. Although this does not prevent unintended traffic from being processed by the services PIC, it prevents the creation of flows, objects, and states that are not consistent with the expected traffic and are likely to be problematic.

**Related
Documentation**

- *Configuring Address Pools for Network Address Port Translation (NAPT) Overview*

CHAPTER 8

NAT Configuration Statements

- [address \(Services NAT Pool\) on page 122](#)
- [address-allocation on page 123](#)
- [address-range on page 123](#)
- [allow-overlapping-nat-pools on page 124](#)
- [app-mapping-timeout on page 124](#)
- [application-sets \(Services NAT\) on page 125](#)
- [applications \(Services NAT\) on page 125](#)
- [cgn-pic on page 126](#)
- [destination-address on page 126](#)
- [destination-address-range on page 127](#)
- [destination-pool on page 127](#)
- [destination-port range on page 128](#)
- [destination-prefix on page 128](#)
- [destination-prefix-list on page 129](#)
- [destined-port on page 129](#)
- [deterministic-port-block-allocation on page 130](#)
- [dns-alg-pool on page 131](#)
- [dns-alg-prefix on page 131](#)
- [ei-mapping-timeout on page 132](#)
- [eif-flow-limit on page 132](#)
- [from \(Services NAT\) on page 133](#)
- [ipv6-multicast-interfaces on page 134](#)
- [mapping-refresh on page 134](#)
- [mapping-timeout on page 135](#)
- [match-direction on page 135](#)
- [no-translation on page 136](#)
- [overload-pool on page 136](#)
- [overload-prefix on page 137](#)

- [pool](#) on page 138
- [port](#) on page 139
- [port-forwarding](#) on page 140
- [port-forwarding-mappings](#) on page 140
- [ports-per-session](#) on page 141
- [rule](#) on page 142
- [rule-set](#) on page 143
- [secure-nat-mapping](#) on page 143
- [secured-port-block-allocation](#) on page 144
- [server \(pcp\)](#) on page 145
- [services \(NAT\)](#) on page 146
- [service-set \(Services\)](#) on page 147
- [source-address \(NAT\)](#) on page 149
- [source-address-range](#) on page 149
- [source-pool](#) on page 150
- [source-prefix](#) on page 150
- [source-prefix-list](#) on page 151
- [syslog](#) on page 151
- [translated-port](#) on page 152
- [term](#) on page 153
- [then](#) on page 154
- [translated](#) on page 155
- [translation-type](#) on page 156

address (Services NAT Pool)

Syntax	<code>address <i>ip-prefix</i></<i>prefix-length</i>>;</code>
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. <i>prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5.
Description	Specify the NAT pool prefix value.
Options	<i>prefix</i> —Specify an IPv4 or IPv6 prefix value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Source and Destination Addresses Network Address Translation Overview</i>

address-allocation

Syntax	address-allocation round-robin;
Hierarchy Level	[edit services nat pool <i>pool-name</i>]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	When you use round-robin allocation, one port is allocated from each 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.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Pools of Addresses and Ports for Network Address Translation Overview</i>

address-range

Syntax	address-range low <i>minimum-value</i> high <i>maximum-value</i> ;
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv6 addresses in Junos OS Release 8.5.
Description	Specify the NAT pool address range.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Source and Destination Addresses Network Address Translation Overview</i>

allow-overlapping-nat-pools

Syntax	allow-overlapping-nat-pools;
Hierarchy Level	[edit services nat]
Release Information	Statement introduced with Junos OS Release 12.1.
Description	Specify that NAT source or destination pools can be shared between multiple service sets.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Service Sets for Network Address Translation</i>

app-mapping-timeout

Syntax	app-mapping-timeout <i>app-mapping-timeout</i> ;
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	mapping-timeout statement introduced in JUNOS Release 12.3.
Description	Specify the duration for address pooling paired (AP-P) mappings that use the specified NAT pool. If this option is not configured and a timeout value is configured for mapping-timeout , the timeout value configured for mapping-timeout is used. If neither option is specified, the default value of 300 seconds is used.
Options	app-mapping-timeout —Lifetime of AP-P mappings in seconds. Default: 300 Range: 120 through 864,000
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring Source and Destination Addresses Network Address Translation Overview</i>

application-sets (Services NAT)

Syntax	<code>applications-sets <i>set-name</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more target application sets.
Options	<i>set-name</i> —Name of the target application set.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Network Address Translation Rules Overview</i>

applications (Services NAT)

Syntax	<code>applications [<i>application-names</i>];</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define one or more application protocols to which the NAT services apply.
Options	<i>application-name</i> —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	<ul style="list-style-type: none"> • <i>Network Address Translation Rules Overview</i>

cg-n-pic

Syntax	cg-n-pic;
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	Restrict usage of the service PIC to carrier-grade NAT (CGN) or associated services (intrusion detection, stateful firewall, and software). All memory is available for CGN or related services and can be used for CGN scaling.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Carrier-Grade NAT Feature Comparison for Junos Address Aware by Type of Interface Card on page 9

destination-address

Syntax	destination-address (<i>address</i> any-unicast) <except>;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced before Junos OS Release 7.4. any-unicast and except options introduced in Junos OS Release 7.6. address option enhanced to support IPv6 and addresses in Junos OS Release 8.5.
Description	Specify the destination address for rule matching.
Options	address —Destination IPv4 or IPv6 address or prefix value. any-unicast —Any unicast packet. except —(Optional) Prevent the specified address, prefix, or unicast packets from being translated.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Network Address Translation Rules Overview

destination-address-range

Syntax	<code>destination-address-range low <i>minimum-value</i> high <i>maximum-value</i> <except>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in Junos OS Release 7.6. <i>minimum-value</i> and <i>maximum-value</i> options enhanced to support IPv6 addresses in Junos OS Release 8.5.
Description	Specify the destination address range for rule matching. If the translation-type statement in the then statement of the nat rule is set to stateful-nat-64 , the destination address range for rule matching must be within the range specified by the destination-prefix statement in the then statement.
Options	<i>minimum-value</i> —Lower boundary for the IPv4 or IPv6 address range. <i>maximum-value</i> —Upper boundary for the IPv4 or IPv6 address range. <i>except</i> —(Optional) Prevent the specified address range from being translated.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Network Address Translation Rules Overview

destination-pool

Syntax	<code>destination-pool <i>nat-pool-name</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the destination address pool for translated traffic.
Options	<i>nat-pool-name</i> —Destination pool name.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> Network Address Translation Rules Overview

destination-port range

Syntax	<code>destination-port range <i>high</i> <i>low</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from (Services NAT)]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the destination port range for rule matching.
Options	<i>high</i> —Upper limit of port range for matching. <i>low</i> —Lower limit of port range for matching.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Port Forwarding for Static Destination Address Translation on page 65

destination-prefix

Syntax	<code>destination-prefix <i>destination-prefix</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in Junos OS Release 7.6. <i>destination-prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5.
Description	Specify the destination prefix for translated traffic.
Options	<i>destination-prefix</i> —IPv4 or IPv6 destination prefix value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Network Address Translation Rules Overview


destination-prefix-list

Syntax	<code>destination-prefix-list <i>list-name</i> <except>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in Junos OS Release 8.2.
Description	<p>Specify the destination prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level.</p> <p>If the translation-type statement in the then statement of the nat rule is set to stateful-nat-64, the destination prefix list for rule matching must be within the range specified by the destination-prefix statement in the then statement.</p>
Options	<p><i>list-name</i>—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Network Address Translation Rules Overview</i> <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>

destined-port

Syntax	<code>destined-port <i>port id</i>;</code>
Hierarchy Level	[edit services nat port-forwarding <i>map-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the port from where traffic has to be forwarded.
Options	<i>port id</i> —The destination port number from where traffic will be forwarded.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> port-forwarding on page 140 translated-port on page 152

deterministic-port-block-allocation

Syntax	<pre>deterministic-port-block-allocation { block-size <i>block-size</i>; include-boundary-addresses; }</pre>
Hierarchy Level	[edit services nat pool <i>pool-name</i> port]
Release Information	Statement introduced in Junos OS Release 12.1.
Description	Configure algorithm-based allocation of blocks of destination ports. By specifying this method, you ensure that an incoming (source) IP address and port always map to the same destination IP address and port block, thus eliminating the need for logging address translations.
Options	<i>block-size</i> —Maximum number of ports that can be allocated to a user.
<hr/>	
<div> NOTE: When a block-size of 0 is specified, block size is calculated according to the formula: $(64512 * \text{Number of IP addresses in the NAT Pool}) / \text{Number of subscribers}$ where</div> <ul style="list-style-type: none">• 64512 is derived from (65535 - 1023) because the regular port assignments start from 1024.• Number of subscribers is derived from the from clause of the applicable NAT rule. <div><hr/></div>	
Default: 256	
Range: 0 through 32,000	
include-boundary-addresses —(Optional) Specifies that the lowest and highest addresses in the source address range of a NAT rule should be translated when the NAT pool is used.	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Deterministic Port Block Allocation on page 51

dns-alg-pool

Syntax	<code>dns-alg-pool <i>dns-alg-pool</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the Network Address Translation (NAT) pool for destination translation.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

dns-alg-prefix

Syntax	<code>dns-alg-prefix <i>dns-alg-prefix</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Set the Domain Name System (DNS) application-level gateway (ALG) 96-bit prefix for mapping IPv4 addresses to IPv6 addresses.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

ei-mapping-timeout

Syntax	mapping-timeout <i>seconds</i> ;
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	ei-mapping-timeout statement introduced in JUNOS Releases 12.3.
Description	Specify the duration for endpoint independent translations that use the specified NAT pool. This includes endpoint-independent mapping (EIM) and endpoint-independent filtering (EIF).
Options	seconds —Lifetime of endpoint independent mappings in seconds. Default: 300 Range: 120 through 864,000
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Network Address Translation Configuration Overview on page 23

eif-flow-limit

Syntax	eif-flow-limit <i>number-of-flows</i>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated secure-nat-mapping]
Release Information	Statement introduced in Junos OS Release 12.3
Description	Specify the maximum number of inbound flows allowed on EIF mapping to the configured value. This limit is per EIF mapping and is per given instance of time. For example, if eif-flow-limit is configured as n, then only n inbound connections are allowed at a given instance of time. The n+1 and subsequent connections arriving when n connections are alive are dropped . A new inbound connection is allowed only when one of the n connections times out or is closed. This limit is applied for all type of flows.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

from (Services NAT)

Syntax	<pre> from { application-sets set-name; applications [application-names]; destination-address (address any-unicast) <except>; destination-address-range low minimum-value high maximum-value <except>; source-address address (address any-unicast) <except>; source-address-range low minimum-value high maximum-value <except>; } </pre>
Hierarchy Level	[edit services nat rule rule-name term term-name]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify input conditions for the NAT term.
Options	<p>For information on match conditions, see the description of firewall filter match conditions in the <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> Network Address Translation Rules Overview

ipv6-multicast-interfaces

Syntax	ipv6-multicast-interfaces (all <i>interface-name</i>) { disable; }
Hierarchy Level	[edit services nat], [edit services software]
Release Information	Statement introduced in Junos OS Release 9.1.
Description	Enable multicast filters on Ethernet interfaces when IPv6 NAT is used for neighbor discovery.
Options	all —Enable filters on all interfaces. disable —Disable filters on the specified interfaces. interface-name —Enable filters on a specific interface only.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Configuring IPv6 Multicast Interfaces</i>

mapping-refresh

Syntax	mapping-refresh nbound outbound inbound-outbound);
Hierarchy Level	[edit services nat rule rule-name term term-name then translated secure-nat-mapping]
Release Information	Statement introduced in Junos OS Release 12.3
Description	Specify how the flow timer should be refreshed based on the mapping refresh configured for all types of fwnat flows. For TCP flows, if tcp-tickles is configured, then tickles are sent only on the flow matching the mapping-refresh direction. For inbound-outbound mapping, refresh tickles will be sent on both the flows (default behavior).
Options	inbound —Refresh the flow timer for inbound flows only. inbound-outbound —Refresh the flow timer for all flows. outbound —Refresh the flow timer for outbound flows only.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">•

mapping-timeout

Syntax	mapping-timeout <i>seconds</i> ;
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	mapping-timeout statement introduced in JUNOS Release 10.1.



NOTE: This configuration option has been replaced by [app-mapping-timeout](#). This option is currently retained only for backward compatibility.

Description	Specify the duration for mappings that use the specified NAT pool.
Options	<p>seconds—Lifetime of mappings in seconds.</p> <p>Default: 300</p> <p>Range: 120 through 864,000</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Source and Destination Addresses Network Address Translation Overview</i>

match-direction

Syntax	match-direction (input output);
Hierarchy Level	[edit services nat rule <i>rule-name</i>]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the direction in which the rule match is applied.
Options	<p>input—Apply the rule match on input.</p> <p>output—Apply the rule match on output.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Network Address Translation Rules Overview</i>

no-translation

Syntax	no-translation;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Specify that traffic is not to be translated.
Options	none
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

overload-pool

Syntax	overload-pool <i>overload-pool-name</i> ;
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Specify an address pool that can be used if the source pool becomes exhausted.
Options	<i>overload-pool-name</i> —Name of the overload pool.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

overload-prefix

Syntax	<code>overload-prefix <i>overload-prefix</i>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in Junos OS Release 7.6.
Description	Specify the prefix that can be used if the source pool becomes exhausted.
Options	<i>overload-prefix</i> —Prefix value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

pool

Syntax `pool nat-pool-name {
 address ip-prefix </prefix-length>;
 address-allocation round-robin;
 address-range low minimum-value high maximum-value;
 app- mapping-timeout app-mapping-timeout;
 ei-mapping-timeout ei-mapping-timeout;
 mapping-timeout mapping-timeout;
 pgcp {
 hint [hint-strings];
 ports-per-session ports;
 remotely-controlled;
 }
 port (automatic | range low minimum-value high maximum-value) {
 preserve-parity;
 preserve-range;
 secured-port-block-allocation {
 active-block-timeout timeout-seconds;
 block-size block-size;
 max-blocks-per-user max-blocks;
 }
 }
}`

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

Release Information Statement introduced before Junos OS Release 7.4.
pgcp statement added in Junos OS Release 8.4.
remotely-controlled and **ports-per-session** statements added in Junos OS Release 8.5.
hint statement added in Junos OS Release 9.0.
address-allocation statement added in Junos OS Release 11.2.

Description Specify the NAT name and properties.

Options *nat-pool-name*—Identifier for the NAT address pool.

The remaining statements are explained separately.

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

Related Documentation

- *Configuring Pools of Addresses and Ports for Network Address Translation Overview*

port

Syntax	<pre> port (automatic range low <i>minimum-value</i> high <i>maximum-value</i> random-allocation) { preserve-parity; preserve-range; deterministic-port-block-allocation <block-size <i>block-size</i>> <include-boundary-addresses>; secured-port-block-allocation { active-block-timeout <i>timeout-seconds</i>; block-size <i>block-size</i>; max-blocks-per-user <i>max-blocks</i>; } } </pre>
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i>]
Release Information	<p>port statement introduced before Junos OS Release 7.4.</p> <p>random-allocation statement introduced in Junos OS Release 9.3.</p> <p>secured-port-block-allocation statement introduced in Junos OS Release 11.2.</p> <p>deterministic-port-block-allocation statement introduced in Junos OS Release 12.1.</p>
Description	Specify the NAT pool port or range. You can configure an automatically assigned port or specify a range with minimum and maximum values.
Options	<p>automatic—Router-assigned port.</p> <p><i>minimum-value</i>—Lower boundary for the port range.</p> <p><i>maximum-value</i>—Upper boundary for the port range.</p> <p>preserve-parity—Allocate ports with same parity as the original port.</p> <p>preserve-range—Preserve privileged port range after translation.</p> <p>random-allocation—Allocate ports within a specified range randomly.</p> <p>Other options are described separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • <i>Configuring Source and Destination Addresses Network Address Translation Overview</i> • <i>Configuring Address Pools for Network Address Port Translation (NAPT) Overview</i>

port-forwarding

Syntax	<code>port-forwarding map-name { destined-port; translated-port; }</code>
Hierarchy Level	[edit services nat]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the mapping for port forwarding.
Options	<i>map-name</i> —Identifier for the port forwarding map.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Port Forwarding for Static Destination Address Translation on page 65• Configuring Port Forwarding Without Destination Address Translation on page 68

port-forwarding-mappings

Syntax	<code>port-forwarding-mappings map-name;</code>
Hierarchy Level	[edit services nat rule rule-name term term-name then]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the name for mapping port forwarding in a Network Address Translation configuration.
Options	<i>map-name</i> —Identifier for the port forwarding mapping.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Port Forwarding for Static Destination Address Translation on page 65• Configuring Port Forwarding Without Destination Address Translation on page 68

ports-per-session

Syntax	<code>ports-per-session <i>ports</i>;</code>
Hierarchy Level	[edit services nat pool <i>nat-pool-name</i> 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	<i>number-of-ports</i> —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.

rule

```
Syntax  rule rule-name {
        match-direction (input | output);
        term term-name {
            from {
                application-sets set-name;
                applications [ application-names ];
                destination-address (address | any-unicast) <except>;
                destination-address-range low minimum-value high maximum-value <except>;
                source-address (address | any-unicast) <except>;
                source-address-range low minimum-value high maximum-value <except>;
            }
            then {
                no-translation;
                translated {
                    address-pooling paired;
                    destination-pool nat-pool-name;
                    destination-prefix destination-prefix; destination-prefix;
                    dns-alg-pool dns-alg-pool;
                    dns-alg-prefix dns-alg-prefix;
                    filtering-type endpoint-independent;
                    mapping-type endpoint-independent;
                    overload-pool overload-pool;
                    overload-prefix overload-prefix;
                    source-pool nat-pool-name;
                    source-prefix source-prefix;
                    translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44
                                   | napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |
                                   twice-dynamic-nat-44 | twice-napt-44);
                }
            }
            syslog;
        }
    }
```

Hierarchy Level [edit [services nat](#)],
[edit [services nat rule-set rule-set-name](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule the router uses when applying this service.



NOTE: You are limited to a maximum of 200 terms for a NAT rule that is applied to an inline services (type si) interface. If you specify more than 200 terms, you will receive following error when you commit the configuration:

```
[edit]
' service-set service-set-name'
  NAT rule rule-name with more than 200 terms is disallowed for
  si-n/n/n.n
error: configuration check-out failed
```


Options *rule-name*—Identifier for the collection of terms that make up this rule.

The remaining statements are explained separately.

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

Related Documentation • *Network Address Translation Rules Overview*

rule-set

Syntax `rule-set rule-set-name {
 [rule rule-names];
}`

Hierarchy Level [edit *services* nat]

Release Information Statement introduced before Junos OS Release 7.4.

Description Specify the rule set the router uses when applying this service.

Options *rule-set-name*—Identifier for the collection of rules that constitute this rule set.

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

Related Documentation • *Network Address Translation Rules Overview*

secure-nat-mapping

Syntax `secure-nat-mapping {
 mapping-refresh (inbound | outbound | inbound-outbound);
 if-flow-limit number-of-flows'
}`

Hierarchy Level [edit *services* nat *rule* rule-name *term* term-name *then translated*]

Release Information Statement introduced in Junos OS Release 12.3

Options The statements are explained separately.

