

Junos® OS

Network Address Translation User Guide

Published
2023-12-13

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

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

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

Junos® OS Network Address Translation User Guide

Copyright © 2023 Juniper Networks, Inc. All rights reserved.

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

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

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

Table of Contents

[About This Guide | xi](#)

1

Overview

[NAT Overview | 2](#)

[Introduction to NAT | 2](#)

[Understanding NAT Rule Sets and Rules | 3](#)

[NAT Configuration Overview | 8](#)

[Configuring NAT Using the NAT Wizard | 8](#)

[Example: Configuring NAT for Multiple ISPs | 9](#)

[Requirements | 9](#)

[Overview | 9](#)

[Configuration | 9](#)

[Verification | 24](#)

[Configuring Proxy ARP for NAT \(CLI Procedure\) | 24](#)

[Configuring NAT trace options | 25](#)

[Monitoring NAT Incoming Table Information | 27](#)

[Monitoring Interface NAT Port Information | 29](#)

2

Types of NAT

[Source NAT | 32](#)

[Understanding Source NAT | 33](#)

[Understanding Central Point Architecture Enhancements for NAT | 34](#)

[Port Overflow Burst Mode | 35](#)

[Optimizing Source NAT Performance | 36](#)

[Monitoring Source NAT Information | 39](#)

[Source NAT Configuration Overview | 48](#)

[Example: Configuring Source NAT for Egress Interface Translation | 48](#)

Requirements | 49

Overview | 49

Configuration | 51

Verification | 53

Example: Configuring Source NAT for Single Address Translation | 54

Requirements | 54

Overview | 55

Configuration | 57

Verification | 60

Example: Configuring Source and Destination NAT Translations | 61

Requirements | 61

Overview | 62

Configuration | 64

Verification | 69

Understanding Source NAT Rules | 71

Example: Configuring Source NAT with Multiple Rules | 72

Requirements | 72

Overview | 72

Configuration | 75

Verification | 80

Understanding Source NAT Pools | 81

Understanding Source NAT Pool Capacities | 83

Understanding Persistent Addresses for Source NAT Pools | 84

Example: Configuring Capacity for Source NAT Pools with PAT | 85

Requirements | 85

Overview | 85

Configuration | 85

Verification | 87

Understanding Source NAT Pools with Address Pooling | 88

Understanding Source NAT Pools with Address Shifting | 88

Example: Configuring Source NAT Pools with Address Shifting | 89

Requirements | 89

Overview | 90

Configuration | 92

Verification | 95

Understanding Source NAT Pools with PAT | 96

Example: Configuring Source NAT for Multiple Addresses with PAT | 97

Requirements | 97

Overview | 98

Configuration | 100

Verification | 103

Understanding Source NAT Pools Without PAT | 104

Example: Configuring a Single IP Address in a Source NAT Pool Without PAT | 105

Requirements | 106

Overview | 106

Configuration | 106

Verification | 109

Example: Configuring Multiple Addresses in a Source NAT Pool Without PAT | 110

Requirements | 110

Overview | 111

Configuration | 113

Verification | 116

Understanding Shared Addresses in Source NAT Pools without PAT | 117

Understanding NAT Session Persistence | 118

Configure Port Block Allocation Size | 120

Configuring the NAT Session Hold Timeout and NAT Session Persistence Scan | 122

Understanding NAT Configuration Check on Egress Interfaces after Reroute | 123

Destination NAT | 125

Understanding Destination NAT | 125

Understanding Destination NAT Address Pools | 126

Understanding Destination NAT Rules | 127

Destination NAT Configuration Overview | 127

Example: Configuring Destination NAT for Single Address Translation | 128

Requirements | 128

Overview | 128

Configuration | 131

Verification | 135

Example: Configuring Destination NAT for IP Address and Port Translation | 139

Requirements | 139

Overview | 139

Configuration | 141

Verification | 145

Example: Configuring Destination NAT for Subnet Translation | 146

Requirements | 146

Overview | 147

Configuration | 149

Verification | 152

Monitoring Destination NAT Information | 153

Static NAT | 157

Understanding Static NAT | 157

Understanding Static NAT Rules | 158

Static NAT Configuration Overview | 159

Example: Configuring Static NAT for Single Address Translation | 159

Requirements | 159

Overview | 159

Configuration | 161

Verification | 164

Example: Configuring Static NAT for Subnet Translation | 165

Requirements | 166

Overview | 166

Configuration | 168

Verification | 171

Example: Configuring Static NAT for Port Mapping | 172

Requirements | 173

Overview | 173

Configuration | 175

Verification | 179

Troubleshooting | 180

Monitoring Static NAT Information | 181

3

NAT Configuration Options

Persistent NAT and NAT64 | 186

Understanding Persistent NAT and NAT64 | 186

Understanding Session Traversal Utilities for NAT (STUN) Protocol | 188

Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation | 189

Persistent NAT and NAT64 Configuration Overview | 191

Example: Configuring Address Persistent NAT64 Pools | 193

Requirements | 193

Overview | 193

Configuration | 193

Verification | 196

Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT | 196

Requirements | 197

Overview | 197

Configuration | 199

Verification | 203

Example: Configuring Address-Dependent Filtering for IPv6 Clients | 205

Requirements | 205

Overview | 205

Configuration | 206

Verification | 209

Example: Configuring Endpoint-Independent Filtering for IPv6 Clients | 210

Requirements | 210

- Overview | 210
- Configuration | 211
- Verification | 215

Example: Setting Maximum Persistent NAT Bindings | 216

- Requirements | 216
- Overview | 216
- Configuration | 217
- Verification | 218

Persistent NAT Hairpinning Overview | 218

Example: Configuring Persistent NAT Hairpinning with Source NAT Pool with Address Shifting | 220

- Requirements | 220
- Overview | 221
- Configuration | 223
- Verification | 226

NAT for Multicast Flows | 228

Understanding NAT for Multicast Flows | 228

Example: Configuring NAT for Multicast Flows | 229

- Requirements | 229
- Overview | 230
- Configuration | 232
- Verification | 240

IPv6 NAT | 242

IPv6 NAT Overview | 242

IPv6 NAT PT Overview | 245

IPv6 NAT-PT Communication Overview | 246

Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination
Address Prefix Static Mapping | 247

- Requirements | 247
- Overview | 247
- Configuration | 248
- Verification | 251

Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination
Address One-to-One Mapping | 252

Requirements | 252

Overview | 252

Configuration | 253

Verification | 256

Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination
Address Prefix Static Mapping | 257

Requirements | 257

Overview | 257

Configuration | 258

Verification | 261

Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination
Address One-to-One Mapping | 263

Requirements | 263

Overview | 263

Configuration | 264

Verification | 267

IPv6 Dual-Stack Lite | 268

Understanding IPv6 Dual-Stack Lite | 268

Example: Configuring IPv6 Dual-Stack Lite | 271

Requirements | 272

Overview | 272

Configuration | 272

Verification | 273

NAT for VRF Routing Instance | 274

NAT Overview | 275

Example: Configuring Source NAT to convert the private IP address of a VRF instance to the
private IP address of another VRF instance | 275

Requirements | 275

Overview | 276

Configuration | 276

Example: Configuring Destination NAT to Convert Public IP Address to VRF's Single Private IP Address of a VRF instance | **283**

Requirements | **283**

Overview | **283**

Configuration | **284**

Verification | **289**

Example: Configuring Static NAT to Convert the Private IP Address of a VRF Instance to Public IP Address | **290**

Requirements | **290**

Overview | **290**

Configuration | **291**

Verification | **295**

NAT for VRF group | 297

Overview | **297**

Example: Configuring Source NAT to convert the private IP address of a VRF Group to the private IP address of different VRF instance | **297**

Requirements | **298**

Overview | **298**

Configuration | **298**

Example: Configuring Destination NAT to Convert Public IP Address of a VRF Group to the private IP address of different VRF instance | **303**

Requirements | **304**

Overview | **304**

Configuration | **305**

Verification | **308**

Configuration Statements and Operational Commands

Junos CLI Reference Overview | 311

About This Guide

Use this guide to configure Network Address Translation (NAT) functionality for translating IP addresses in Junos OS on NFX Series and SRX Series Firewalls.

1

CHAPTER

Overview

[NAT Overview | 2](#)

[NAT Configuration Overview | 8](#)

NAT Overview

IN THIS SECTION

- [Introduction to NAT | 2](#)
- [Understanding NAT Rule Sets and Rules | 3](#)

Network Address Translation (NAT) is a mechanism to translate the IP address of a computer or group of computers into a single public address when the packets are sent out to the internet. By translating the IP address, only one IP address is publicized to the outside network. Since only one IP address is visible to the outside world, NAT provides additional security and it can have only one public address for the entire network instead of having multiple IP addresses.