—

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

Related Documentation •

secured-port-block-allocation

Syntax	<pre>secured-port-block-allocation { active-block-timeout <i>timeout-seconds</i>; block-size <i>block-size</i>; max-blocks-per-address <i>max-blocks</i>; }</pre>
Hierarchy Level	[edit services nat pool <i>pool-name</i> port]
Release Information	Statement introduced in Junos OS Release 11.2.
Description	When you use block allocation, one or more blocks of ports in a NAT pool address range are available for assignment to a subscriber.
Options	<p><i>block-size</i>—Number of ports included in a block. Default: 128 Range: 1 to 32,000</p> <p><i>max-blocks</i>—Maximum number of blocks that can be allocated to a user address. Default: 8 Range: 1 to 512</p> <p><i>timeout-seconds</i>—Interval, in seconds, during which a block is active. After timeout, a new block is allocated, even if ports are available in the active block. Default: 0—The default timeout of the active block is 0 (infinite). In this case, the active block transitions to inactive only when it runs out of ports and a new block is allocated. Any inactive block without any ports in use will be freed to the NAT pool. Range: Any value greater than or equal to 120.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring Address Pools for Network Address Port Translation (NAPT) Overview

server (pcp)

Syntax `server server-name {
 ipv4-address ipv4-address;
 ipv6-address ipv6-address;
 software-concentrator software-concentrator-name;
 mapping-lifetime-min mapping-lifetime-min;
 mapping-lifetime-max mapping-lifetime-max;
 short-lifetime-error short-lifetime-error;
 long-lifetime-error long-lifetime-error;
 nat-options {
 pool pool-name ;
 }
 pcp-options {
 third-party
 prefer-failure
 }
 max-mapping-per-client max-mapping-per-client;
}`

Hierarchy Level [edit services pcp]

Release Information Statement introduced in Junos OS Release 13.2R1.

Description Configure PCP server options.

Options *ipv4-address*—IPv4 address of the PCP server.

ipv6-address—IPv6 address of the PCP server.

software-concentrator-name—Softwire concentrator name whose softwire-address is used in creating PCP mappings. The PCP server address must be the same as the softwire-concentrator address.

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

mapping-lifetime-max *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 2147483647 seconds

short-lifetime-error *short-lifetime-error*—Certain error opcodes mentioned in section 2 are classified as short lifetime errors. In case of these errors, the PCP server will use the value configured with this option to respond to the PCP client.

Default: 30 seconds

Range: 15 through 300 seconds

long-lifetime-error—Certain error opcodes mentioned in section 2 are classified as long lifetime errors. In case of these errors, the PCP server will use the value configured with this option to respond to the PCP client.

Default: 1800 seconds

Range: 900 through 18,000 seconds

max-mapping-per-client number-of-mappings—Maximum number of PCP mappings that the PCP client can request.

Default: 32

Range: 1 through 32

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	<ul style="list-style-type: none">• Configuring Port Control Protocol on page 71
------------------------------	--

services (NAT)

Syntax	services nat { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Define the service rules to be applied to traffic.
Options	nat —Identifies the NAT set of rules statements.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service-set (Services)

```

Syntax  service-set service-set-name {
        allow-multicast;
        extension-service service-name {
            provider-specific-rules-configuration;
        }
        (ids-rules rule-name | ids-rule-sets rule-set-name);
        interface-service {
            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;
            passive-mode-tunneling;
            trusted-ca [ ca-profile-names ];
            tunnel-mtu bytes;
        }
        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;
        }
        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;
        }
        (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;
            enable-asymmetric-traffic-processing;
            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;
      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;
  }
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 7.4.
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.

Description Define the service set.

Options ***service-set-name***—Name of the service set.

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

- *Service Set Properties*

source-address (NAT)

Syntax	<code>source-address (address any-unicast) <except>;</code>
Hierarchy Level	[edit services nat rule rule-name term term-name from]
Release Information	Statement introduced before Junos OS Release 7.4. any-unicast and except options introduced in Junos OS Release 7.6. address option enhanced to support IPv6 addresses in Junos OS Release 8.5.
Description	Specify the source address for rule matching.
Options	address —Source IPv4 or IPv6 address or prefix value. any-unicast —Any unicast packet. except —(Optional) Prevent the specified address or unicast packets from being translated.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Network Address Translation Rules Overview</i>

source-address-range

Syntax	<code>source-address-range low minimum-value high maximum-value <except>;</code>
Hierarchy Level	[edit services nat rule rule-name term term-name from]
Release Information	Statement introduced in Junos OS Release 7.6. minimum-value and maximum-value options enhanced to support IPv6 addresses in Junos OS Release 8.5.
Description	Specify the source address range for rule matching.
Options	minimum-value —Lower boundary for the IPv4 or IPv6 address range. maximum-value —Upper boundary for the IPv4 or IPv6 address range. except —(Optional) Prevent the specified address range from being translated.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • <i>Network Address Translation Rules Overview</i>

source-pool

Syntax	<code>source-pool nat-pool-name;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Specify the source address pool for translated traffic.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

source-prefix

Syntax	<code>source-prefix source-prefix;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	Statement introduced in Junos OS Release 7.6. <i>source-prefix</i> option enhanced to support IPv6 addresses in Junos OS Release 8.5.
Description	Specify the source prefix for translated traffic.
Options	<i>source-prefix</i> —IPv4 or IPv6 source prefix value.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

source-prefix-list

Syntax	<code>source-prefix-list <i>list-name</i> <except>;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> from]
Release Information	Statement introduced in Junos OS Release 8.2.
Description	Specify the source prefix list for rule matching. You configure the prefix list by including the prefix-list statement at the [edit policy-options] hierarchy level.
Options	<p>list-name—Destination prefix list.</p> <p>except—(Optional) Exclude the specified prefix list from rule matching.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Network Address Translation Rules Overview</i> <i>Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices</i>

syslog

Syntax	<code>syslog;</code>
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then]
Release Information	Statement introduced before Junos OS Release 7.4.
Description	Enable system logging. The system log information from the Multiservices PIC is passed to the kernel for logging in the /var/log directory.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> <i>Network Address Translation Rules Overview</i>

translated-port

Syntax	<code>translated-port <i>port id</i>;</code>
Hierarchy Level	[edit services nat port-forwarding <i>map-name</i>]
Release Information	Statement introduced in Junos OS Release 11.4.
Description	Specify the port to which all traffic will be translated.
Options	<i>port id</i> —The port number to which traffic will be translated.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• port-forwarding on page 140• destined-port on page 129

term

```
Syntax  term term-name {
        from {
            application-sets set-name;
            applications [ application-names ];
            destination-address (address | any-unicast) <except>;
            destination-address-range low minimum-value high maximum-value <except>;
            source-address (address | any-unicast) <except>;
            source-address-range low minimum-value high maximum-value <except>;
        }
        then {
            no-translation;
            translated {
                address-pooling paired;
                destination-pool nat-pool-name;
                destination-prefix destination-prefix;
                dns-alg-pool dns-alg-pool;
                dns-alg-prefix dns-alg-prefix;
                filtering-type endpoint-independent;
                mapping-type endpoint-independent;
                source-pool nat-pool-name;
                source-prefix source-prefix;
                translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44
                    | napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |
                    twice-dynamic-nat-44 | twice-napt-44);
            }
        }
        syslog;
    }
```

Hierarchy Level [edit [services](#) nat [rule](#) *rule-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the NAT term properties.

Options *term-name*—Identifier for the term.

The remaining statements are explained separately.

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

Related Documentation

- *Network Address Translation Rules Overview*

then

```
Syntax  then {
        no-translation;
        translated {
            address-pooling paired;
            destination-pool nat-pool-name;
            destination-prefix destination-prefix;
            dns-alg-pool dns-alg-pool;
            dns-alg-prefix dns-alg-prefix;
            filtering-type endpoint-independent;
            mapping-type endpoint-independent;
            source-pool nat-pool-name;
            source-prefix source-prefix;
            translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44
                | napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 |
                twice-dynamic-nat-44 | twice-napt-44);
        }
    }
    syslog;
}
```

Hierarchy Level [edit [services](#) nat [rule](#) *rule-name* [term](#) *term-name*]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define the NAT term actions.

Options The remaining statements are explained separately.

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

Related Documentation

- *Network Address Translation Rules Overview*

translated

Syntax translated {
 address-pooling paired;
 destination-pool nat-pool-name;
 destination-prefix destination-prefix;
 dns-alg-pool dns-alg-pool;
 dns-alg-prefix dns-alg-prefix;
 filtering-type endpoint-independent;
 mapping-type endpoint-independent;
 overload-pool overload-pool-name;
 overload-prefix;
 source-pool nat-pool-name;
 translation-type (basic-nat-pt | basic-nat44 | basic-nat66 | dnat-44 | dynamic-nat44 |
 napt-44 | napt-66 | napt-pt | stateful-nat64 | twice-basic-nat-44 | twice-dynamic-nat-44
 | twice-napt-44)
 }
 }

Hierarchy Level [edit [services](#) nat [rule](#) rule-name [term](#) term-name [then](#)]

Release Information Statement introduced before Junos OS Release 7.4.

Description Define properties for translated traffic.

Options The remaining statements are explained separately.

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

Related Documentation

- *Network Address Translation Rules Overview*

translation-type

Syntax	translation-type (basic-nat-pt basic-nat44 basic-nat66 nat-44 deterministic-napt44 dnat-44 dynamic-nat44 napt-44 napt-66 napt-pt stateful-nat64 twice-basic-nat-44 twice-dynamic-nat-44 twice-napt-44)
Hierarchy Level	[edit services nat rule <i>rule-name</i> term <i>term-name</i> then translated]
Release Information	<p>Statement introduced before Junos OS Release 7.4.</p> <p>The following options introduced in Junos OS Release 11.2, replacing all previous options:</p> <ul style="list-style-type: none">• basic-nat44• basic-nat66• basic-nat-pt• deterministic-napt44• dnat-44• dynamic-nat44• napt-44• napt-66• napt-pt• stateful-nat64 <p>twice-basic-nat-44 option introduced in Junos OS Release 11.4.</p> <p>twice-dynamic-nat-44 option introduced in Junos OS Release 11.4.</p> <p>twice-napt-44 option introduced in Junos OS Release 11.4.</p> <p>deterministic-napt44 option introduced in Junos OS Release 12.1.</p>
Description	Specify the NAT translation types.
Options	<ul style="list-style-type: none">• basic-nat44—Translate the source address statically (IPv4 to IPv4).• basic-nat66—Translate the source address statically (IPv6 to IPv6).• basic-nat-pt—Translate the addresses of IPv6 hosts as they originate sessions to the IPv4 hosts in the external domain. The basic-nat-pt option is always implemented with DNS ALG.• deterministic-napt44—Translate as napt-44, and use deterministic port block allocation for port translation.• dnat-44—Translate the destination address statically (IPv4 to IPv4).• dynamic-nat44—Translate only the source address by dynamically choosing the NAT address from the source address pool.

- **napt-44**—Translate the transport identifier of the IPv4 private network to a single IPv4 external address.
- **napt-66**—Translate the transport identifier of the IPv6 private network to a single IPv6 external address.
- **napt-pt**—Bind addresses in an IPv6 network with addresses in an IPv4 network and vice versa to provide transparent routing for the datagrams traversing between the address realms.
- **stateful-nat64**—Implement dynamic address and port translation for source IP addresses (IPv6-to-IPv4) and prefix removal translation for the destination IP addresses (IPv6-to-IPv4).
- **twice-basic-nat-44**—Translate the source and destination addresses statically (IPv4 to IPv4).
- **twice-dynamic-nat-44**—Translate the source address by dynamically choosing the NAT address from the source address pool. Translate the destination address statically.
- **twice-dynamic-napt-44**—Translate the transport identifier of the IPv4 private network to a single IPv4 external address. Translate the destination address statically.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• <i>Network Address Translation Rules Overview</i>

CHAPTER 9

Softwire Configuration Tasks

- [Configuring a DS-Lite Softwire Concentrator on page 159](#)
- [Configuring a 6rd Softwire Concentrator on page 160](#)
- [Configuring Stateful Firewall Rules for 6rd Softwire on page 161](#)
- [Configuring Softwire Rules on page 161](#)
- [Configuring Service Sets for Softwire on page 162](#)

Configuring a DS-Lite Softwire Concentrator

To configure a DS-Lite softwire concentrator:

1. Assign a name to the DS-Lite softwire concentrator.

```
[edit services softwire software-concentrator]  
user@host# edit ds-lite ds-lite-softwire-concentrator
```

2. Specify the address of the softwire tunnel.

```
[edit services softwire software-concentrator ds-lite ds-lite-softwire-concentrator]  
user@host# set software-address address
```

3. Specify the MTU for the softwire tunnel.

```
[edit services softwire software-concentrator ds-lite ds-lite-softwire-concentrator]  
user@host# set mtu-v6 mtu-v6
```



NOTE: This option sets the maximum transmission unit when encapsulating IPv4 packets into IPv6. If the final length is greater than the MTU, the IPv6 packet will be fragmented. This option is mandatory since it depends on other network parameters under administrator control.

4. To copy DSCP information from the IPv6 header into the decapsulated IPv4 header, include the **copy-dscp** statement.

```
[edit services softwire software-concentrator ds-lite ds-lite-softwire-concentrator]  
user@host# set copy-dscp
```

5. Specify the maximum number of flows for the softwire

```
[edit services softwire software-concentrator ds-lite ds-lite-softwire-concentrator]  
user@host# set flow-limit 1000
```

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
 - [Configuring Software Rules on page 161](#)
 - [Configuring IPv6 Multicast Interfaces](#)
 - [Configuring Service Sets for Software on page 162](#)
 - [Example: Basic DS-Lite Configuration on page 165](#)
 - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 176](#)

Configuring a 6rd Software Concentrator

To configure a 6rd software concentrator:

1. Assign a name to the 6rd software concentrator.

```
[edit services software software-concentrator]  
user@host# edit v6rd v6rd-software-concentrator
```

2. Specify the address of the software tunnel.

```
[edit services software software-concentrator v6rd v6rd-software-concentrator]  
user@host# set software-address address
```

3. Specify the MTU for the software tunnel.

```
[edit services software software-concentrator v6rd v6rd-software-concentrator]  
user@host# set mtu-v4 mtu-v4
```



TIP: In this release there is no support for fragmentation and reassembly, therefore the MTUs on the IPv6 and IPV4 network must be properly configured by the administrator.

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
 - [Configuring Software Rules on page 161](#)
 - [Configuring Stateful Firewall Rules for 6rd Software on page 161](#)
 - [Configuring Service Sets for Software on page 162](#)
 - [Example: Basic 6rd Configuration on page 171](#)
 - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 176](#)

Configuring Stateful Firewall Rules for 6rd Software

You must configure a stateful firewall rule for use with 6rd softwires. The stateful firewall service is used only to direct packets to the software, not for firewalling purposes. The 6rd software service itself must be stateless. To support stateless processing, you must include an **allow** term in both directions of the stateful firewall policy.

To include a stateful firewall rule for 6rd software processing:

1. Assign a name to the rule.

```
[edit services stateful-firewall]
user@host# edit rule rule-name
```

2. Specify the match direction.

```
[edit services stateful-firewall rule-name]
user@host# set match-direction input-output
```

3. Assign a name for the term.

```
[edit services stateful-firewall rule-name]
user@host# edit term term-name
```

4. Specify that all traffic in both directions should be accepted for the software process.

```
[edit services stateful-firewall rule-name term term-name]
user@host# set then accept
```

Related Documentation

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
- [Configuring a 6rd Software Concentrator on page 160](#)
- [Configuring Software Rules on page 161](#)
- [Configuring Service Sets for Software on page 162](#)
- [Example: Basic 6rd Configuration on page 171](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 176](#)

Configuring Software Rules

You configure software rules to instruct the router how to direct traffic to the addresses specified for 6rd or DS-Lite 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 6rd or DS-Lite software concentrator.

You can create a software rule consisting of one or more terms and associate a particular 6rd or DS-Lite software concentrator with each term. You can include the software rule in service sets along with other services rules.

To configure a software rule:

1. Assign a name to the rule.

```
[edit services software ]  
user@host# edit rule rule-name
```

2. Specify the match direction.

```
[edit services software rule rule-name]  
user@host# set match-direction (input | output)
```

3. Assign a name for the first term.

```
[edit services software rule rule-name]  
user@host# edit term term-name
```

4. Associate a 6rd or DS-Lite software concentrator with this term.

```
[edit services software rule rule-name term term-name]  
user@host# set then ds-lite name
```

Or

```
user@host# set then v6rd v6rd-software-concentrator
```

5. Repeat Steps 3 and 4 for as many additional terms as needed.

Related Documentation

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
- [Configuring a 6rd Software Concentrator on page 160](#)
- [Configuring a DS-Lite Software Concentrator on page 159](#)
- [Configuring IPv6 Multicast Interfaces](#)
- [Configuring Service Sets for Software on page 162](#)

Configuring Service Sets for Software

You must include software rules or a software rule set in a service set to enable software processing. You must include a stateful firewall rule for DS-Lite.

To configure service sets for software:

1. Include a software rule or rule set in the service set.

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

2. When using a 6rd software, include a stateful-firewall rule.

```
[edit services service-set service-set-name]  
user@host# set stateful-firewall-rules software-rule-name
```

3. You can include a NAT rule for flows originated by DS-Lite softwares.

**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 concentrator, if you configure stateful firewall rules without configuring NAT rules, using an interface service set causes the ICMP echo reply messages to be not 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.

For further information, see “*Configuring Service Rules*.”

Related Documentation

- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
- [Configuring Software Rules on page 161](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 176](#)

CHAPTER 10

Softwire Configuration Examples

- [Example: Basic DS-Lite Configuration on page 165](#)
- [Example: Basic 6rd Configuration on page 171](#)
- [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 176](#)

Example: Basic DS-Lite Configuration

- [Requirements on page 165](#)
- [Configuration Overview and Topology on page 165](#)
- [Configuration on page 166](#)

Requirements

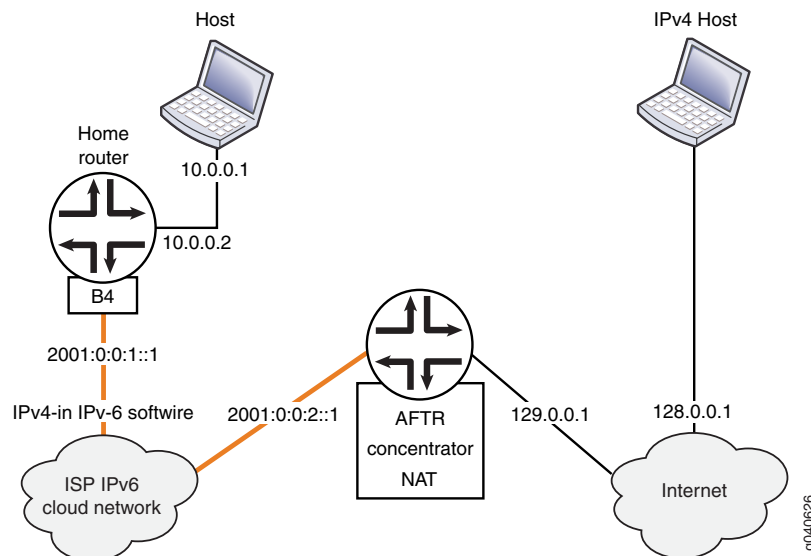
The following hardware components can perform DS-Lite:

- M Series Multiservice Edge routers with Multiservices PICs
- T Series Core routers with Multiservices PICs
- MX Series 3D Universal Edge routers with Multiservices DPCs

Configuration Overview and Topology

This example describes how to configure an MX Series router with an MS-DPC as an AFTR to facilitate the flow shown in [Figure 11 on page 166](#).

Figure 11: DS-Lite Topology



In this example, the DS-Lite softwire concentrator, or AFTR, is an MX Series router with two Gigabit interfaces and a Services DPC. The interface facing the B4 element is ge-3/1/5 and the interface facing the Internet is ge-3/1/0.

Configuration

- [Chassis Configuration on page 166](#)
- [Interfaces Configuration on page 166](#)
- [Network Address and Port Translation Configuration on page 168](#)
- [Softwire Configuration on page 169](#)
- [Service Set Configuration on page 170](#)

Chassis Configuration

Step-by-Step Procedure

To configure the service PIC (FPC 0 Slot 0) with the Layer 3 service package:

1. Enter the **edit chassis** hierarchy level.

```
user@host# edit chassis
```
2. Configure the Layer 3 service package.

```
[edit chassis]
user@host# set fpc 0 pic 0 adaptive-services service-package layer-3
```

Interfaces Configuration

Step-by-Step Procedure

To configure interfaces facing the B4 (softwire initiator) and facing the Internet:

1. Go the **[edit interfaces]** edit hierarchy level for ge-3/1/0, which faces the Internet.

```
host# edit interfaces ge-3/1/0
```


2. Define the interface.

```
[edit interfaces ge-3/1/0]
user@host# set description AFTR-Internet
user@host# set unit 0 family inet address 128.0.0.2/24
```

3. Go to the **[edit interfaces]** hierarchy level for ge-3/1/5, which faces the B4.

```
user@host# up 1
[edit]
user@host# edit interfaces ge-3/1/5
```

4. Define the interface.

```
[edit interfaces ge-3/1/5]
user@host# set description AFTR-B4
user@host# set unit 0 family inet
user@host# edit unit 0 family inet6
[edit unit 0 family inet6]
user@host# set service input service-set sset
user@host# set service output service-set sset
user@host# set address 2001:0:0:2::1/48
```

5. Go to the **[edit interfaces]** hierarchy level for sp-0/0/0, used to host the DS-Lite AFTR.

```
[edit]
user@host# edit interfaces sp-0/0/0
```

6. Define the interface.

```
[edit interfaces sp-0/0/0]
user@host# set description AFTR-B4
user@host# set unit 0 family inet
user@host# edit unit 0 family inet6
```

Results

```
user@host# show interfaces ge-3/1/0
description AFTR-Internet;
unit 0 {
    family inet {
        address 128.0.0.2/24;
    }
}

user@host# show interfaces ge-3/1/5
description AFTR-B4;
unit 0 {
    family inet;
    family inet6 {
        service {
            input {
                service-set sset;
            }
            output {
                service-set sset;
            }
        }
        address 2001:0:0:2::1/48;
    }
}

user@host# show interfaces sp-o/o/o
unit 0 {
    family inet;
    family inet6;
}
```

Network Address and Port Translation Configuration

Step-by-Step Procedure

To configure NAPT:

1. Go to the **[edit services nat]** hierarchy level.

```
user@host# edit services nat
[edit services nat]
```
2. Define a NAT pool p1.

```
user@host# set pool p1 address 129.0.0.1/32 port automatic
```
3. Define a NAT rule, beginning with the match direction.

```
[edit services nat]
user@host# set rule r1 match-direction input
```
4. Define a **term** for the rule, beginning with a **from** clause.

```
[edit services nat]
user@host# set rule r1 term t1 from source-address 10.0.0.0/16
```
5. Define the desired translation in a **then** clause. In this case, use dynamic source translation.

```
[edit services nat]
user@host# set rule r1 term t1 then translated source-pool p1 translation-type
napt-44
```
6. (Optional) Configure logging of translation information for the rule.

```
[edit services nat]
user@host# set rule r1 term t1 then syslog
```

Results

```
user@host# show services nat
pool p1 {
  address 129.0.0.1/32;
  port {
    automatic;
  }
}
rule r1 {
  match-direction input;
  term t1 {
    from {
      source-address {
        10.0.0.0/16;
      }
    }
    then {
      translated {
        source-pool p1;
        translation-type {
          napt-44;
        }
      }
      syslog;
    }
  }
}
```

Softwire Configuration

Step-by-Step Procedure To configure the DS-Lite softwire concentrator and associated rules:

1. Go to the **[edit services softwire]** hierarchy level.

```
user@host# edit services softwire
```
2. Define the DS-Lite softwire concentrator.

```
[edit services softwire]
user@host# set softwire-concentrator ds-lite ds-1 softwire-address 1001::1 mtu-v6 1460
```
3. Define the softwire rule.

```
[edit services softwire]
user@host# set rule r1 match-direction input term t1 then ds-lite ds1.
```

Results `user@host# show services software`

```
software-concentrator {
  ds-lite ds1 {
    software-address 1001::1;
    mtu-v6 1460;
  }
}
rule r1 {
  match-direction input;
  term t1 {
    then {
      ds-lite ds1;
    }
  }
}
```

Service Set Configuration

Step-by-Step Procedure Configure a service set that includes software and NAT rules and specifies either interface-service or next-hop service. This example uses a next-hop service.

1. Go to the `[edit services service-set]` hierarchy level, naming the service set.

```
user@host# edit services service-set sset
```

2. Define the NAT rule to be used for IPv4-to-IPv4 translation.

```
[edit services service-set sset]
user@host# set nat-rules r1
```

3. Define the software rule to define the software tunnel.

```
[edit services service-set sset]
user@host# set software-rules r1
```

4. Define the interface service,

```
[edit services service-set sset]
user@host# set interface-service service-interface sp-0/0/0.0
```



TIP: In order to avoid or minimize IPv6 fragmentation, you can configure a TCP maximum segment size (MSS) for your service set.

5. (Optional) Define a TCP MSS.

```
[edit services service-set sset]
user@host# set tcp-mss 1024
```

Results `user@host# show services service-set`

```

syslog {
  host local {
    services any;
  }
}
softwire-rules r1;
nat-rules r1;
interface-service {
  service-interface sp-0/0/0;
}
}

```

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
 - [Configuring a DS-Lite Softwire Concentrator on page 159](#)
 - [Configuring Softwire Rules on page 161](#)
 - [Configuring Service Sets for Softwire on page 162](#)
 - [Example: Basic 6rd Configuration on page 171](#)
 - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 176](#)

Example: Basic 6rd Configuration

- [Requirements on page 171](#)
- [Overview on page 171](#)
- [Configuration on page 171](#)

Requirements

This example describes how a 6rd concentrator can be configured for a 6rd domain, D1, to provide IPv6 Internet connectivity.

The following hardware components can perform 6rd:

- M Series Multiservice Edge routers with Multiservices PICs
- T Series Core routers with Multiservices PICs
- MX Series 3D Universal Edge routers with Multiservices DPCs

Overview

This configuration example describes how to configure a basic 6rd tunneling solution.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```

set interfaces ge-1/2/0 unit 0 family inet service input service-set v6rd-dom1-service-set
set interfaces ge-1/2/0 unit 0 family inet service output service-set v6rd-dom1-service-set
set interfaces ge-1/2/0 unit 0 family inet address 10.10.10.1/24
set interfaces ge-1/2/0 unit 0 family inet6 service input service-set v6rd-dom1-service-set
set interfaces ge-1/2/0 unit 0 family inet6 service output service-set v6rd-dom1-service-set
set interfaces ge-1/2/2 unit 0 family inet6 address 3abc::1/16
set interfaces sp-0/2/0 unit 0 family inet
set interfaces sp-0/2/0 unit 0 family inet6
set services software software-concentrator v6rd v6rd-dom1 software-address 30.30.30.1
set services software software-concentrator v6rd v6rd-dom1 ipv4-prefix 10.10.10.0/24
set services software software-concentrator v6rd v6rd-dom1 v6rd-prefix 3040::0/16
set services software rule v6rd-dom1 match-direction input
set services software rule v6rd-dom1 term t1 then v6rd v6rd-dom1
set services service-set v6rd-dom1-service-set software-rules v6rd-dom1
set services service-set v6rd-dom1-service-set stateful-firewall-rules r1
set services service-set v6rd-dom1-service-set interface-service service-interface sp-0/2/0
set services stateful-firewall rule r1 match-direction input-output
set services stateful-firewall rule r1 term t1 then accept

```

Chassis Configuration

Step-by-Step Procedure

To configure the chassis:

1. Define the ingress interface.

```

user@host# edit interfaces ge-1/2/0

```
2. Configure the ingress interface logical unit and input/output service options.

```

[edit interfaces ge-1/2/0]
user@host# set unit 0 family inet service input service-set v6rd-dom1-service-set
user@host# set unit 0 family inet service output service-set v6rd-dom1-service-set
user@host# set unit 0 family inet6 service input service-set v6rd-dom1-service-set
user@host# set unit 0 family inet6 service output service-set v6rd-dom1-service-set

```
3. Configure the address of the ingress interface.

```

[edit interfaces ge-1/2/0]
user@host# set unit 0 family inet address 10.10.10.1/24

```
4. Define the egress interface.

```

user@host# up
[edit interfaces]
user@host# edit ge-1/2/2

```
5. Define the logical unit and address for the egress interface.

```

[edit interfaces ge-1/2/2]
user@host# set unit 0 family inet6 address 3ABC::1/16

```
6. Define the services PIC.

```

[edit interfaces ge-1/2/2]
user@host# up
[edit interfaces]
user@host# edit sp-0/2/0

```
7. Configure the logical unit for the services PIC.

```
[edit interfaces sp-0/2/0]
user@host# up
[edit interfaces]
user@host# set unit 0 family inet
user@host# set unit 0 family inet6
```

Results

```
[edit interfaces]
user@router# show
sp-0/2/0 {
  unit 0 {
    family inet;
    family inet6;
  }
}
ge-1/2/0 {
  unit 0 {
    family inet {
      service {
        input {
          service-set v6rd-dom1-service-set;
        }
        output {
          service-set v6rd-dom1-service-set;
        }
      }
      address 10.10.10.1/24;
    }
    family inet6 {
      service {
        input {
          service-set v6rd-dom1-service-set;
        }
        output {
          service-set v6rd-dom1-service-set;
        }
      }
    }
  }
}
ge-1/2/2 {
  unit 0 {
    family inet6 {
      address 3abc::1/16;
    }
  }
}
```

Software Concentrator, Software Rule, and Stateful Firewall Rule Configuration

Step-by-Step Procedure To configure the software concentrator, software rule, and stateful firewall rule:

1. Define the 6rd software concentrator.

```
user@host# top
user@host# edit services software software-concentrator v6rd v6rd-dom1
```

2. Configure the software concentrator properties. Here, software address 30.30.30.1 is the software concentrator IPv4 address, 10.10.10.0/24 is the IPv4 prefix of the CE WAN side, and 3040::0/16 is the IPv6 prefix of the 6rd domain D1.

```
[edit services software software-concentrator v6rd v6rd-dom1]
user@host# set software-address 30.30.30.1
user@host# set ipv4-prefix 10.10.10.0/24
user@host# set v6rd-prefix 3040::0/16
user@host# set mtu-v4 9192
```

3. Define the software rule.

```
[edit services software software-concentrator v6rd v6rd-dom1]
user@host# up 2
[edit services software]
user@host# edit rule v6rd-dom1
[edit services software rule v6rd-dom1]
user@host# set match-direction input
[edit services software rule v6rd-dom1]
user@host# set term t1 then v6rd v6rd-dom1
```

4. Define a stateful firewall rule and properties. You must configure a stateful firewall rule that accepts all traffic in both the input and output direction in order for 6rd to work; however, this is not enforced through the CLI. This is because in IPv6, gratuitous IPv6 packets are expected (due to Anycast) and should not be dropped. The service PIC can handle reverse traffic without seeing all forward traffic. This can also happen with service PIC switchover in the middle of a session. By default, the stateful firewall on the service PIC will drop all traffic unless a rule is configured explicitly to allow it.

```
[edit services software software-concentrator v6rd v6rd-dom1]
user@host# up 3
[edit services]
user@host# edit services stateful-firewall
[edit services stateful-firewall]
user@host# edit rule r1
[edit services stateful-firewall rule r1]
user@host# set match-direction input-output
user@host# set term t1 then accept
```


Results [edit services software]
 user@router# show
 software-concentrator {
 v6rd v6rd-dom1 {
 software-address 30.30.30.1;
 ipv4-prefix 10.10.10.0/24;
 v6rd-prefix 3040::0/16;
 mtu-v4 9192;
 }
 }
 rule v6rd-dom1-r1 {
 match-direction input;
 term t1 {
 then {
 v6rd v6rd-dom1;
 }
 }
 }

Service Set Configuration

Step-by-Step Procedure To configure the service set:

1. Define the service set for 6rd processing.

```
user@host# top
user@host# edit services service-set v6rd-dom1-service-set
```
2. Define the software and stateful firewall rules for the service set.

```
[edit services service-set v6rd-dom1-service-set]
user@host# set software-rules v6rd-dom1
user@host# set stateful-firewall-rules r1
```
3. Define the interface-service for the service set.

```
[edit services service-set v6rd-dom1-service-set]
user@host# set interface-service service-interface sp-0/2/0
```

Results [edit service-set v6rd-dom1-service-set]
 user@host# show
 software-rules v6rd-dom1-r1
 interface-service {
 service-interface sp-0/2/0;
 }

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
 - [Configuring a 6rd Software Concentrator on page 160](#)
 - [Configuring Software Rules on page 161](#)
 - [Configuring Stateful Firewall Rules for 6rd Software on page 161](#)
 - [Configuring Service Sets for Software on page 162](#)
 - [Example: Basic DS-Lite Configuration on page 165](#)
 - [Example: Configuring DS-Lite and 6rd in the Same Service Set on page 176](#)

Example: Configuring DS-Lite and 6rd in the Same Service Set

- [Requirements on page 176](#)
- [Overview on page 176](#)
- [Configuration on page 176](#)

Requirements

The following hardware components can perform DS-Lite:

- M Series Multiservice Edge routers with Multiservices PICs
- T Series Core routers with Multiservices PICs
- MX Series 3D Universal Edge routers with Multiservices DPCs

Overview

This example describes a software solution that includes DS-Lite and 6rd in the same service set.

Configuration

Chassis Configuration

Step-by-Step Procedure

To configure the chassis:

1. Configure the ingress interface.

```
user@host# edit interfaces ge-1/2/0
[edit interfaces ge-1/2/0]
user@host# set unit 0 family inet service input service-set v6rd-dslite-service-set
user@host# set unit 0 family inet service output service-set v6rd-dslite-service-set
user@host# set unit 0 family inet address address 10.10.10.1/24
user@host# set unit 0 family inet6 service input service-set v6rd-dslite-service-set
user@host# set unit 0 family inet6 service output service-set v6rd-dslite-service-set
user@host# set unit 0 family inet6 address address address 2001::1/16
```

Here the service set is applied on the inet (IPv4) and inet6 (IPv6) families of subunit 0. Both DS-Lite IPv6 traffic and 6rd IPv4 traffic hits the service filter and is sent to the services PIC.

2. Configure the egress interface (IPv6 Internet). The IPv4 server that the DS-Lite clients are trying to reach is at 200.200.200.2/24, and the IPv6 server is at 3ABC::2/16.

```
user@host# edit interfaces ge-1/2/2
[edit interfaces ge-1/2/2]
user@host# set unit 0 family inet address 200.200.200.1/24
user@host# set unit 0 family inet6 address 3ABC::1/16
```

3. Configure the services PIC.

```
user@host# edit interfaces sp-3/0/0
[edit interfaces sp-3/0/0]
```

```

user@host# set unit 0 family inet
user@host# set unit 0 family inet6

```

Results [edit interfaces]

```

user@host# show
ge-1/2/0 {
  unit 0 {
    family inet {
      service {
        input {
          service-set v6rd-dslite-service-set;
        }
        output {
          service-set v6rd-dslite-service-set;
        }
      }
      address 10.10.10.1/24;
    }
    family inet6 {
      service {
        input {
          service-set v6rd-dslite-service-set;
        }
        output {
          service-set v6rd-dslite-service-set;
        }
      }
      address 2001::1/16;
    }
  }
}
ge-1/2/2 {
  unit 0 {
    family inet {
      address 200.200.200.1/24;
    }
    family inet6 {
      address 3ABC::1/16;
    }
  }
}
sp-3/0/0 {
  unit 0 {
    family inet;
    family inet6;
  }
}

```

Software Concentrator, Software Rule, Stateful Firewall Rule Configuration

Step-by-Step Procedure To configure the software concentrator, software rule, and stateful firewall rule:

1. Configure the DS-Lite and 6rd software concentrators.

```

user@host# edit services software software-concentrator ds-lite ds1
[edit services software software-concentrator ds-lite ds1]
user@host# set software-address 1001::1
user@host# mtu-v6 9192
usert@host# up 1

```

```

user@host# edit v6rd v6rd-dom1
[edit services software-concentrator v6rd v6rd-dom1]
user@host# set software-address 30.30.30.1
user@host# set ipv4-prefix 10.10.10.0/24
user@host# set v6rd-prefix 3040::0/16
user@host# set mtu-v4 9192

```

2. Configure the software rules.

```

user@host# edit services software rule v6rd-r1]
[edit services software rule v6rd-r1]
user@host# set match-direction input
user@host# set term t1 then v6rd v6rd-dom1
user@host# up 1
user@host# edit services software]
[edit services software]
user@host# edit rule dslite-r1
[edit services software rule dslite-r1]
user@host# set term dslite-t1 then ds-lite ds1

```

The following routes are added by the services PIC daemon on the Routing Engine:

```

user@router# run show route 30.30.30.1
inet.0: 43 destinations, 46 routes (42 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

30.30.30.1/32      *[Static/786432] 00:24:11
                  Service to v6rd-dslite-service-set

[edit]
user@router# run show route 3040::0/16

inet6.0: 23 destinations, 33 routes (23 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3040::/16         *[Static/786432] 00:24:39
                  Service to v6rd-dslite-service-set

user@router# run show route 1001::1
inet6.0: 33 destinations, 43 routes (33 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1001::1/128       *[Static/1] 1w2d 22:05:41
                  Service to v6rd-dslite-service-set

```

3. Configure a stateful firewall rule.

```

user@host# edit services stateful-firewall rule r1
[edit services stateful-firewall rule r1]
user@host# set match-direction input-output
user@host# set term t1 then accept

[edit services stateful-firewall]
rule r1 {
    match-direction input-output;
    term t1 {
        then {
            accept;
        }
    }
}

```

```

Results [edit services softwire]
user@host# show
software-concentrator {
    ds-lite ds1 {
        software-address 1001::1;
        mtu-v6 9192;
    }
    v6rd v6rd-dom1 {
        software-address 30.30.30.1;
        ipv4-prefix 10.10.10.0/24;
        v6rd-prefix 3040::0/16;
        mtu-v4 9192;
    }
}
rule v6rd-r1 {
    match-direction input;
    term t1 {
        then {
            v6rd v6rd-dom1;
        }
    }
}
rule dslite-r1 {
    match-direction input;
    term dslite-t1 {
        then {
            ds-lite ds1;
        }
    }
}

[edit services stateful-firewall]
user@host# show
rule r1 {
    match-direction input-output;
    term t1 {
        then {
            accept;
        }
    }
}

```

NAT Configuration for DS-Lite

Step-by-Step Procedure To configure NAT for DS-Lite:

1. Configure a NAT pool for DS-Lite.


```

user@host# edit services nat pool dslite-pool
[edit services nat pool dslite-pool]
user@host# set address-range low 33.33.33.1 high 33.33.33.32
user@host# set port automatic
      
```
2. Configure a NAT rule.


```

user@host# up 1
[edit services nat rule dslite-nat-r1]
user@host# set match-direction input
      
```

```
user@host# set term dslite-nat-t1 from source-address 20.20.0.0/16 then translated  
translation-type napt-44
```

Results

```
[edit services nat]
user@host# show
pool dslite-pool {
    address-range low 33.33.33.1 high 33.33.33.32;
    port {
        automatic;
    }
}
rule dslite-nat-r1 {
    match-direction input;
    term dslite-nat-t1 {
        from {
            source-address {
                20.20.0.0/16;
            }
        }
        then {
            translated {
                source-pool dslite-pool;
                translation-type {
                    source dynamic;
                }
            }
        }
    }
}
```

Because of this NAT rule, the following NAT routes are installed for the reverse DS-Lite traffic:

```
user@router# run show route 33.33.33.0/24
inet.0: 48 destinations, 52 routes (47 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

33.33.33.1/32      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.2/31      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.4/30      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.8/29      *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.16/28     *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
33.33.33.32/32     *[Static/1] 1w2d 23:08:38
                  Service to v6rd-dslite-service-set
```

The NAT rule triggers address translation for the traffic coming from 20.20.0.0/16 to public address range 33.33.33.1 to 33.33.33.32.