Introduction to NAT

Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. NAT can include the translation of port numbers as well as IP addresses.

NAT is described in RFC 1631 to solve IP (version 4) address depletion problems. Since then, NAT has been found to be a useful tool for firewalls, traffic redirect, load sharing, network migrations, and so on.

The following types of NAT are supported on Juniper Networks devices:

- Static NAT
- Destination NAT
- Source NAT

NOTE: SRX Series Firewalls perform both policy lookup and service lookup based on the translated destination port.

You can use the NAT Wizard to perform basic NAT configuration. To perform more advanced configuration, use the J-Web interface or the CLI.

Starting from Junos OS Release 19.3R1, SRX5000 line devices with SRX5K-SPC3 card, SRX4100, SRX4200, and vSRX Virtual Firewall instances support NAT features such as source NAT, destination NAT, and static NAT for both IPv4 and IPv6 traffic in PowerMode IPsec (PMI) mode. NAT64 is not supported in PMI mode. However, NAT64 works properly in normal mode, when PMI is enabled.

SEE ALSO

[Source NAT | 32](#)

[Destination NAT | 125](#)

[Static NAT | 157](#)

Understanding NAT Rule Sets and Rules

IN THIS SECTION

- [NAT Rule Sets | 3](#)
- [NAT Rules | 4](#)
- [Rule Processing | 5](#)
- [NAT Rule Capacity | 6](#)

NAT processing centers on the evaluation of NAT rule sets and rules. A rule set determines the overall direction of the traffic to be processed. For example, a rule set can select traffic from a particular interface or to a specific zone. A rule set can contain multiple rules. Once a rule set is found that matches specific traffic, each rule in the rule set is evaluated for a match. Each rule in the rule set further specifies the traffic to be matched and the action to be taken when traffic matches the rule.

This topic includes the following sections:

NAT Rule Sets

A rule set specifies a general set of matching conditions for traffic. For static NAT and destination NAT, a rule set specifies one of the following:

- Source interface
- Source zone

- Source routing instance

For source NAT rule sets, you configure both source and destination conditions:

- Source interface, zone, or routing instance
- Destination interface, zone, or routing instance

It is possible for a packet to match more than one rule set; in this case, the rule set with the more specific match is used. An interface match is considered more specific than a zone match, which is more specific than a routing instance match. If a packet matches both a destination NAT rule set that specifies a source zone and a destination NAT rule set that specifies a source interface, the rule set that specifies the source interface is the more specific match.

Source NAT rule set matching is more complex because you specify both source and destination conditions in a source NAT rule set. In the case where a packet matches more than one source NAT rule set, the rule set chosen is based on the following source/destination conditions (in order of priority):

1. Source interface/destination interface
2. Source zone/destination interface
3. Source routing instance/destination interface
4. Source interface/destination zone
5. Source zone/destination zone
6. Source routing instance/destination zone
7. Source interface/destination routing instance
8. Source zone/destination routing instance
9. Source routing instance/destination routing instance

For example, you can configure rule set A, which specifies a source interface and a destination zone, and rule set B, which specifies a source zone and a destination interface. If a packet matches both rule sets, rule set B is the more specific match.

NOTE: You cannot specify the same source and destination conditions for source NAT rule sets.

NAT Rules

Once a rule set that matches the traffic has been found, each rule in the rule set is evaluated in order for a match. NAT rules can match on the following packet information:

- Source and destination address
- Source port (for source and static NAT only)
- Destination port

The first rule in the rule set that matches the traffic is used. If a packet matches a rule in a rule set during session establishment, traffic is processed according to the action specified by that rule.

You can use the **show security nat source rule** and **show security nat destination rule** and the **show security nat static rule** commands to view the number of sessions for a specific rule.


Rule Processing

The NAT type determines the order in which NAT rules are processed. During the first packet processing for a flow, NAT rules are applied in the following order:

1. Static NAT rules
2. Destination NAT rules
3. Route lookup
4. Security policy lookup
5. Reverse mapping of static NAT rules
6. Source NAT rules

[Figure 1 on page 6](#) illustrates the order for NAT rule processing.

Figure 1: NAT Rule Processing



Static NAT and destination NAT rules are processed before route and security policy lookup. Static NAT rules take precedence over destination NAT rules. Reverse mapping of static NAT rules takes place after route and security policy lookup and takes precedence over source NAT rules. Source NAT rules are processed after route and security policy lookup and after reverse mapping of static NAT rules.