Service Set Configuration

Step-by-Step Procedure

This service set has a stateful firewall rule and 6rd rule for 6rd service. The service set also includes a software rule for DS-Lite and a NAT rule to perform address translation for all DS-Lite traffic. The NAT rule performs NAPT translation in the forward direction on the source address and port of the DS-Lite traffic.

To configure the service set:

1. Define the service set.

```
user@host# edit services service-set v6rd-dslite-service-set
```
2. Configure the service set rules.

```
[edit services service-set v6rd-dslite-service-set]  
user@host# set software-rules dslite-r1  
user@host# set stateful-firewall-rules r1  
user@host# set nat-rules dslite-nat-r1
```
3. Configure the service set interface-service.

```
[edit services service-set v6rd-dslite-service-set]  
user@host# set interface-service service-interface sp-3/0/0
```

Results

```
[edit services service-set]  
user@host# show  
v6rd-dslite-service-set {  
  software-rules v6rd-r1;  
  software-rules dslite-r1;  
  stateful-firewall-rules r1;  
  nat-rules dslite-nat-r1;  
  interface-service {  
    service-interface sp-3/0/0;  
  }  
}
```

- Related Documentation**
- [Tunneling Services for IPv4-to-IPv6 Transition Overview on page 15](#)
 - [Configuring Service Sets for Software on page 162](#)
 - [Example: Basic DS-Lite Configuration on page 165](#)
 - [Example: Basic 6rd Configuration on page 171](#)

6to4 Configuration

- [Configuring a 6to4 Provider-Managed Tunnel on page 183](#)

Configuring a 6to4 Provider-Managed Tunnel

When configuring a 6to4 provider-managed tunnel (PMT), replace the Anycast destination with the address of a managed relay in the provider network.

To configure a 6to4 PMT:

1. Configure the ingress interface for 6to4 traffic. Include the name of the service set that identifies the rules for input and output service on this interface.

```
[edit interfaces ge-0/2/1]
user@host# set unit logical-unit-number family family service input service-set-name
user@host# set unit logical-unit-number family family service output service-set-name
user@host# set unit logical-unit-number family family address address
```

For example:

```
[edit interfaces ge-0/2/1]
user@host# set unit 0 family inet service input service-set v6to4-pmt
user@host# set unit 0 family inet service output service-set v6to4-pmt
user@host# set unit 0 family inet address 130.130.130.1/24
```

2. Configure the egress interface.

```
[edit interfaces ge-0/2/2]
user@host# set unit logical-unit-number family family address address
```

For example:

```
[edit interfaces ge-0/2/2]
user@host# set unit 0 family inet6 address 4ABC::1/16
```

3. Configure the service interface that contains the rules for processing incoming traffic. Include a syslog option and associate a logical unit.

```
[edit interfaces sp-2/0/0]
user@host# edit services-options syslog host host-name services any
user@host# edit unit logical-unit-number family family
user@host# edit unit 0 family family
```

For example:

```
[edit interfaces sp-2/0/0]
```

```
user@host# set services-options syslog host local services any
user@host# set unit 0 family inet
user@host# set unit 0 family inet6
```

4. Configure the softwire concentrator and softwire rule for 6to4. In the Junos OS, 6to4 PMT configuration uses the same options as 6rd.

```
[edit services softwire softwire-concentrator v6rd v6to4]
user@host# set softwire-address softwire-address
user@host# set ipv4-prefix ipv4-prefix
user@host# set v6rd-prefix v6rd-prefix
user@host# set mtu-v4 mtu-v4
```

For example:

```
[edit services softwire softwire-concentrator v6rd v6to4]
user@host# set softwire-address 192.88.99.1
user@host# set ipv4-prefix 130.130.130.2/32
user@host# set v6rd-prefix 2002::0/16
user@host# set mtu-v4 9192
```

5. Define the softwire rule that will process traffic on the ingress interface.

```
[edit services softwire rule v6to4-r1]
user@host# set match-direction input
user@host# set term term-name then v6rd softwire-concentrator
```

For example:

```
[edit services softwire rule v6to4-r1]
user@host# set match-direction input
user@host# set term t1 then v6rd v6to4
```

6. Define a stateful firewall rule that will accept all incoming traffic on the ingress interface.

```
[edit services stateful-firewall rule sfw-r1]
user@host# set match-direction direction
user@host# set term term-name then accept
user@host# set term term-name then syslog
```

For example:

```
[edit services stateful-firewall rule sfw-r1]
user@host# set match-direction input-output
user@host# set term t1 then accept
user@host# set term t1 then syslog
```

7. Define the NAT pool to be used for IPv6 NAT translation. This pool supports translation of the Anycast 6to4 relay addresses to addresses at the provider-managed relay.

```
[edit services nat pool v6to4-pmt]
user@host# set address address
user@host# port automatic
```

For example:

```
[edit services nat pool v6to4-pmt]
user@host# set address 3ABC::1/128
user@host# set port automatic
```

8. Define the NAT rule for translation.

```
[edit services nat rule rule-name]
user@host# set match-direction input
user@host# set term term-name then translated source-pool pool-name
user@host# set term t1 then translated translation-type translation-type
```

For example:

```
[edit services nat rule v6to4-pmt-r1]
user@host# set match-direction input
user@host# set term t1 then translated source-pool v6to4-pmt
user@host# set term t1 then translated translation-type napt-66
```

9. Define the service set that specifies the software rule and NAT rule.

```
[edit services service-set v6to4-pmt]
user@host# set software-rules rule-name
user@host# set stateful-firewall-rules rule-name
user@host# set nat-rules rule-name
user@host# set interface-service service-interface interface-name
```

For example:

```
[edit services service-set v6to4-pmt]
user@host# set software-rules v6to4-r1
user@host# set stateful-firewall-rules sfw-r1
user@host# set nat-rules v6to4-pmt-r1
user@host# set interface-service service-interface sp-2/0/0
```


CHAPTER 12

Softwire Configuration Statements

- [ds-lite on page 188](#)
- [rule \(Softwire\) on page 189](#)
- [rule-set \(Softwire\) on page 189](#)
- [softwire-concentrator on page 190](#)
- [softwire-options on page 191](#)
- [softwire-rules on page 191](#)
- [v6rd on page 192](#)

ds-lite

Syntax	<pre>ds-lite ds-lite-software-concentrator { auto-update-mtu; copy-dscp; flow-limit flow-limit session-limit-per-prefix session-limit-per-prefix; mtu-v6 mtu-v6; software-address software-address; }</pre>
Hierarchy Level	[edit services software software-concentrator]
Release Information	<p>Statement introduced in Junos OS Release 10.4.</p> <p>auto-update-mtu option introduced in Junos OS Release 10.4.</p> <p>copy-dscp option introduced in Junos OS Release 11.2.</p> <p>mtu-v6 option introduced in Junos OS Release 10.4.</p> <p>software-address option introduced in Junos OS Release 10.4.</p>
Description	Configure settings for a DS-Lite concentrator used to process IPv4 packets encapsulated in IPv6.
Options	<p>ds-lite-software-concentrator—Name applied to a DS-Lite software concentrator.</p> <p>auto-update-mtu—This option is not currently supported.</p> <p>copy-dscp—Copy DSCP information to IPv4 headers during decapsulation.</p> <p>flow-limit—Maximum number of IPv4 flows per software.</p> <p>Range: 0 through 16384 flows</p> <p>mtu-v6—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.</p> <p>Range: 0 through 9192 bytes</p> <p>session-limit-per-prefix—Maximum number of sessions per B4 subnet prefix. (0 through 16384).</p> <p>Range: 0 through 16384 sessions</p> <p>software-address—Address of the DS-Lite software concentrator.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring a DS-Lite Software Concentrator on page 159

rule (Software)

Syntax	<pre>rule <i>rule-name</i> { match-direction (input output); term <i>term-name</i> { then { (ds-lite <i>ds-lite-software-concentrator</i> v6rd <i>v6rd-software-concentrator</i>); } } }</pre>
Hierarchy Level	[edit services software], [edit services software rule-set <i>rule-set-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Configure a rule to apply a software concentrator for a flow.
Options	<p><i>rule-name</i>—Identifier for the collection of terms that constitute this rule.</p> <p>input—Apply the rule match on the input side of the interface.</p> <p>output—Apply the rule match on the output side of the interface.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Software Rules on page 161

rule-set (Software)

Syntax	<pre>rule-set <i>rule-set-name</i> { rule <i>rule-name</i>; }</pre>
Hierarchy Level	[edit services software]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the rule set the router uses when applying this service.
Options	<p><i>rule-set-name</i>—Identifier for the collection of rules that constitute this rule set.</p>
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none"> • Configuring Software Rules on page 161

software-concentrator

Syntax	<pre>software-concentrator { ds-lite ds-lite-software-concentrator { auto-update-mtu; flow-limit <i>flow-limit</i> session-limit-per-prefix <i>session-limit-per-prefix</i>; mtu-v6 <i>mtu-v6</i>; software-address <i>address</i>; } v6rd v6rd-software-concentrator { ipv4-prefix <i>ipv4-prefix</i>; v6rd-prefix <i>ipv6-prefix</i>; mtu-v4 <i>mtu-v4</i>; } }</pre>
Hierarchy Level	[edit services software]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Configure settings for a software concentrator.
Options	The remaining statements are explained separately.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Documentation	<ul style="list-style-type: none">• Configuring a DS-Lite Software Concentrator on page 159• Configuring a 6rd Software Concentrator on page 160

software-options

Syntax	software-options { dslite-ipv6-prefix-length <i>dslite-ipv6-prefix-length</i> ; }
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
Release Information	Statement introduced in Junos OS Release 14.1.
Description	Specify the IPv6 prefix length associated with a subscriber's basic broadband bridging device that is subject to a limited number of sessions.
Options	<p><i>dslite-ipv6-prefix-length</i>—Subnet prefix representing the size of the subnet subject to session limitation.</p> <p>Values: 56, 64, 96, 128</p> <p>Default: 0—no limitation.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • DS-Lite Per Subnet Limitation Overview on page 226

software-rules

Syntax	(software-rule <i>rule-name</i> software-rule-sets <i>rule-set-name</i>);
Hierarchy Level	[edit services service-set <i>service-set-name</i>]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Specify the DS-Lite or 6rd rules or rule set included in this service set. You can configure multiple rules; however, you can only configure one rule set for each service set.
Options	<p><i>rule-name</i>—Identifier for the collection of terms that constitute this rule.</p> <p><i>rule-set-name</i>—Identifier for the set of rules to be included.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none"> • Configuring Service Rules

v6rd

Syntax	<pre>v6rd v6rd-softwire-concentrator { ipv4-prefix <i>ipv4-prefix</i>; v6rd-prefix <i>ipv6-prefix</i>; mtu-v4 <i>mtu-v4</i>; softwire-address <i>ipv4-address</i>; }</pre>
Hierarchy Level	[edit services softwire softwire-concentrator]
Release Information	Statement introduced in Junos OS Release 10.4.
Description	Configure settings for a 6rd concentrator used to process IPv6 packets encapsulated in IPv4 packets.
Options	<p><i>ipv4-prefix</i>—IPv4 prefix of the customer edge (CE) network</p> <p><i>ipv6-prefix</i>—IPv6 prefix of the 6rd domain.</p> <p><i>mtu-v4</i>—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.</p> <p><i>address</i>—IPv4 address of a softwire concentrator. This is an IPv4 address independent of any interface and on a different prefix.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Documentation	<ul style="list-style-type: none">• Configuring a 6rd Softwire Concentrator on page 160

PART 3

Administration

- [Monitoring CGN and Softwire Tunnels on page 195](#)
- [Logging on page 203](#)
- [High Availability and Load Balancing on page 207](#)
- [Protecting Against Denial of Service Attacks on page 225](#)
- [Network Address Translation Operational Mode Commands on page 229](#)

CHAPTER 13

Monitoring CGN and Software Tunnels

- [Monitoring CGN, Stateful Firewall, and Software Flows on page 195](#)
- [Monitoring Stateful Firewall Conversations on page 196](#)
- [Monitoring Global Stateful Firewall Statistics on page 196](#)
- [Monitoring NAT Pool Usage on page 197](#)
- [Monitoring Port Control Protocol Operations on page 197](#)
- [Monitoring Software Statistics on page 199](#)
- [Ping and Traceroute for DS-Lite on page 201](#)

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			
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			

- Related Documentation**
- *NAT Objects MIB in SNMP MIBs and Traps Reference*
 - *Network Address Translation Resources—Monitoring MIB in SNMP MIBs and Traps Reference*
 - *Juniper Networks Enterprise-Specific NAT Traps on SRX Series Services Gateways in SNMP MIBs and Traps Reference*

Monitoring Stateful Firewall Conversations

Purpose Use the **show services stateful-firewall conversations** command to show conversations, or collections of related flows.

Action user@host# **show services stateful-firewall conversations**

Interface: sp-0/0/0, Service set: sset

Conversation: ALG protocol: tcp

Number of initiators: 1, Number of responders: 1

Flow State Dir Frm

count

TCP 10.0.0.1:1025 -> 128.0.0.1:80 Forward I 372755

NAT source 10.0.0.1:1025 -> 129.0.0.1:1024

Software 2001:0:0:1::1 -> 1001::1

TCP 128.0.0.1:80 -> 129.0.0.1:1024 Forward O 794083

NAT dest 129.0.0.1:1024 -> 10.0.0.1:1025

Software 2001:0:0:1::1 -> 1001::1

Monitoring Global Stateful Firewall Statistics

Purpose Use the **show services stateful-firewall statistics** command to observe statistics for service sets containing software rules.

Action user@host# **show services stateful-firewall statistics**
 Interface Service set Accept Discard Reject Errors
 sp-0/0/0 dslite-svc-set2 118991296 0 0 0
 sp-0/1/0 dslite-svc-set1 237615050 0 0 0

Monitoring NAT Pool Usage

Purpose Use the **show services nat pool detail** command to find global NAT statistics related to pool usage. This command is frequently used in conjunction with the **show services stateful-firewall statistics** command.

Action user@host# **show services nat pool detail**

```
Interface: ms-1/0/0, Service set: s1
  NAT pool: dest-pool, Translation type: DNAT-44
    Address range: 10.10.10.2-10.10.10.2
  NAT pool: napt-pool, Translation type: NAPT-44
    Address range: 50.50.50.1-50.50.50.254
    Port range: 1024-63487, Ports in use: 0, Out of port errors: 0, Max ports
used: 0
  NAT pool: source-dynamic-pool, Translation type: DYNAMIC NAT44
    Address range: 40.40.40.1-40.40.40.254
    Out of address errors: 0, Addresses in use: 0
  NAT pool: source-static-pool, Translation type: BASIC NAT44
    Address range: 30.30.30.1-30.30.30.254
```

- Related Documentation**
- *NAT Objects MIB in SNMP MIBs and Traps Reference*
 - *Network Address Translation Resources—Monitoring MIB in SNMP MIBs and Traps Reference*
 - *Juniper Networks Enterprise-Specific NAT Traps on SRX Series Services Gateways in SNMP MIBs and Traps Reference*

Monitoring Port Control Protocol Operations

You can monitor Port Control Protocol (PCP) operations with the following operational commands:

- **show services nat mappings pcip**
- **show services nat mappings endpoint-independent**
- **show services pcip statistics protocol**

The following are examples of the output of these commands.

```
user@host> show services nat mappings pcip
Interface: sp-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

```

```

user@host> show services nat mappings endpoint-independent

```

```

Interface: sp-0/0/0, Service set: in

```

NAT pool: p

```

Mapping         : 10.1.1.2              :57400 --> 8.8.8.8      : 1024
Session Count   : 0
Mapping State    : Timeout
PCP Client      : 10.1.1.2                PCP lifetime : 991
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 : 190
Mapping         : 88.1.1.3              : 4001 --> 70.70.70.2 :58989
Session Count   : 1
Mapping State    : Active
B4 Address      : 2222::1

```

```

user@host> show services pcg statistics protocol

```

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

Monitoring Softwire Statistics

Purpose You can review softwire global statistics by using the [show services softwire](#) or [show services softwire 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
-----
Software 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 Software ID :0
No Flow Extension :0
ICMPv4 Dropped Packets :0
```

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.

CHAPTER 14

Logging

- [Log Generation on page 203](#)
- [Configuring NAT Session Logs on page 204](#)

Log Generation

The Multiservices PIC uses the system logging protocol to generate session logging. System log messages can be sent directly from the services PIC to an external system logging server. This requires that the services PIC interface have an IP address and appropriate system logging options configured, as in this example:

```
[edit interfaces sp-5/0/0]
services-options {
  syslog {
    host 130.0.0.1 {
      services any;
    }
  }
}
unit 0 {
  family inet {
    address 150.0.0.1/32;
  }
}
```

Log Format

For each session, three logs are generated. The three logs allow correlation of start and end times for each session.

```
Jun 28 15:29:20 cypher (FPC Slot 5, PIC Slot 0) {sset2}[FWNAT]:
ASP_SFW_CREATE_ACCEPT_FLOW: proto 6 (TCP) application: any,
ge-1/3/5.0:10.0.0.1:8856 -> 128.0.0.2:80, creating forward or watch flow ; source
address and port translate to 129.0.0.1:1028
Jun 28 15:29:23 cypher (FPC Slot 5, PIC Slot 0)
{sset2}[FWNAT]:ASP_NAT_POOL_RELEASE: natpool release 129.0.0.1:1028[1]
Jun 28 15:29:23 cypher (FPC Slot 5, PIC Slot 0) {sset2}[FWNAT]:
ASP_SFW_DELETE_FLOW: proto 6 (TCP) application: any, (null)(null)10.0.0.1:8856
-> 128.0.0.2:80, deleting forward or watch flow ; source address and port translate
to 129.0.0.1:1028
```

Log format varies somewhat depending on interface card. The example show is for the MS-DPC.

System Log Throttling

You can limit logging with the **message-rate-limit** command.



NOTE: in next-hop based service sets, the log sent to remote syslog server comes to the packet forwarding engine (PFE) via the output services PIC interface. This means that you must configure routing properly in the routing-instance where the output interface is configured.

Related Documentation

- *message-rate-limit*
- *Configuring System Logging for Service Sets*

Configuring NAT Session Logs

You can configure session logs for NAT from the CLI. By default, session open and close logs are produced. However, you can request that only one type of log be produced.

To configure NAT session logs:

1. Go to the **[edit services service-set service-set-name syslog host class classname]** hierarchy level.

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

2. Configure NAT logging using the **nat-logs** configuration statement.

```
[edit services service-set service-set-name syslog host class classname]  
user@host# set nat-logs.
```

3. Configure session logging using the **session-logs** statement. Open and close logs are produced by default. Specify **open** or **close** to produce only one type of log.

```
[edit services service-set service-set-name syslog host class classname]  
user@host# set session-logs.
```

Or

```
[edit services service-set service-set-name syslog host class classname]  
user@host# set session-logs open.
```

Or

```
[edit services service-set service-set-name syslog host class classname]  
user@host# set session-logs close.
```

4. For NAT sessions that use secured port block allocation (PBA), enter the **pba-interim-logging interval** option.

```
[edit services service-set service-set-name syslog host class classname]  
user@host# top.  
[edit]  
user@host# set interfaces interface-name service-options  
pba-interim-logging-intervale.
```

- Related Documentation**
- *Configuring System Logging for Service Sets*
 - *Interim Logging for Port Block Allocation*

High Availability and Load Balancing

- [Inter-Chassis High Availability for MS-MIC and MS-MPC on page 207](#)
- [High Availability and Load Balancing for 6rd Softwires on page 219](#)

Inter-Chassis High Availability for MS-MIC and MS-MPC

Inter-chassis high availability supports stateful synchronization of services using a switchover to a backup services PIC on a different chassis. The feature is described in the following topics:

- [Inter-Chassis High Availability for Stateful Firewall and NAPT44 Overview \(MS-MIC, MS-MPC\) on page 207](#)
- [Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 \(MS-MPC, MS-MIC\) on page 208](#)
- [Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall \(MS-MIC, MS-MPC\) on page 209](#)

Inter-Chassis High Availability for Stateful Firewall and NAPT44 Overview (MS-MIC, MS-MPC)

Carrier-grade NAT (CGN) deployments can use dual-chassis implementations to provide a redundant data path and redundancy for key components in the router. Although intra-chassis high availability can be used in dual-chassis environments, it deals only with service PIC failures. If traffic is switched to a backup router due to some other failure in the router, state is lost. Inter-chassis high availability preserves state and provides redundancy using fewer service PICs than intra-chassis high availability. Only long-lived flows are synchronized between the master and backup chassis in the high availability pair. The service PICs do not replicate state until an explicit CLI command, **request services redundancy (synchronize | no-synchronize)**, is issued to start or stop the state replication. Stateful firewall, NAPT44, and APP state information can be synchronized.



NOTE: When both the master and backup PICs are up, replication starts immediately when the **request services redundancy** command is issued.

In order to use Inter-chassis high availability, you must use service sets configured for next-hop service interfaces. Inter-chassis high availability works with ms- service interfaces

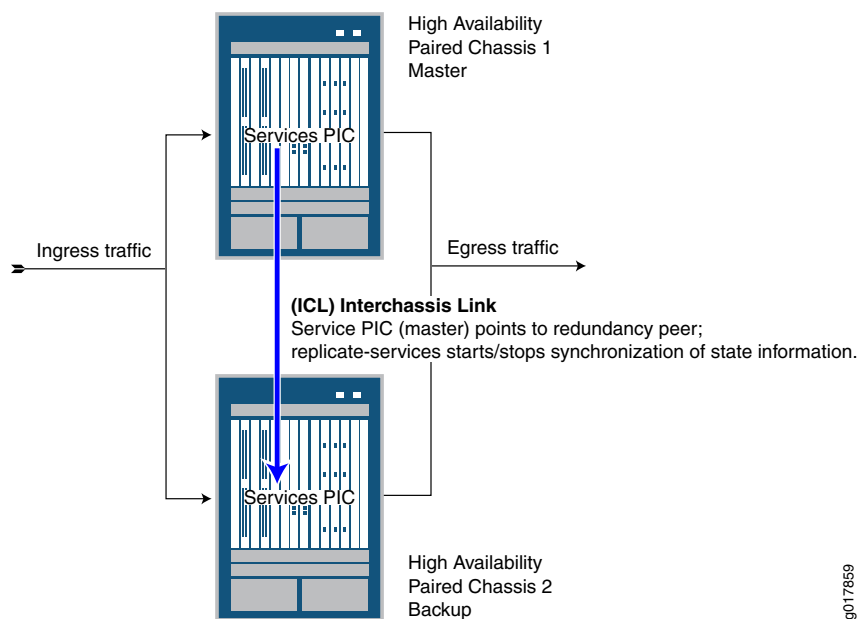
configured on MS-MIC or MS-MPC interface cards. A unit other than unit 0 must be configured with the **ip-address-owner service-plane** option.

The following restrictions apply:

- NAPT44 is the only translation type supported.
- Checkpointing is not supported for ALGs, PBA port block allocation (PBA), endpoint-independent mapping (EIM), or endpoint-independent filters (EIF).

Figure 12 on page 208 shows the inter-chassis high availability topology.

Figure 12: Inter-Chassis High Availability Topology



Configuring Inter-Chassis High Availability for Stateful Firewall and NAPT44 (MS-MPC, MS-MIC)

To configure inter-chassis availability for stateful firewall and NAPT44 on MS-MIC or MS-MPC service PICS, perform the following configuration steps on each chassis of the high availability pair:

1. At the **[edit interfaces *interface-name* redundancy-options]** hierarchy level, set the **ipaddress** for the **redundancy-peer**. This IPv4 address specifies one of the hosted IP addresses of the remote PIC. This address is used by the TCP channel between the HA pairs.

```
[edit interfaces interface-name redundancy-options]
user@host# set redundancy-peer ipaddress ipaddress
```



NOTE: When you enable or disable high availability of MS-MICs or MS-MPCs by configuring or removing the primary and backup adaptive services PICs by using the `redundancy-options redundancy-peer ipaddress address` statement at the `[edit interfaces interface-name]` hierarchy level, the configuration change is treated as a catastrophic event for each service-set that refers to the affected interface at the `[edit services service-set name interface-service service-interface interface-name]` hierarchy level. A catastrophic event at the service-set level has the effect of deactivating the service set, applying the change, and then reactivating the service set.

2. Specify the name of a special routing instance, or VRF, you want applied to the HA synchronization traffic between the high availability pair.

```
[edit interfaces interface-name redundancy-options]
user@host# set routing-instance instance-name
```

3. For the service set defining an interface that is a member of the high availability pair, configure the service replication options using the `replicate-services` option.

```
[edit services service-set service-set-name replicate-services]
user@host# set replication-threshold threshold-value
stateful-firewall
nat
```

Example: Inter-Chassis Stateful High Availability for NAT and Stateful Firewall (MS-MIC, MS-MPC)

This example shows how to configure inter-chassis high availability for stateful firewall and NAT services.

- [Requirements on page 209](#)
- [Overview on page 209](#)
- [Configuration on page 210](#)

Requirements

This example uses the following hardware and software components:

- Two MX480 routers with MS-MPC line cards
- Junos OS Release 13.3 or later

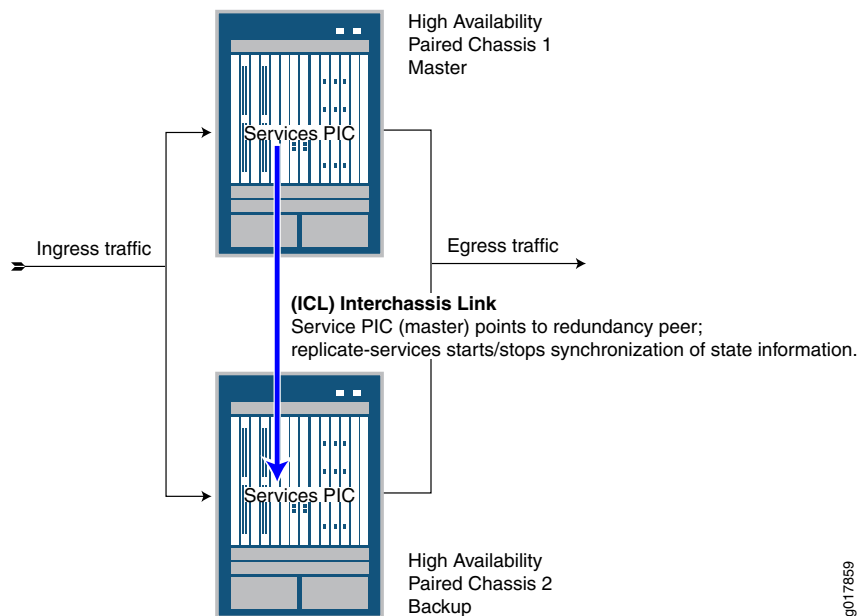
Overview

Two MX 3D routers are identically configured to facilitate stateful failover for firewall and NAT services in case of a chassis failure.

Topology

[Figure 13 on page 210](#) shows the inter-chassis high availability topology.

Figure 13: Inter-Chassis High Availability Topology



g017859

Configuration

To configure inter-chassis high availability for this example, perform these tasks:

- [Configuring Interfaces for Chassis 1 on page 212](#)
- [Configure Routing Information for Chassis 1 on page 213](#)
- [Configuring NAT and Stateful Firewall for Chassis 1 on page 214](#)
- [Configuring the Service Set on page 215](#)
- [Configuring Interfaces for Chassis 2 on page 216](#)
- [Configure Routing Information for Chassis 2 on page 218](#)

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 ms-4/0/0 redundancy-options redundancy-peer ipaddress 5.5.5.2
set interfaces ms-4/0/0 redundancy-options routing-instance HA
set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
set interfaces ms-4/0/0 unit 10 family inet address 5.5.5.1/32
set interfaces ms-4/0/0 unit 20 family inet
set interfaces ms-4/0/0 unit 20 service-domain inside
set interfaces ms-4/0/0 unit 30 family inet
set interfaces ms-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 ms-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 5.5.1/32 next-hop
    ms-4/0/0.10
set routing-instances HA routing-options static route route 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 ms-4/0/0.20
set services service-set ss2 next-hop-service outside-service-interface ms-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 ms-4/0/0 redundancy-options routing-instance HA
set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
set interfaces ms-4/0/0 unit 10 family inet address 5.5.2/32
set interfaces ms-4/0/0 unit 20 family inet
set interfaces ms-4/0/0 unit 20 service-domain inside
set interfaces ms-4/0/0 unit 30 family inet
set interfaces ms-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 ms-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 ms-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 ms-4/0/0.20
set services service-set ss2 next-hop-service outside-service-interface ms-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 ms-4/0/0 redundancy-options redundancy-peer
ipaddress 5.5.5.2
user@host# set interfaces ms-4/0/0 redundancy-options routing-instance HA
user@host# set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
user@host# set interfaces ms-4/0/0 unit 10 family inet address 5.5.5.1/32
user@host# set interfaces ms-4/0/0 unit 20 family inet
user@host# set interfaces ms-4/0/0 unit 20 service-domain inside
user@host# set interfaces ms-4/0/0 unit 30 family inet
user@host# set interfaces ms-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;
        }
    }
}
ms-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 ms-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 ms-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 ms-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 ms-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 MS-PIC.

```
user@host# set services service-set ss2 next-hop-service inside-service-interface
ms-4/0/0.20
user@host# set services service-set ss2 next-hop-service outside-service-interface
ms-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 ms-3/0/0.20;
    outside-service-interface ms-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 ms-4/0/0 on chassis 1 that contains the **ip-address-owner service-plane** statement.

```
[edit interfaces]
set interfaces ms-4/0/0 redundancy-options redundancy-peer ipaddress 5.5.5.1
user@host# set interfaces ms-4/0/0 redundancy-options routing-instance HA
user@host# set interfaces ms-4/0/0 unit 10 ip-address-owner service-plane
user@host# set interfaces ms-4/0/0 unit 10 family inet address 5.5.5.2/32
user@host# set interfaces ms-4/0/0 unit 20 family inet
user@host# set interfaces ms-4/0/0 unit 20 service-domain inside
user@host# set interfaces ms-4/0/0 unit 30 family inet
user@host# set interfaces ms-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
ms-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 ms-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
ms-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 ms-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 ms-4/0/0.10;
            route 5.5.5.1/32 next-hop 20.1.1.1;
        }
    }
}

```

High Availability and Load Balancing for 6rd Softwires

- [Load Balancing a 6rd Domain Across Multiple Services PICs on page 219](#)
- [Example: Load Balancing a 6rd Domain Across Multiple Services PICs on page 219](#)
- [Configuring High Availability for 6rd Using 6rd Anycast on page 224](#)

Load Balancing a 6rd Domain Across Multiple Services PICs

The 6rd domain is an IPv6 network, which can potentially be very large. A single PIC, or network processing unit (NPU) on a Multiservices DPC, might not be able to handle all the traffic for the 6rd domain. To alleviate load problems, you can load-balance the 6rd domain traffic across multiple PICs. To do so, assign the same software rule to different services sets that use different interfaces. Configure explicit routes and equal-cost multipath (ECMP) to load-balance the 6rd traffic.

Example: Load Balancing a 6rd Domain Across Multiple Services PICs

- [Hardware and Software Requirements on page 219](#)
- [Overview on page 220](#)
- [Configuration on page 220](#)

Hardware and Software Requirements

This example requires the following hardware:

- An MX Series 3D Universal Edge router with a services DPC with two available NPUs or an M Series Multiservice Edge router with two services PICs available for 6rd software concentrator processing
- A domain name server (DNS)

This example uses the following software:

- Junos OS Release 11.4 or higher

Overview

Because of anticipated volume, a provider needs to balance 6rd software traffic between two services PICs.

Configuration

- [Chassis Configuration on page 220](#)
- [Software Concentrator and Software Rule Configuration on page 221](#)
- [Stateful Firewall Configuration on page 221](#)
- [Service Set Configuration on page 222](#)
- [Load-Balancing Configuration on page 222](#)

Chassis Configuration

Step-by-Step Procedure

To configure the chassis:

1. Define the ingress interface and its properties.

```
user@host# edit interfaces ge-1/2/0
user@host# set unit 0 family inet address 10.10.10.1/16
```
2. Define the egress interface and its properties. In this example, the IPv6 clients try to reach the IPv6 server at 3abc::2/16.

```
user@host# edit interfaces ge-1/2/2
user@host# set unit 0 family inet6 address 3ABC::1/16
```
3. Define the services PICs for selection as software concentrators by the load-balancing process. This configuration uses two PICs/NPUs: sp-3/0/0 and sp-3/1/0. A next-hop style service set is configured (shown in the next section).

```
user@host# edit interfaces sp-3/0/0
[edit interfaces ge-3/0/0]
user@host# set services-options syslog host local services any
user@host# set unit 0 family inet
user@host# set unit 0 family inet6
user@host# set unit 1 family inet service-domain inside
user@host# set unit 1 family inet service-domain outside
user@host# set unit 2 family inet service-domain inside
user@host# set unit 2 family inet service-domain outside
user@host# up 1
[edit]
user@host# edit interfaces sp-3/1/0
[edit interfaces sp-3/1/0]
user@host# set services-options syslog host local services any
user@host# set unit 0 family inet
user@host# set unit 0 family inet6
user@host# set unit 1 family inet service-domain inside
user@host# set unit 1 family inet service-domain outside
user@host# set unit 2 family inet service-domain inside
user@host# set unit 2 family inet service-domain outside
```

Software Concentrator and Software Rule Configuration

Step-by-Step Procedure The software configuration is straightforward. In this example, the 6rd domain prefix is 3040::0/16, the 6rd software concentrator IPv4 address is 30.30.30.1, and the customer IPv4 network is 10.10.0.0/16. In the customer premises equipment (CPE) network, all customer edge (CE) devices have addresses that belong to the 10.10.0.0/16 network. To configure the software:

1. Go to the **[edit services software]** hierarchy level.

```
user@host# edit services software
```
2. Configure IPv6 multicast.

```
[edit services software]
user@host# set ipv6-multicast-interfaces all
```
3. Go to the software concentrator v6rd hierarchy level and name the software concentrator **shenick01-rd1**.

```
[edit services software]
user@host# edit software-concentrator v6rd shenick01-rd1
```
4. Configure the software concentrator properties.

```
[edit services software software-concentrator v6rdshenick01-rd1 ]
user@host# set software-address 30.30.30.1
user@host# set ipv4-prefix 10.10.0.0/16
user@host# set v6rd-prefix 3040::/16
user@host# set mtu-v4 9192
```
5. Configure a software rule for incoming 6rd traffic.

```
[edit services software software-concentrator v6rd shenick01-rd1 ]
user@host# up 1
[edit services software ]
user@host# edit rule shenick01-r1
[edit services software rule shenick01-r1]
user@host# set match-direction input
user@host# set term t1 then v6rd shenick01-rd1
```

Stateful Firewall Configuration

Step-by-Step Procedure To configure the stateful firewall rule:

1. Go to the stateful firewall hierarchy level and define a rule.

```
user@host# edit services stateful-firewall rule r1
```
2. Set the match direction.

```
[edit services stateful-firewall rule r1]
user@host# set match-direction input-output
```
3. Configure a term that accepts all traffic.

```
[edit services stateful-firewall rule r1]
user@host# set term t1 then accept
```

Service Set Configuration

Step-by-Step Procedure This configuration provides two service sets, each pointing to a different network processing unit (NPU). Both service sets use the same stateful firewall and software rules. Because they use the same software rule, they refer to same 6rd software concentrator. This results in the software concentrator being hosted on both the NPUs.