The configuration of rules and rule sets is basically the same for each type of NAT—source, destination, or static. But because both destination and static NAT are processed before route lookup, you cannot specify the destination zone, interface or routing instance in the rule set.

NAT Rule Capacity

Table 1 on page 6 provides the NAT rule capacity requirements per device. Platform support depends on the Junos OS release in your installation.

Table 1: Number of Rules on SRX Series Firewall

NAT Rule Type	SRX100	SRX300 SRX320	SRX340 SRX345	SRX1500	SRX4100 SRX4200	SRX4600	SRX5400 SRX5600 SRX5800
Source NAT rule	1024	1024	2048	8192	20,480	51,200	30,720

Table 1: Number of Rules on SRX Series Firewall (Continued)

NAT Rule Type	SRX100	SRX300 SRX320	SRX340 SRX345	SRX1500	SRX4100 SRX4200	SRX4600	SRX5400 SRX5600 SRX5800
Destination NAT rule	1024	1024	2048	8192	20,480	51,200	30,720
Static NAT rule	1024	1024	2048	8192	20,480	51,200	30,720

The restriction on the number of rules per rule set is a device-wide limitation on how many rules a device can support. This restriction is provided to help you better plan and configure the NAT rules for the device.

For memory consumption, there is no guarantee to support these numbers (maximum source rule or rule set + maximum destination rule or rule set + maximum static rule or rule-set) at the same time for SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 devices.

[Table 2 on page 7](#) provides the recommended maximum number of rules and rule sets for SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800 devices. Platform support depends on the Junos OS release in your installation.

Table 2: Number of Rules and Rule Sets

Objects	SRX3400 SRX3600	SRX4600	SRX5400 SRX5600 SRX5800
Total NAT rule sets per system	20,480	51,200	30,720
Total NAT rules per rule set	20,480	51,200	30,720

NAT Configuration Overview

IN THIS SECTION

- [Configuring NAT Using the NAT Wizard | 8](#)
- [Example: Configuring NAT for Multiple ISPs | 9](#)
- [Configuring Proxy ARP for NAT \(CLI Procedure\) | 24](#)
- [Configuring NAT trace options | 25](#)
- [Monitoring NAT Incoming Table Information | 27](#)
- [Monitoring Interface NAT Port Information | 29](#)

This topic describes how to configure Network Address Translation (NAT) and multiple ISPs. Also, this topic helps to verify the NAT traffic by configuring the trace options and monitoring NAT table.

Configuring NAT Using the NAT Wizard

You can use the NAT Wizard to perform basic NAT configuration on SRX300, SRX320, SRX340, SRX345, and SRX550M devices. To perform more advanced configuration, use the J-Web interface or the CLI.

To configure NAT using the NAT Wizard:

1. Select **Configure>Tasks>Configure NAT** in the J-Web interface.
2. Click the **Launch NAT Wizard** button.
3. Follow the wizard prompts.

The upper-left area of the wizard page shows where you are in the configuration process. The lower-left area of the page shows field-sensitive help. When you click a link under the **Resources** heading, the document opens in your browser. If the document opens in a new tab, be sure to close only the tab (not the browser window) when you close the document.

Example: Configuring NAT for Multiple ISPs

IN THIS SECTION

- [Requirements | 9](#)
- [Overview | 9](#)
- [Configuration | 9](#)
- [Verification | 24](#)

This example shows how to configure a Juniper Networks device for address translation of multiple ISPs.

Requirements

Before you begin:

1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

In this example, you can configure an SRX Series Firewall by connecting the LAN to the Internet by using NAT feature through two ISP connections. In this configuration, trust is the security zone for the private address space and the two untrust security zones for the public address space are used to connect from LAN to the two ISPs and vice versa. The example is a combination of source NAT rules to connect to Internet from the LAN, and destination and static NAT rules to connect to the LAN from Internet.