To configure the service set:

1. Define a service set for the first NPU.

```
user@host# edit services service-set v6rd-sset1
```
2. Configure the software and stateful firewall rules for the first NPU.

```
[edit services service-set v6rd-sset1]  
user@host# set software-rules shenick01-r1  
user@host# set stateful-firewall-rules r1
```
3. Configure the inside and outside interfaces for the next-hop service.

```
[edit services service-set v6rd-sset1]  
user@host# set next-hop-service inside-service-interface sp-3/0/0.1  
user@host# set next-hop-service outside-service-interface sp-3/0/0.2
```
4. Define a service set for the second NPU.

```
user@host# edit services service-set v6rd-sset2
```
5. Configure the software and stateful firewall rules for the second NPU.

```
[edit services service-set v6rd-sset2]  
user@host# set software-rules shenick01-r1  
user@host# set stateful-firewall-rules r1
```
6. Configure the inside and outside interfaces for the next-hop service.

```
[edit services service-set v6rd-sset1]  
user@host# set next-hop-service inside-service-interface sp-3/1/0.1  
user@host# set next-hop-service outside-service-interface sp-3/1/0.2
```

Load-Balancing Configuration

Step-by-Step Procedure To configure load balancing:
Configure explicit routes and ECMP to load-balance the 6rd traffic. Configure explicit routes for both the 6rd concentrator IPv4 address and the 6rd domain prefix, so that they point to both NPUs.

1. To configure static routes for the 6rd domain using the routing-table inet6.0, go to the **[edit forwarding-options rib inet6.0 static]** hierarchy level and set the routes for the 6rd domain and the 6rd concentrator IPv4 address.

```
user@host edit forwarding-options rib inet6.0 static  
[edit forwarding-options rib inet6.0 static]  
user@host# set route 3040::0/16 next-hop [ sp-3/0/0.2 sp-3/1/0.2 ]  
user@host# set route 30.30.30.1/32 next-hop [ sp-3/0/0.1 sp-3/1/0.1 ]
```


The service PIC daemon (spd) also adds default routes to these addresses pointing to the NPUs. However, the routes added by the spd use different metrics, which are computed based on the FPC, PIC, slot numbers, and subunit of the services PIC if used in the service set configuration. The static routes configured in this sample configuration will have metrics of 5 and therefore a higher preference than the spd-added routes.

The explicitly configured routes are as follows:

```
root@router# run show route 30.30.30.1
inet.0: 37 destinations, 40 routes (36 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

30.30.30.1/32      *[Static/5] 00:00:10
                  > via sp-3/0/0.1
                  via sp-3/1/0.1
                  [Static/786433] 00:23:03
                  > via sp-3/0/0.1
                  [Static/851969] 00:00:09
                  > via sp-3/1/0.1

root@router# run show route 3040::/16
inet6.0: 20 destinations, 33 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3040::/16         *[Static/5] 00:00:15
                  via sp-3/0/0.2
                  > via sp-3/1/0.2
                  [Static/786434] 00:23:08
                  > via sp-3/0/0.2
                  [Static/851970] 00:00:14
                  > via sp-3/1/0.2
```



BEST PRACTICE: The spd-installed routes have higher metric values (hence a low preference) and the metrics are different. If the metrics are different and ECMP is not enabled, even though multiple routes exist for the same destination, only one of the routes is picked up all the time (based on the metric). For ECMP you must configure equal-cost routes, and hence a manual configuration of routes is needed as shown above.

2. Configure equal-cost multipath (ECMP) load balancing by configuring the hash key at the **[edit forwarding-options hash-key]** hierarchy level.

```
user@host# forwarding-options hash-key
[edit forwarding-options hash-key]
user@host# set family inet layer-3 destination-address
user@host# set family inet layer-3 source-address
user@host# set family inet6 layer-3 destination-address
user@host# set family inet6 layer-3 source-address
```

3. Verify your configuration by displaying **forwarding-options**.

```
user@host# show forwarding-options
hash-key {
    family inet { <== IPv4 traffic from CEs uses this
```

```
layer-3 {  
    destination-address;  
    source-address;  
}  
}  
family inet6 { <== IPv6 traffic from Internet uses this  
    layer-3 {  
        destination-address;  
        source-address;  
    }  
}  
}
```



TIP: Both IPv4 and IPv6 hash keys must be configured. The IPv4 hash key is used to distribute the traffic coming from CPE devices to the 6rd branch relay. The IPv6 hash key is used to distribute the traffic coming from the IPv6 Internet to the 6rd domain. Because the hash in the forward and reverse direction is for different families, different flows from the same session can reside on different NPUs. However, 6rd processing is stateless (as far as mapping IPv6 packets to softwires is concerned), so this should not be a problem.

Configuring High Availability for 6rd Using 6rd Anycast

You configure 6rd Anycast by defining two service sets that use the same softwire rule in both service sets, just as you do when you configure load balancing for 6rd. However, you do not configure ECMP, and as a result, the services PIC daemon (spd) installs two routes *each* for the softwire concentrator address and 6rd domain pointing to each service interface. The forwarding plane can select any route based on the priority, which is computed when the spd installs the routes. The priority is computed based on the FPC, PIC, slot numbers, and subunit number used on the sp- interface. *Only one PIC is used* based on the route priority, and that PIC gets all of the 6rd traffic. If the PIC goes down, the route pointing to it is also deleted and the forwarding plane automatically selects the alternate available PIC.

6rd Anycast is completely stateless. The spd installs the route and doesn't run any state machine for the PIC. Because the routes are pre-installed and service sets are already on the PIC, there is no service delay if a failover occurs.

Related Documentation

- [Junos Address Aware Carrier Grade NAT and IPv6 Feature Guide](#)

CHAPTER 16

Protecting Against Denial of Service Attacks

- [Protecting CGN Devices Against Denial of Service \(DOS\) Attacks on page 225](#)
- [DS-Lite Subnet Limitation on page 225](#)

Protecting CGN Devices Against Denial of Service (DOS) Attacks

You can now choose configuration options that help prevent or minimize the effect of attempted denial of service (DOS) attacks.

- [Mapping Refresh Behavior on page 225](#)
- [EIF Inbound Flow Limit on page 225](#)

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.

DS-Lite Subnet Limitation

- [DS-Lite Per Subnet Limitation Overview on page 226](#)
- [Configuring DS-Lite Per Subnet Session Limitation to Prevent Denial of Service Attacks on page 226](#)

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.

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 prefix the length is 128, then each IPv6 address is treated as a B4. Port blocks are allocated per 128-bit V6 address.
- Session limit, defined under the DSLite 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 alloc 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.

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

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@router# set services service-set service-set-name software-options
  dslite-ipv6-prefix-length 56.
```



NOTE: Ensure that all mappings are cleared before changing the prefix length.

2. Configure the maximum number of subscriber sessions allowed per prefix. You can configure from 0 through 16,384 sessions.

[edit]

```
user@router# set services software software-concentrator dslite  
dslite-concentrator-name session-limit-per-prefix 12
```



NOTE: You cannot use flow-limit and session-limit-per-prefix in the same dslite configuration.

CHAPTER 17

Network Address Translation Operational Mode Commands

- `clear services inline nat pool`
- `clear services inline nat statistics`
- `clear services nat flows`
- `clear services nat mappings`
- `clear services nat mappings app`
- `clear services nat mappings eim`
- `clear services nat mappings pcp`
- `clear services nat statistics`
- `show services inline nat pool`
- `show services inline nat statistics`
- `show services nat ipv6-multicast-interfaces`
- `show services nat pool`
- `show services nat mappings`
- `show services nat statistics`
- `show services pcp statistics`
- `show services software`
- `show services software flows`
- `show services software statistics`
- `show services stateful-firewall conversations`
- `show services stateful-firewall flows`
- `show services stateful-firewall statistics`

clear services inline nat pool

Syntax	clear services inline nat pool <i>pool-name</i>
Release Information	Command introduced in Junos OS Release 11.4.
Description	Clear global inline NAT statistics.
Options	pool-name —Name of the NAT pool for which statistic are cleared.
Required Privilege Level	clear
List of Sample Output	clear services inline nat pool on page 230
Output Fields	When you enter this command, the NAT pool statistics are cleared. There is no specific output.

Sample Output

clear services inline nat pool

```
user@host> clear services inline nat pool p1
```


clear services inline nat statistics

Syntax	clear services inline nat statistics <interface <i>interface-name</i> >
Release Information	Command introduced in Junos OS Release 11.4.
Description	Clear global inline NAT statistics.
Options	interface <i>interface-name</i> —(Optional) Clear inline NAT statistics for the specified interface only.
Required Privilege Level	clear
List of Sample Output	clear services inline nat statistics on page 231
Output Fields	When you enter this command, the global inline NAT statistics are cleared. There is no specific output.

Sample Output

clear services inline nat statistics

```
user@host> clear services inline nat statistics
```

clear services nat flows

Syntax	clear services nat flows <b4address b4address> <service-set service-set> <subscriber subscriber-address>
Release Information	Command introduced in Junos OS Release 14.1.
Description	Clear NAT flows.
Options	<p>none—Clear all NAT flows.</p> <p>b4address b4address—(Optional) Clear NAT flows for a particular B4 address.</p> <p>service-set service-set—(Optional) Clear NAT flows for a particular service set.</p> <p>subscriber ip—(Optional) Clear NAT flows for a particular subscriber, identified by IPv4 address.</p>
Required Privilege Level	view
Related Documentation	
List of Sample Output	clear services nat flows subscriber (IPv4 address) on page 232
Output Fields	Table 6 on page 232 lists the output fields for the clear services nat flows command. Output fields are listed in the approximate order in which they appear.

Table 6: clear services nat flows Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Flows removed	Number of flows removed.

Sample Output

clear services nat flows subscriber (IPv4 address)

```

user@host> clear services nat flows subscriber ip 3.3.3.3
Interface  Service set  Flows removed

sp-2/0/0   ss1             0

```

Sample Output

clear services nat mappings

Syntax	clear services nat mappings <app> <eim> <pcp> <service-set <i>service-set</i> >
Release Information	Command introduced in Junos OS Release 14.1.
Description	Clear NAT mappings.
Options	<p>none—Clear all NAT mappings.</p> <p>app—(Optional) Clear address-pooling paired NAT mappings.</p> <p>eim—(Optional) Clear endpoint-independent NAT mappings.</p> <p>pcp—(Optional) Clear Port Control Protocol NAT mappings.</p> <p>service-set <i>service-set</i>—(Optional) Clear NAT mappings for a specified service set..</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show services nat mappings on page 249 • clear services nat mappings app on page 235 • clear services nat mappings eim on page 236 • clear services nat mappings pcp on page 238
List of Sample Output	clear services nat mappings on page 234
Output Fields	Table 7 on page 233 lists the output fields for the clear services nat mappings command. Output fields are listed in the approximate order in which they appear.

Table 7: clear services nat mappings Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings

```
user@host> clear services nat mappings
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1          0                  0
```

clear services nat mappings app

Syntax	clear services nat mappings app <b4address <i>b4address/prefix</i> > <service-set <i>service-set</i> > <subscriber <i>subscriber-ipv4-address</i> >
Release Information	Command introduced in Junos OS Release 14.1.
Description	Clear NAT mappings for address pooling paired (app).
Options	<p>none—Clear all NAT app mappings.</p> <p>b4address <i>b4address/prefix</i>—(Optional) Clear NAT APP mappings for a particular subscriber <i>b4address/prefix</i></p> <p>service-set <i>service-set</i>—(Optional) Clear NAT APP mappings for a specified service set..</p> <p>subscriber <i>subscriber-ipv4-address/prefix</i>—(Optional) Clear NAT APP mappings for a particular subscriber <i>ipv4-address/prefix</i></p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show services nat mappings on page 249
List of Sample Output	clear services nat mappings app on page 235
Output Fields	Table 8 on page 235 lists the output fields for the clear services nat mappings app command. Output fields are listed in the approximate order in which they appear.

Table 8: clear services nat mappings app Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings app

```

user@host> clear services nat mappings app
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1           0                  0

```

clear services nat mappings eim

Syntax	clear services nat mappings eim <b4address <i>b4address/prefix</i> > <subscriber <i>subscriber-ipv4-address</i> >
Release Information	Command introduced in Junos OS Release 14.1.
Description	Clear endpoint independent (EIM) and port control protocol (PCP) mappings .
Options	<p>none—Clear all EIM and PCP mappings.</p> <p>b4address <i>b4address/prefix</i>—(Optional) Clear EIM and PCP mappings for a particular subscriber <i>b4address/prefix</i></p> <p>internal-host <i>ipv4address/prefix</i>—(Optional) Clear EIM and PCP mappings matching the specified <i>b4address</i> and internal-host..</p> <p>port <i>port</i>—(Optional) Clear EIM and PCP mappings matching the specified <i>b4address</i>, internal host, and port.</p> <p>service-set <i>service-set</i>—(Optional) Clear EIM and PCP mappings for the specified service set.</p> <p>subscriber <i>subscriber-ipv4-address/prefix</i>—(Optional) Clear EIM and PCP mappings for a particular subscriber <i>ipv4-address/prefix</i></p> <ul style="list-style-type: none"> port <i>port</i>—(Optional) Clear EIM and PCP mappings matching the specified <i>ipv4-address/prefix</i> and port. service-set <i>service-set</i>—(Optional) Clear EIM and PCP mappings for the specified service set.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> show services nat mappings on page 249
List of Sample Output	clear services nat mappings eim on page 237
Output Fields	<p>Table 9 on page 236 lists the output fields for the clear services nat mappings eim command. Output fields are listed in the approximate order in which they appear.</p>

Table 9: clear services nat mappings eim Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.

Table 9: clear services nat mappings eim Output Fields (*continued*)

Field Name	Field Description
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings eim

```
user@host> clear services nat mappings eim
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1          0                  0
```

clear services nat mappings pcg

Syntax	clear services nat mappings pcg <b4address <i>b4address/prefix</i> > <subscriber <i>subscriber-ipv4-address</i> >
Release Information	Command introduced in Junos OS Release 14.1.
Description	Clear NAT mappings for Port Control Protocol (PCP).
Options	<p>none—Clear all NAT PCP mappings.</p> <p>b4address <i>b4address/prefix</i>—(Optional) Clear NAT PCP mappings for a particular subscriber <i>b4address/prefix</i></p> <p>port <i>port</i>—(Optional) Clear NAT PCP mappings matching the specified <i>b4address</i> internal host, and port.</p> <p>service-set <i>service-set</i>—(Optional) Clear NAT PCP mappings for the specified service set.</p> <p>subscriber <i>ipv4-address/prefix</i>—(Optional) Clear NAT PCP mappings for a particular subscriber <i>ipv4-address/prefix</i></p> <p>port <i>port</i>—(Optional) Clear NAT PCP mappings matching the specified <i>ipv4-address/prefix</i>, and port.</p>
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none"> • show services nat mappings on page 249
List of Sample Output	clear services nat mappings pcg on page 239
Output Fields	Table 10 on page 238 lists the output fields for the clear services nat mappings pcg command. Output fields are listed in the approximate order in which they appear.

Table 10: clear services nat mappings pcg Output Fields

Field Name	Field Description
Interface	Name of a services interface.
Service set	Name of the service set from which flows are being cleared.
Mappings removed	Number of mappings removed.
Flows removed	Number of flows removed.

Sample Output

clear services nat mappings pcp

```
user@host> clear services nat mappings pcp
Interface  Service set  Mappings removed  Flows removed
sp-2/0/0   ss1          0                  0
```

clear services nat statistics

Syntax	clear services nat statistics <interface <i>interface-name</i> > <service-set <i>service-set-name</i> >
Release Information	Command introduced in Junos OS Release 11.4.
Description	Clear global NAT statistics.
Options	interface <i>interface-name</i> —(Optional) Clear NAT statistics for the specified interface only. service-set <i>service-set-name</i> —(Optional) Clear NAT statistics for the specified service set only.
Required Privilege Level	clear
Related Documentation	<ul style="list-style-type: none">• show services nat statistics on page 253
Output Fields	When you enter this command, the global NAT statistics are cleared. There is no specific output.

Sample Output

clear services nat statistics

```
user@host> clear services nat statistics
```

show services inline nat pool

Syntax	<code>show services inline nat pool</code> <code><pool pool--name></code>
Release Information	Command introduced in Junos OS Release 11.4.
Description	Display information about inline Network Address Translation (NAT) pool.
Options	<i>pool-name</i> —Display information about the specified services-inline interface NAT pool.
Required Privilege Level	view
List of Sample Output	show services inline nat pool on page 241
Output Fields	Table 11 on page 241 lists the output fields for the show services inline nat pool command. Output fields are listed in the order in which they appear.

Table 11: show services inline nat pool Output Fields

Field Name	Field Description
Interface	Name of an si interface hosted on a Trio-based line card.
NAT pool	Name of the pool used for address translations.
Translation type	Translation type specified in the applicable NAT rule for the service set.
Address range	Starting and ending public NAT addresses available for translation.
NATed packets	Number of packets translated for the specified pool.
un-NATed packets	Number of received packets that were not translated.
Errors	Number of packets with translation errors.

Sample Output

show services inline nat pool

```

user@host> show services inline nat pool p1
Interface: si-5/0/0, Service set: ss-inat
NAT pool: p1, Translation type: BASIC NAT44
Address range: 20.20.20.0-20.20.20.255
NATed packets: 0, Un-NATed packets: 0, Errors: 0

```

show services inline nat statistics

Syntax	<code>show services inline nat statistics</code> <code><interface <i>interface-name</i>></code>
Release Information	Command introduced in Junos OS Release 11.4.
Description	Display information about inline Network Address Translation (NAT) address translations.
Options	<i>interface-name</i> —(Optional) Display information about the specified NAT services-inline interface only. When a specific interface is not specified, statistics for all services-inline interfaces are shown.
Required Privilege Level	view
List of Sample Output	show services inline nat statistics on page 242
Output Fields	Table 12 on page 242 lists the output fields for the <code>show services inline nat statistics</code> command. Output fields are listed in the order in which they appear.

Table 12: show services inline nat statistics Output Fields

Field Name	Field Description	Level of Output
Service PIC	Name of an <code>si</code> interface hosted on a Trio-based line card.	All levels
Slow path packets received	Number of ICMP exception packets received for NAT translation.	All levels
Slow path packets dropped	Number of received ICMP exception packets that were dropped.	All levels

Sample Output

show services inline nat statistics

```

user@host> show services inline nat statistics
Service PIC Name                               :si-5/0/0