Configuration

IN THIS SECTION

- [Configuring NAT for Multiple ISPs | 10](#)

Configuring NAT for Multiple ISPs

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, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter `commit` from configuration mode.

```
set routing-instances isp1 instance-type virtual-router
set routing-instances isp1 interface ge-0/0/2.0
set routing-instances isp1 routing-options static route 10.0.0.0/8 next-table inet.0
set routing-instances isp1 routing-options static route 0.0.0.0/0 next-hop 192.0.2.20
set routing-instances isp2 instance-type virtual-router
set routing-instances isp2 interface ge-0/0/3.0
set routing-instances isp2 routing-options static route 10.0.0.0/8 next-table inet.0
set routing-instances isp2 routing-options static route 0.0.0.0/0 next-hop 198.51.100.251
set routing-options interface-routes rib-group inet isp
set routing-options static route 10.0.0.0/8 next-hop 10.0.21.254
set routing-options rib-groups isp import-rib inet.0
set routing-options rib-groups isp import-rib isp1.inet.0
set routing-options rib-groups isp import-rib isp2.inet.0
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol match source-address any
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol match destination-address any
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol match application any
set security policies from-zone trust to-zone untrust1 policy tr-untr1-pol then permit
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol match source-address any
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol match destination-address any
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol match application any
set security policies from-zone trust to-zone untrust2 policy tr-untr2-pol then permit
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match source-address any
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match destination-address any
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match application any
set security policies from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol then reject
set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match source-address any
```

```

set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
destination-address any
set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match
application any
set security policies from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol then reject
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match source-address
any
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match destination-
address ftp-ser
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match destination-
address telnet-ser
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match application
junos-ftp
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol match application
junos-telnet
set security policies from-zone untrust1 to-zone trust policy untr1-tr-pol then permit
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match source-address
any
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match destination-
address 10.171.9.23/32
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match destination-
address http-ser
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match destination-
address 10.103.12.0/24
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match application
junos-http
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match application
junos-icmp-all
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol match application
junos-dhcp-server
set security policies from-zone untrust2 to-zone trust policy untr2-tr-pol then permit
set security nat source pool pool_1 address 192.0.2.40/32 to 192.0.2.190/32
set security nat source pool pool_2 address 192.0.2.250/32
set security nat source pool pool_3 address 198.51.100.20/32 to 198.51.100.30/32
set security nat source address-persistent
set security nat source pool-utilization-alarm raise-threshold 90
set security nat source pool-utilization-alarm clear-threshold 80
set security nat source rule-set SR_SET_1 from zone trust
set security nat source rule-set SR_SET_1 to zone untrust1
set security nat source rule-set SR_SET_1 rule rule1 match source-address 10.11.0.0/16
set security nat source rule-set SR_SET_1 rule rule1 match source-address 10.147.0.0/16
set security nat source rule-set SR_SET_1 rule rule1 match destination-address 0.0.0.0/0
set security nat source rule-set SR_SET_1 rule rule1 then source-nat pool pool_1

```

```

set security nat source rule-set SR_SET_1 rule rule2 match source-address 10.148.1.0/27
set security nat source rule-set SR_SET_1 rule rule2 match destination-address 0.0.0.0/0
set security nat source rule-set SR_SET_1 rule rule2 then source-nat interface
set security nat source rule-set SR_SET_2 from zone trust
set security nat source rule-set SR_SET_2 to zone untrust2
set security nat source rule-set SR_SET_2 rule rule3 match source-address 10.140.21.0/27
set security nat source rule-set SR_SET_2 rule rule3 then source-nat pool pool_3
set security nat source rule-set SR_SET_2 rule rule4 match source-address 10.150.45.0/24
set security nat source rule-set SR_SET_2 rule rule4 then source-nat off
set security nat destination pool dppol_1 address 10.101.1.10/32
set security nat destination pool dppol_1 address port 21
set security nat destination pool dppol_2 address 10.101.1.11/32
set security nat destination pool dppol_2 address port 2101
set security nat destination pool dppol_3 address 10.103.12.251/32
set security nat destination pool dppol_3 address port 23
set security nat destination pool dppol_4 address 10.103.12.241/32
set security nat destination pool dppol_4 address port 23
set security nat destination pool dppol_5 address 10.103.1.11/32
set security nat destination pool dppol_5 address port 22
set security nat destination rule-set DR_SET1 from routing-instance isp1
set security nat destination rule-set DR_SET1 rule rule1 match destination-address
192.168.0.10/32
set security nat destination rule-set DR_SET1 rule rule1 match destination-port 7230
set security nat destination rule-set DR_SET1 rule rule1 then destination-nat pool dppol_1
set security nat destination rule-set DR_SET1 rule rule2 match destination-address 192.169.1.0/24
set security nat destination rule-set DR_SET1 rule rule2 then destination-nat pool dppol_2
set security nat destination rule-set DR_SET2 from routing-instance isp2
set security nat destination rule-set DR_SET2 rule rule3 match destination-address 192.168.2.2/32
set security nat destination rule-set DR_SET2 rule rule3 match destination-port 7351
set security nat destination rule-set DR_SET2 rule rule3 then destination-nat pool dppol_3
set security nat destination rule-set DR_SET2 rule rule4 match destination-address
192.168.4.171/32
set security nat destination rule-set DR_SET2 rule rule4 match destination-port 3451
set security nat destination rule-set DR_SET2 rule rule4 then destination-nat pool dppol_4
set security nat static rule-set ST_SET1 from zone trust
set security nat static rule-set ST_SET1 rule rule1 match destination-address 10.0.10.0/24
set security nat static rule-set ST_SET1 rule rule1 then static-nat prefix 192.168.5.0/24
set security nat static rule-set ST_SET2 from routing-instance isp1
set security nat static rule-set ST_SET2 rule rule2 match destination-address 192.168.6.0/24
set security nat static rule-set ST_SET2 rule rule2 then static-nat prefix 10.107.30.0/24
set security nat static rule-set ST_SET2 rule rule3 match destination-address 192.168.0.10/32
set security nat static rule-set ST_SET2 rule rule3 then static-nat prefix 10.171.9.23/32

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the Junos OS CLI User Guide.

1. Configure routing instances.