Slow path packets received                     :0
Slow path packets dropped                      :0

```

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 mutlicast.
Required Privilege Level	view
List of Sample Output	show services nat ipv6-multicast-interfaces on page 243
Output Fields	Table 13 on page 243 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 13: 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      Operational
                   State      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 pool

Syntax	show services nat pool <brief detail> <pool-name> pgcp <ports-per-session remotely-controlled>
Release Information	Command introduced before Junos OS Release 7.4. pgcp option added in Junos OS Release 8.5.
Description	Display information about Network Address Translation (NAT) pools.
Options	<p>none—Display standard information about all NAT pools.</p> <p>brief detail—(Optional) Display the specified level of output.</p> <p>pool-name—(Optional) Display information about the specified NAT pool.</p> <p>pgcp—(Optional) Display information about a NAT pool that is exclusive to the BGF.</p> <p>ports-per-session—(Optional) Display the number of ports allocated per session from the NAT pool.</p> <p>remotely-controlled—(Optional) Display if the NAT pool is explicitly specified by the gateway controller.</p>
Required Privilege Level	view
List of Sample Output	show services nat pool brief on page 247 show services nat pool detail on page 247 show services nat pool for Secured Port Block Allocation on page 247 show services nat pool detail for Deterministic Port Block Allocation on page 248 show services nat pool for Deterministic Port Block Allocation on page 248 show services nat pool detail for Port Block Allocation on page 248
Output Fields	Table 14 on page 245 lists the output fields for the show services nat pool command. Output fields are listed in the approximate order in which they appear.

Table 14: show services nat pool Output Fields

Field Name	Field Description	Level of Output
DetNat subscriber exceeded port limits	The number of times a subscriber exceeded its port limits for a NAT pool that uses deterministic port block allocation.	All levels.
MAX number of port blocks used	The maximum number of port blocks used.	All levels.
Port block memory allocation errors	The number of port block allocation errors.	All levels.

Table 14: show services nat pool Output Fields (*continued*)

Field Name	Field Description	Level of Output
Current number of port blocks in use	Current count of the port blocks that are being used.	
Unique pool users	The number of different users of the NAT pools.	All levels.
Interface	Name of an adaptive services interface.	All levels
Service set	Name of a service set. Individual empty service sets are not displayed, but if none of the service sets has any flows, a flow table header is printed for each service set.	All levels
NAT pool	Name of the Network Address Translation pool.	All levels
Type or Translation type	Address translation type: basic-nat-pt , basic-nat44 , basic-nat66 , deterministic-napt44 , dnat-44 , dynamic-nat44 , napt44 , napt-66 , napt-pt , stateful-nat64 , twice-basic-nat-44 , twice-dynamic-nat-44 , twice-dynamic-napt-44 .	All levels
Address or Address range	IPv4 address range of the pool.	All levels
Port or Port range	Port range of the pool. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	All levels
Ports used' or Ports in use	Number of ports allocated in this pool with this name. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	All levels
Port block type	Type of port block allocation: secured or deterministic	All levels
Out of port errors	Number of port allocation errors. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	detail
Max ports used	Maximum number of ports used. Applicable only for dynamic NAT pools. Not displayed for static NAT pools.	detail
Addresses in use	Number of addresses in use for dynamic source address NAT pools.	detail
Out of Port Errors	No more ports available to allocate.	Detail
Max Ports Used	The maximum number of ports in use at any time since the services PIC was started.	Detail
AP-P out of port errors	When address pooling paired (AP-P) is configured, a private IP is paired to a public IP. This is counter of translation errors where there are free ports available in the NAT pool, but none for the NAT IP to which the private IP is paired.	Detail
Current EIF Inbound flows count	Current count of EIF inbound flows, including all EIF flows per pool.	

Table 14: show services nat pool Output Fields (*continued*)

Field Name	Field Description	Level of Output
EIF flow limit exceeded drops	Current number of flow drops due to exceeded flow limit. This number is per pool, not per EIF mapping.	

Sample Output

show services nat pool brief

```
user@host> show services nat pool brief
```

```
Interface: ms-1/0/0, Service set: s1
NAT pool      Type      Address                      Port      Ports used
dest-pool     DNAT-44  10.10.10.2-10.10.10.2
napt-pool     NAPT-44  50.50.50.1-50.50.50.254    1024-63487  0
source-dynamic-pool DYNAMIC NAT44 40.40.40.1-40.40.40.254
source-static-pool BASIC NAT44 30.30.30.1-30.30.30.254
```

show services nat pool detail

```
user@host> show services nat pool detail
```

```
Interface: ms-1/0/0, Service set: s1
NAT pool: dest-pool, Translation type: DNAT-44
Address range: 10.10.10.2-10.10.10.2
NAT pool: napt-pool, Translation type: NAPT-44
Address range: 50.50.50.1-50.50.50.254
Port range: 1024-63487, Ports in use: 0, Out of port errors: 0, Max ports
used: 0
NAT pool: source-dynamic-pool, Translation type: DYNAMIC NAT44
Address range: 40.40.40.1-40.40.40.254
Out of address errors: 0, Addresses in use: 0
NAT pool: source-static-pool, Translation type: BASIC NAT44
Address range: 30.30.30.1-30.30.30.254
```

show services nat pool for Secured Port Block Allocation

```
user@host> show services nat pool
```

```
Interface: sp-2/0/0, Service set: in
NAT pool      Type      Address                      Port      Ports used
mypool        dynamic  3.3.3.3-3.3.3.10           512-65535  0
               3.3.3.15-3.3.3.20
               3.3.3.25-3.3.3.30
               3.3.3.95-3.3.3.200
Port block size: 64, Max port blocks per address: 1, Active block timeout: 86400,
Effective port range: 1024-65471,
Effective number of port blocks: 126882, Effective number of ports: 8120448, Port
block efficiency: nan
```

```
Interface: sp-2/1/0, Service set: in1
NAT pool      Type      Address                      Port      Ports used
mypool1        dynamic  9.9.9.1-9.9.9.254          512-65535  0
Port block size: 64, Max port blocks per address: 1, Active block timeout: 86400,
Effective port range: 1024-65471,
Effective number of port blocks: 255778, Effective number of ports: 16369792,
Port block efficiency: nan
```

show services nat pool detail for Deterministic Port Block Allocation

```
user@host> show services nat pool detail
Interface: sp-2/0/0, Service set: ss1
  NAT pool: napt_pool, Translation type: dynamic
    Address range: 5.5.5.1-5.5.5.254
    Port range: 2000-2002, Ports in use: 2, Out of port errors: 0, Max ports used:
2
    AP-P out of port errors: 188
    Max number of port blocks used: 1, Current number of port blocks in use: 1,
Port block allocation errors: 0,
  Port block memory allocation errors: 0
  DetNAT subscriber exceeded port limits: 1  <<<<<<<<
  Unique pool users: 1
```

show services nat pool for Deterministic Port Block Allocation

```
user@host> show services nat pool

Interface: sp-2/0/0, Service set: ss2
NAT pool      Type      Address                               Port      Ports Used
pba           dynamic  33.33.33.1-33.33.33.128             512-65535 6604
Port block type: Deterministic port block, Port block size: 200
```

show services nat pool detail for Port Block Allocation

```
user@host> show services nat pool detail

Interface: sp-2/0/0, Service set: s
  NAT pool: napt_pool, Translation type: dynamic
    Address range: 44.1.1.1-44.1.1.1
    Port range: 1024-65535, Ports in use: 0, Out of port errors: 0,
    Max ports used: 0
    AP-P out of port errors: 0
    Current EIF Inbound flows count: 0
    EIF flow limit exceeded drops: 0
```

Sample Output

show services nat mappings

Syntax	<pre>show services nat mappings <brief detail summary> <pool-name> <address-pooling-paired endpoint-independent pcg></pre>
Release Information	<p>Command introduced in Junos OS Release 10.1.</p> <p>summary option introduced in Junos OS Release 11.1.</p> <p>address-pooling paired option introduced in Junos OS Release 13.2.</p> <p>endpoint-independent option introduced in Junos OS Release 13.2.</p> <p>pcg option introduced in Junos OS Release 13.2.</p>
Description	Display information about Network Address Translation (NAT) address, port, and port control protocol (PCP) mappings.
Options	<p>none—Display standard information about all NAT pools.</p> <p>brief detail summary—(Optional) Display the specified level of output.</p> <p>pool-name—(Optional) Display detailed information about a specific NAT pool. Used only with detail level output.</p> <p>address-pooling-paired—(Optional) Display only information about address-pooling paired mappings.</p> <p>endpoint-independent—(Optional) Display only information about endpoint-independent mappings.</p> <p>pcg—(Optional) Display only information about port control protocol mappings.</p>



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 mappings brief on page 251](#)
[show services nat mapping detail on page 251](#)
[show services nat mappings pool-name on page 251](#)
[show services nat mappings summary on page 251](#)
[show services nat mappings address-pooling-paired on page 252](#)
[show services nat mappings address-pooling-paired \(mapping of active B4 for a subscriber\) on page 252](#)
[show services nat mappings endpoint-independent on page 252](#)
[show services nat mappings pcg on page 252](#)

Output Fields [Table 15 on page 250](#) lists the output fields for the **show services nat mappings** command. Output fields are listed in the approximate order in which they appear.

Table 15: show services nat mappings Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a service interface.	All levels
Service set	Name of a service set. Individual empty service sets are not displayed, but if none of the service sets has any flows, a flow table header is printed for each service set.	All levels
NAT pool	Name of the NAT pool.	All levels
Address Mapping or Mapping	Mapping performed by NAT to conceal the network address.	All levels
No. of port mappings	Number of port mappings.	All levels
Port mapping	Port mapping performed by NAT.	detail
Flow Count	Number of flows.	detail
Total number of address mappings	Total number of address mappings, by service interface.	summary
Total number of endpoint independent port mappings:	Total number of port mappings by service interface.	summary
Total number of endpoint independent filters	Total number of independent filters that filter out only packets that are not destined to the internal address and port, regardless of the external IP address and port source, by service interface.	summary

Table 15: show services nat mappings Output Fields (*continued*)

Field Name	Field Description	Level of Output
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 <i>pool-name</i>] hierarchy level, lapses, the mapping is deleted. 	
Ports In Use	The number of ports used for a specific address-pooling paired mapping.	
PCP Lifetime	Elapsed PCP lifetime in seconds.	
PCP Client	Address of the PCP client sending the PCP request.	
Session Count	Number of sessions currently using the mapping.	

Sample Output

show services nat mappings brief

```

user@host> show services nat mappings brief
Interface: sp-2/3/0, Service set: s1

NAT pool: p1
Address Mapping: 2.1.20.10 ---> 34.34.34.34
No. of port mappings: 1

```

show services nat mapping detail

```

user@host> show services nat mapping detail
Interface: sp-2/3/0, Service set: s1

NAT pool: p1
Address Mapping: 2.1.20.10 ---> 34.34.34.34, No. of port mappings: 1
Port mapping: 49604 --> 1024, Flow Count: 2

```

show services nat mappings pool-name

```

user@host> show services nat mappings pool-name p1
Interface: sp-2/3/0, Service set: s1

NAT pool: p1
Address Mapping: 2.1.20.10 ---> 34.34.34.34
No. of port mappings: 1

```

show services nat mappings summary

```

user@host> show services nat mapping summary
Service Interface:                               sp-1/0/0
Total number of address mappings:                 790
Total number of endpoint independent port mappings: 1580
Total number of endpoint independent filters:      1580

Service Interface:                               sp-1/1/0

```

Total number of address mappings:	914
Total number of endpoint independent port mappings:	1828
Total number of endpoint independent filters:	1828
Service Interface: sp-4/0/0	
Total number of address mappings:	688
Total number of endpoint independent port mappings:	1376
Total number of endpoint independent filters:	1376
Service Interface: sp-4/1/0	
Total number of address mappings:	648
Total number of endpoint independent port mappings:	1296
Total number of endpoint independent filters:	1296

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-p1
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-p1
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 statistics

Syntax	show services nat statistics <interface <i>interface</i>>
Release Information	Command introduced in Junos OS Release 13.2.
Description	Display the NAT statistics for the multiservices interfaces present on the broadband gateway.
Options	interface <i>interface</i> —Name of the extension provider interface.
Required Privilege Level	view
List of Sample Output	show services nat statistics on page 258
Output Fields	Table 16 on page 253 lists the output fields for the show services nat statistics command. Output fields are listed in the approximate order in which they appear. Some of these fields are used internally by Juniper's engineers for troubleshooting.

Table 16: show services nat statistics Output Fields

Field Name	Field Description
Interface	Name of the multiservices interface.
Session Statistics	
Total Session Interest events	Total number of Session Interest events.
Total Session Create events	Total number of Session Create events.
Total Session Destroy events	Total number of Session Destroy events.
Total Session Pub Req events	Total number of Session Pub Req events.
Total Session Accepts	Total number of sessions accepted.
Total Session Discards	Total number of sessions discarded.
Total Session Ignores	Total number of sessions ignored.
Session interest thru pub event	Session interest through pub event.

Table 16: show services nat statistics Output Fields (*continued*)

Field Name	Field Description
ALG Session interest	Application-level gateway (ALG) session interest.
ALG Session Create	ALG Session Create
Packet Dst in NAT route	Sessions discarded due to packet destination in the NAT route.
Session Ext Alloc Failures	Session extension allocation failures.
Session Ext Set Failures	Session extension set failures.
Session Created for EIF	Number of sessions created for Endpoint Independent Filtering (EIF).
Session Created for EIM	Number of sessions created for Endpoint Independent Mapping (EIM).
NAT rule lookup failures	Number of NAT rule lookup failures.
NAT Allocation Statistics	
NAT allocation Successes	Number of successful NAT map allocations.
NAT allocation Failures	Number of NAT map allocation failures.
NAT Free Successes	NAT free successes.
NAT Free Failures	NAT free failures.
NAT EIM mapping reused	Number of NAT EIM mappings reused.
NAT EIM mapping allocation failures	Number of NAT EIM mapping allocation failures.
NAT EIM mapping Duplicate entry	Number of duplicate NAT EIM mappings.
NAT EIM mapping create failed	Number of failed NAT EIM mappings.
NAT EIM mapping Created	Number of NAT EIM mappings created.

Table 16: show services nat statistics Output Fields (*continued*)

Field Name	Field Description
NAT EIF mapping Free	Number of free NAT EIF mappings.
NAT EIM mapping Free	Number of free NAT EIM mappings.
NAT EIM mapping updated	Time of last EIM update.
NAT EIM waiting for init	Number of NAT EIM mappings waiting for initialization.
NAT EIM waiting for init failed	Number of NAT EIM mappings that failed initialization.
NAT EIM lookup and hold success	Number of successful NAT EIM lookups and holds.
NAT EIM lookup entry in timeout	NAT EIM lookup entry in timeout.
NAT EIM lookup timer cleared for timeout entry	NAT EIM lookup timer cleared for timeout entry.
NAT EIM lookup timeout entry without timer	NAT EIM lookup timeout entry without timer.
NAT EIM release without entry	NAT EIM release without entry.
NAT EIM release entry in timeout	NAT EIM release entry in timeout.
NAT EIM release race	NAT EIM release race.
NAT EIM release set entry for timeout	NAT EIM release set entry for timeout.
NAT EIM timer entry updated	Time of NAT EIM timer update.
NAT EIM timer invalid timer started	NAT EIM timer invalid timer started.
NAT EIM timer entry freed	NAT EIM timer entry freed.

Table 16: show services nat statistics Output Fields (*continued*)

Field Name	Field Description
Packet Statistics	
Total Packets Processed	Total number of packets processed.
Total Packets Forwarded	Total number of packets forwarded.
Total Packets Discarded	Total number of packets discarded.
Total Packets Translated	Total number of packets translated.
Total Packets Restored	Total number of packets restored.
Translation Statistics	
Src IPv4 Translations	Number of source IPv4 translations.
Src IPv4 Restorations	Number of source IPv4 restorations.
Dst IPv4 Translations	Number of destination IPv4 translations.
Dst IPv4 Restorations	Number of destination IPv4 restorations.
Src Port Translations	Number of source port translations.
Src Port Restorations	Number of source port restorations.
Dst Port Translations	Number of destination port translations.
Dst Port Restorations	Number of destination port restorations.
ICMP ID Translations	Number of Internet Control Message Protocol (ICMP) translations.
ICMP ID Restorations	Number of ICMP restorations.
ICMP Error Translations	Number of ICMP error packets after translations.

Table 16: show services nat statistics Output Fields (*continued*)

Field Name	Field Description
TCP Port Translations	Number of TCP port translations.
TCP Port Restorations	Number of TCP port restorations.
UDP Port Translations	Number of UDP port translations.
UDP Port Restorations	Number of UDP port restorations.
GRE Call ID Translations	Number of generic routing encapsulation (GRE) call ID translations.
GRE Call ID Restorations	Number of GRE call ID restorations.
SRC IP restored in ICMP Error	Source IP restored in ICMP Error.
DST IP restored in ICMP Error	DST IP restored in ICMP Error.
SRC IP translated in ICMP Error	SRC IP translated in ICMP Error.
DST IP translated in ICMP Error	Destination IP translated in ICMP Error.
New SRC IP translated in ICMP Error	New source IP translated in ICMP Error.
Inner SRC IP restored in ICMP Error	Inner source IP restored in ICMP Error.
Inner SRC port restored in ICMP Error	Inner source port restored in ICMP Error.
Inner DST IP restored in ICMP Error	Inner destination IP restored in ICMP Error.
Inner SRC IP Translated in ICMP Error	Inner source IP translated in ICMP Error.

Table 16: show services nat statistics Output Fields (*continued*)

Field Name	Field Description
Inner SRC port Translated in ICMP Error	Inner source port translated in ICMP Error.
Inner DST IP Translated in ICMP Error	Inner destination IP translated in ICMP Error.
Misc Errors	
NAT error - no policy	Number of NAT errors of the no policy type.
NAT error - xlate free called with null ext	Number of NAT errors of the xlate free called with null ext type.
NAT error - ext free failed	Number of NAT errors of the ext free failed type.
NAT error - policy add failed	Number of NAT errors of the policy add failed type.
NAT error - policy delete failed	Number of NAT errors of the policy delete failed type.

Sample Output

show services nat statistics

```

user@host> show services nat statistics
Interface: ms-0/0/0

Session statistics
  Total Session Interest events      :12315
  Total Session Create events       :2
  Total Session Destroy events      :12315
  Total Session Pub Req events      :0
  Total Session Accepts             :12315
  Total Session Discards            :0
  Total Session Ignores             :0
  Session interest thru pub event   :0
  ALG Session interest              :0
  ALG Session Create                :0
  Packet Dst in NAT route          :0
  Session Ext Alloc Failures        :0
  Session Ext Set Failures          :0
  Session Created for EIF           :1
  Session Created for EIM           :12314
  NAT rule lookup failures          :0

NAT Allocation statistics
  NAT allocation Successes          :12314

```

```

NAT allocation Failures           :0
NAT Free Successes                :0
NAT Free Failures                 :0
NAT EIM mapping reused            :12312
NAT EIM mapping allocation failures :0
NAT EIM mapping Duplicate entry   :0
NAT EIM mapping create failed     :0
NAT EIM mapping Created           :2
NAT EIF mapping Free              :1
NAT EIM mapping Free              :12314
NAT EIM waiting for init          :0
NAT EIM waiting for init failed   :0
NAT EIM lookup and hold success   :12313
NAT EIM lookup entry in timeout   :0
NAT EIM lookup timer cleared for timeout entry :0
NAT EIM lookup timeout entry without timer :0
NAT EIM release without entry     :0
NAT EIM release entry in timeout  :0
NAT EIM release race              :0
NAT EIM release set entry for timeout :2
NAT EIM timer entry refreshed     :0
NAT EIM timer invalid timer started :2
NAT EIM timer entry freed         :2

Packet statistics
  Total Packets Processed          :2715735062
  Total Packets Forwarded          :2715735062
  Total Packets Discarded          :0
  Total Packets Translated         :1818000836
  Total Packets Restored           :897734226

Translation statistics
  Src IPv4 Translations            :400996
  Src IPv4 Restorations            :897734226
  Dst IPv4 Translations            :1817599840
  Dst IPv4 Restorations            :0
  Src Port Translations            :400996
  Src Port Restorations            :897734226
  Dst Port Translations            :1817599840
  Dst Port Restorations            :0
  ICMP ID Translations             :0
  ICMP ID Restorations             :0
  ICMP Error Translations           :0
  TCP Port Translations             :0
  TCP Port Restorations            :0
  UDP Port Translations             :1818000836
  UDP Port Restorations            :897734226
  GRE CallID Translations           :0
  GRE CallID Restorations           :0
  SRC IP restored in ICMP Error     :0
  DST IP restored in ICMP Error     :0
  SRC IP translated in ICMP Error   :0
  DST IP translated in ICMP Error   :0
  New SRC IP translated in ICMP Error :0
  Inner SRC IP restored in ICMP Error :0
  Inner SRC port restored in ICMP Error :0
  Inner DST IP restored in ICMP Error :0
  Inner SRC IP Translated in ICMP Error :0
  Inner SRC port Translated in ICMP Error :0
  Inner DST IP Translated in ICMP Error :0

```

```
Misc Errors
  NAT error - no policy                                :0
  NAT error - xlate free called with null ext          :0
  NAT error - ext free failed                          :0
  NAT error - policy add failed                        :0
  NAT error - policy delete failed                     :0

Interface: ms-1/1/0

Session statistics
  Total Session Interest events                        :6
  Total Session Create events                         :6
  Total Session Destroy events                        :7
  Total Session Pub Req events                       :0
  Total Session Accepts                              :6
  Total Session Discards                             :0
  Total Session Ignores                              :0
  Session interest thru pub event                     :0
  ALG Session interest                               :0
  ALG Session Create                                 :0
  Packet Dst in NAT route                            :0
  Session Ext Alloc Failures                          :0
  Session Ext Set Failures                           :0
  Session Created for EIF                             :0
  Session Created for EIM                             :6
  NAT rule lookup failures                            :0

NAT Allocation statistics
  NAT allocation Successes                            :6
  NAT allocation Failures                             :0
  NAT Free Successes                                  :0
  NAT Free Failures                                   :0
  NAT EIM mapping reused                              :3
  NAT EIM mapping allocation failures                  :0
  NAT EIM mapping Duplicate entry                     :0
  NAT EIM mapping create failed                       :0
  NAT EIM mapping Created                             :3
  NAT EIF mapping Free                                :0
  NAT EIM mapping Free                                :7
  NAT EIM waiting for init                            :0
  NAT EIM waiting for init failed                     :0
  NAT EIM lookup and hold success                     :2
  NAT EIM lookup entry in timeout                     :1
  NAT EIM lookup timer cleared for timeout entry      :1
  NAT EIM lookup timeout entry without timer          :0
  NAT EIM release without entry                       :0
  NAT EIM release entry in timeout                    :0
  NAT EIM release race                                :0
  NAT EIM release set entry for timeout                :5
  NAT EIM timer entry refreshed                       :0
  NAT EIM timer invalid timer started                 :4
  NAT EIM timer entry freed                           :4

Packet statistics
  Total Packets Processed                             :2733886586
  Total Packets Forwarded                             :2733886586
  Total Packets Discarded                             :0
  Total Packets Translated                            :1836152360
  Total Packets Restored                              :897734226

Translation statistics
```

```

Src IPv4 Translations :1836152360
Src IPv4 Restorations :0
Dst IPv4 Translations :0
Dst IPv4 Restorations :897734226
Src Port Translations :1836152360
Src Port Restorations :0
Dst Port Translations :0
Dst Port Restorations :897734226
ICMP ID Translations :0
ICMP ID Restorations :0
ICMP Error Translations :0
TCP Port Translations :0
TCP Port Restorations :0
UDP Port Translations :1836152360
UDP Port Restorations :897734226
GRE CallID Translations :0
GRE CallID Restorations :0
SRC IP restored in ICMP Error :0
DST IP restored in ICMP Error :0
SRC IP translated in ICMP Error :0
DST IP translated in ICMP Error :0
New SRC IP translated in ICMP Error :0
Inner SRC IP restored in ICMP Error :0
Inner SRC port restored in ICMP Error :0
Inner DST IP restored in ICMP Error :0
Inner SRC IP Translated in ICMP Error :0
Inner SRC port Translated in ICMP Error :0
Inner DST IP Translated in ICMP Error :0

Misc Errors
NAT error - no policy :0
NAT error - xlate free called with null ext :0
NAT error - ext free failed :0
NAT error - policy add failed :0
NAT error - policy delete failed :0

```

show services pcp statistics

Syntax	show services pcp statistics
Release Information	Command introduced in Junos OS Release 13.2
Description	Display information PCP mappings.
Required Privilege Level	view
List of Sample Output	show services pcp statistics pcp on page 263
Output Fields	Table 17 on page 262 lists the output fields for the show services pcp statistics command. Output fields are listed in the approximate order in which they appear.

Table 17: show services pcp 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.
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.

Table 17: show services pcp statistics Output Fields (*continued*)

Field Name	Field Description
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.
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 pcp statistics pcp

```
user@host> show services pcp statistics pcp
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 software

Syntax	show services software <count>
Release Information	Command introduced in Junos OS Release 10.4. <count> option added in Junos OS Release 11.2.
Description	Display information about software services. Information is displayed on both 6rd and DS-Lite services.
Options	count <i>interface-name</i> — (Optional) Display the current software counts for a service set for both DS-Lite and 6rd.
Required Privilege Level	view
List of Sample Output	show services software on page 265 show services software count on page 265
Output Fields	Table 18 on page 265 lists the output fields for the command-name command. Output fields are listed in the approximate order in which they appear.

Table 18: 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
10.10.10.2    ->    30.30.30.1    Direction    Flow count
                                I                                13

```

show services software count

```

user@host> show services software count
Interface    Service set    DS-Lite    6RD
sp-0/0/0    dslite-svc-set1    2          0

```

show services software flows

Syntax	<pre>show services software flows (<interface <i>interface-name</i>> <service-set <i>service-set-name</i>> count <interface <i>interface-name</i>> <service-set <i>service-set-name</i>> ds-lite <B4 <i>b4-address</i>> <AFTR <i>aftr-address</i>> v6rd <initiator <i>initiator-ip-address</i>><concentrator <i>concentrator-ip-address</i>>)</pre>
Release Information	Command introduced in Junos OS Release 10.2.
Description	Display statistics information about the software flows.
Options	<p>interface <i>interface-name</i>—(Optional) Display statistics information about the specified interface only.</p> <p>service-set <i>service-set-name</i>—(Optional) Display statistics information about the specified service set only.</p> <p>count <interface <i>interface-name</i>> <service-set <i>service-set-name</i>> —(Optional) Display flow count information only, with optional filtering by interface and service set.</p> <p>ds-lite <B4 <i>b4-address</i>> <AFTR <i>aftr-address</i>> —(Optional) Display DS-Lite flow information, with optional filtering by B4 (software initiator) and AFTR (software concentrator).</p> <p>v6rd <initiator <i>initiator-ip-address</i>><concentrator <i>concentrator-ip-address</i>>)—(Optional) Display v6rd flow information, with optional filtering by the software initiator and software concentrator.</p>
Required Privilege Level	view
List of Sample Output	<p>show services software flows on page 267</p> <p>show services software flows count on page 267</p> <p>show services software flows ds-lite B4 on page 267</p> <p>show services software flows ds-lite AFTR on page 268</p> <p>services software flows ds-lite AFTR and B4 on page 268</p>
Output Fields	<p>Table 19 on page 266 lists the output fields for the show services software flows command. Output fields are listed in the approximate order in which they appear.</p>

Table 19: 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.

Table 19: show services software flows Output Fields (*continued*)

Field Name	Field Description
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 software statistics

Syntax	<pre>show services software statistics <ds-lite> <ds-lite> <interface interface-name> <v6rd></pre>
Release Information	Command introduced in Junos OS Release 10.4.
Description	Display information about software services.
Options	<p>ds-lite—(Optional) Display only DS-Lite.</p> <p>interface interface-name —(Optional) Name of the interface servicing the software. When you omit this option, data for all interfaces are shown.</p> <p>v6rd—(Optional) Display only 6rd statistics.</p>
Required Privilege Level	view
List of Sample Output	show services software statistics on page 272 show services software statistics ds-lite on page 273
Output Fields	Table 20 on page 269 lists the output fields for the command-name command. Output fields are listed in the approximate order in which they appear.

Table 20: command-name Output Fields

Field Name	Field Description	Level of Output
Service PIC Name	Name of service PIC for which statistics are shown.	statistics
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 software 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

Table 20: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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

Table 20: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
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
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 Softwire ID	Number of times a softwire 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
Packet not IPv6-in-IPv4	Number of IPv6 packets not encapsulated in IPv4.	statistics for v6rd only

Table 20: command-name Output Fields (*continued*)

Field Name	Field Description	Level of Output
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

```
user@host> show services software statistics
DS-Lite Statistics:
```

```
Service PIC Name:                               :sp-0/0/0
```

Statistics

```
-----
```

```

Software Created                               :0
Software Created for EIF/HP                     :0
Software Deleted                               :0
Software Flows Created                         :0
Software 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
Software 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-liteuser@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
Software 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

Software Creation Failed	:0
Software 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
Software 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

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 278](#)
[show services stateful-firewall conversations destination-port on page 278](#)

Output Fields Table 21 on page 277 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 21: 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.
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.

Table 21: show services stateful-firewall conversations Output Fields (*continued*)

Field Name	Field Description
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.
Timeout	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    0      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     0      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 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
- **iioip**—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*

List of Sample Output [show services stateful-firewall flows on page 282](#)
[show services stateful-firewall flows \(For Software Flows\) on page 282](#)
[show services stateful-firewall flows brief on page 283](#)
[show services stateful-firewall flows extensive on page 283](#)
[show services stateful-firewall flows count on page 283](#)
[show services stateful-firewall flows destination port on page 283](#)
[show services stateful-firewall flows source port on page 283](#)
[show services stateful-firewall flows \(Twice NAT\) on page 283](#)

Output Fields [Table 22 on page 281](#) 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 22: 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.

Table 22: show services stateful-firewall flows Output Fields (*continued*)

Field Name	Field Description
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).
Frm count	Number of frames in the flow.

Sample Output

show services stateful-firewall flows

```
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  0
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@router> show services stateful-firewall flows destination-port 21
Interface: ms-0/3/0, Service set: svc_set_trust
Flow
Interface: ms-0/3/0, Service set: svc_set_untrust
Flow
TCP      10.50.10.2:2143  ->    10.50.20.2:21      Watch   Dir  Frm count
                                State   Dir  Frm count
                                0      0

```

show services stateful-firewall flows source port

```

user@router> show services stateful-firewall flows source-port 2143
Interface: ms-0/3/0, Service set: svc_set_trust
Flow
Interface: ms-0/3/0, Service set: svc_set_untrust
Flow
TCP      10.50.10.2:2143  ->    10.50.20.2:21      Watch   Dir  Frm count
                                State   Dir  Frm count
                                0      0

```

show services stateful-firewall flows (Twice NAT)

```

user@router> 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 statistics

Syntax	<pre>show services stateful-firewall statistics <application-protocol <i>protocol</i>> <brief detail extensive summary> <interface <i>interface-name</i>> <service-set <i>service-set</i>></pre>
Release Information	Command introduced before Junos OS Release 7.4.
Description	Display stateful firewall statistics.
Options	<p>none—Display standard information about all stateful firewall statistics.</p> <p>brief detail extensive summary—(Optional) Display the specified level of output.</p> <p>interface <i>interface-name</i>—(Optional) Display information about a particular interface. On M Series and T Series routers, the <i>interface-name</i> can be <i>ms-fpc/pic/port</i> or <i>rspnumber</i>.</p> <p>service-set <i>service-set</i>—(Optional) Display information about a particular service set.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> <i>clear services stateful-firewall statistics</i>
List of Sample Output	show services stateful-firewall statistics extensive on page 292
Output Fields	Table 23 on page 285 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 23: 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.

Table 23: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
Hairpinning Counters	<p>Hairpinning counters:</p> <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	<p>Drop counters:</p> <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	<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. • Non-IP packets—Total non-IPv4 errors. • ALG—Total application-level gateway (ALG) errors

Table 23: 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 23: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
TCP Errors	

Table 23: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	TCP protocol errors:
	<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. • 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

Table 23: show services stateful-firewall statistics Output Fields (*continued*)

Field Name	Field Description
	<p>RST is received either from the client or server with a non-matching sequence number.</p> <ul style="list-style-type: none"> • 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. • 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 23: 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 • IIOB—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
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
Haripinning 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
```


PART 4

Index

- [Index on page 297](#)

Index

Symbols

#, comments in configuration statements.....	xvi
(), in syntax descriptions.....	xvi
6rd flows	
statistics.....	266
< >, in syntax descriptions.....	xvi
[], in configuration statements.....	xvi
{ }, in configuration statements.....	xvi
(pipe), in syntax descriptions.....	xvi

A

address statement	
NAT.....	122
address-allocation statement.....	123
address-range statement	
NAT.....	123
ALGs	
default.....	12
allow-overlapping-nat-pools statement.....	124
app-mapping-timeout statement.....	124
application-sets statement	
NAT.....	125
applications statement	
NAT.....	125

B

basic-nat-pt option	
configuring.....	54
basic-nat44 option	
configuring.....	25
basic-nat66 option	
configuring.....	31
braces, in configuration statements.....	xvi
brackets	
angle, in syntax descriptions.....	xvi
square, in configuration statements.....	xvi

C

cgn-pic statement.....	126
CGNAT	
ALGs.....	12

clear services inline nat pool command.....	230
clear services inline nat statistics command.....	231
clear services nat flows command.....	232
clear services nat mappings app	
command.....	235, 236, 238
clear services nat mappings command.....	233
clear services nat statistics command.....	240
comments, in configuration statements.....	xvi
configuring NAT-PT with DNS application-level	
gateways	
example.....	82
conventions	
text and syntax.....	xv
curly braces, in configuration statements.....	xvi
customer support.....	xvii
contacting JTAC.....	xvii

D

destination NAT	
configuring.....	35, 65, 68
destination-address statement	
NAT.....	126
destination-address-range statement	
NAT.....	127
destination-pool statement.....	127
destination-port range statement	
NAT.....	128
destination-prefix statement.....	128
destination-prefix-list statement	
NAT.....	129
destined-port statement	
NAT.....	129
deterministic-port-block-allocation	
statement.....	130
dnat-44 option	
usage guidelines.....	35, 65, 68
documentation	
comments on.....	xvii
ds-lite	
subnet session limitation	
configuring.....	226
DS-Lite flows	
statistics.....	266
ds-lite statement.....	188
usage guidelines.....	159
dynamic address-only source translation	
configuring.....	39
dynamic NAT	
configuring.....	39

dynamic-nat44 option		
usage guidelines.....	39	
E		
ei-mapping-timeout statement.....	132	
F		
font conventions.....	xv	
from statement		
NAT.....	133	
I		
inline NAT		
statistics, displaying.....	241, 242	
IPv4		
napt-44 option.....	43	
translation type		
basic-nat-pt option.....	54	
basic-nat44 option.....	25	
basic-nat66 option.....	31	
IPv4 dynamic source translation		
configuring.....	43	
IPv6		
napt-66 option.....	47	
IPv6 dynamic source translation		
configuring.....	47	
ipv6-multicast-interfaces statement.....	134	
M		
manuals		
comments on.....	xvii	
mapping-timeout statement.....	135	
match-direction statement		
NAT.....	135	
MS-MPC		
configuration example		
napt.....	77	
N		
NAPT		
configuring.....	43, 47	
IPv4.....	43	
IPv6.....	47	
napt		
configuration example.....	77	
napt-44 option		
usage guidelines.....	43	
napt-66 option		
usage guidelines.....	47	
napt-pt option		
example.....	82	
NAT		
ALGs.....	12	
destination NAT.....	35, 65	
dynamic address-only source translation.....	39	
dynamic NAT.....	39	
dynamic source translation.....	43, 47	
inline		
configuring.....	96	
inter-chassis high availability.....	207	
ipv6-multicast-interfaces information,		
displaying.....	243	
mapping information, address-pooling		
paired.....	249	
mapping information, displaying.....	249	
mapping information, endpoint-independent		
.....	249	
mapping information, pcp.....	249	
NAT-PT example.....	82	
session logging.....	204	
static destination address translation.....	35, 65	
status information, displaying.....	245	
nat		
flows		
clearing.....	232	
mappings		
clearing.....	233, 235, 236, 238	
network address translation		
configuration example		
napt.....	77	
network address translation See NAT		
no-translation statement.....	136	
O		
overload-pool statement.....	136	
overload-prefix statement.....	137	
P		
parentheses, in syntax descriptions.....	xvi	
pool statement.....	138	
port block allocation		
deterministic		
configuring.....	51	
interim syslog messages.....	204	
secured		
configuring.....	49	

Port Control Protocol	
Configuring.....	71
Configuring a Service Set to Apply PCP.....	73
Configuring PCP Server Options.....	71, 72
port forwarding	
configuring.....	68
dnat-44.....	65
static destination address translation.....	65
without destination address translation.....	68
port forwarding without static destination address translation	
configuring.....	68
port statement	
NAT.....	139
port-forwarding	
example.....	69
port-forwarding statement	
destined-port statement.....	129
NAT.....	140
translated-port statement.....	152
port-forwarding-mappings statement.....	140
ports-per-session statement.....	141
R	
random-allocation statement.....	139
rule statement	
NAT.....	142
software.....	161, 189
rule-set statement	
NAT.....	143
software.....	189
S	
secure-nat-mapping statement.....	143
secured-port-block-allocation statement.....	144
server (PCP) statement.....	145
service-set statement.....	147
services statement	
NAT.....	146
session logging.....	204
show services inline nat pool command.....	241
show services inline nat statistics command.....	242
show services nat ipv6-multicast-interfaces command.....	243
show services nat mappings command.....	249
show services nat pool command.....	245
show services nat statistics command.....	253
show services pcg statistics command.....	262
show services software command.....	265
show services software flows command.....	266
show services software statistics command.....	269
show services stateful-firewall conversations command.....	275
show services stateful-firewall flows command.....	279
show services stateful-firewall statistics command.....	285
software flows	
statistics.....	266
software-concentrator statement.....	190
software-rules statement.....	191
source-address statement	
NAT.....	149
source-address-range statement	
NAT.....	149
source-pool statement.....	150
source-prefix statement.....	150
source-prefix-list statement	
NAT.....	151
stateful firewall	
conversations	
displaying.....	275
flows	
displaying.....	279
statistics	
displaying.....	285
stateful NAT64	
configuring.....	52
static destination address translation	
configuring.....	35, 65
subnet session limitation	
ds-lite	
configuring.....	226
support, technical See technical support	
syntax conventions.....	xv
syslog statement	
NAT.....	151
T	
technical support	
contacting JTAC.....	xvii
term statement	
NAT.....	153
then statement	
NAT.....	154
translated statement.....	155
translated-port statement	
NAT.....	152

translation-type statement.....	156
basic-nat-pt option.....	54
basic-nat44 option.....	25
basic-nat66 option.....	31
dnat-44 option, configuring.....	35, 65
dynamic-nat44, configuring.....	39
napt-44 option, configuring.....	43
napt-66 option, configuring.....	47
napt-pt option, example.....	82
stateful-nat64 option, configuring.....	52
twice-napt-44 option	
example.....	69
 V	
v6rd statement.....	192
usage guidelines.....	160