```
[edit ]
user@host# set routing-instances isp1 instance-type virtual-router
user@host# set routing-instances isp1 interface ge-0/0/2.0
user@host# set routing-instances isp1 routing-options static route 10.0.0.0/8 next-table
inet.0
user@host# set routing-instances isp1 routing-options static route 0.0.0.0/0 next-hop
192.0.2.20
user@host# set routing-instances isp2 instance-type virtual-router
user@host# set routing-instances isp2 interface ge-0/0/3.0
user@host# set routing-instances isp2 routing-options static route 10.0.0.0/8 next-table
inet.0
user@host# set routing-instances isp2 routing-options static route 0.0.0.0/0 next-hop
198.51.100.251
```

2. Configure rib groups and routing options.

```
[edit ]
user@host# set routing-options interface-routes rib-group inet isp
user@host# set routing-options static route 10.0.0.0/8 next-hop 10.0.21.254
user@host# set routing-options rib-groups isp import-rib inet.0
user@host# set routing-options rib-groups isp import-rib isp1.inet.0
user@host# set routing-options rib-groups isp import-rib isp2.inet.0
```

3. Configure security policies.

```
[edit security policies]
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol match source-address any
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol match destination-address
any
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol match application any
user@host# set from-zone trust to-zone untrust1 policy tr-untr1-pol then permit
user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol match source-address any
```



```

user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol match destination-address
any
user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol match application any
user@host# set from-zone trust to-zone untrust2 policy tr-untr2-pol then permit
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match source-
address any
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match destination-
address anyfrom-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match destination-
address any
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol match application
any
user@host# set from-zone untrust1 to-zone untrust2 policy untr1-untr2-pol then reject
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match source-
address any
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match destination-
address any
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol match application
any
user@host# set from-zone untrust2 to-zone untrust1 policy untr2-untr1-pol then reject
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match source-address any
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match destination-address
ftp-ser
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match destination-address
telnet-ser
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match application junos-
ftp
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol match application junos-
telnet
user@host# set from-zone untrust1 to-zone trust policy untr1-tr-pol then permit
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match source-address any
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match destination-address
10.171.9.23/32
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match destination-address
http-ser
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match destination-address
10.103.12.0/24
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match application junos-
http
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match application junos-
icmp-all
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol match application junos-

```

dhcp-server

```
user@host# set from-zone untrust2 to-zone trust policy untr2-tr-pol then permit
```

4. Configure source NAT pools and rules.

```
[edit security nat]
user@host# set source pool pool_1 address 192.0.2.40/32 to 192.0.2.190/32
user@host# set source pool pool_2 address 192.0.2.250/32
user@host# set source pool pool_3 address 198.51.100.20/32 to 198.51.100.30/32
user@host# set source address-persistent
user@host# set source pool-utilization-alarm raise-threshold 90
user@host# set source pool-utilization-alarm clear-threshold 80
user@host# set source rule-set SR_SET_1 from zone trust
user@host# set source rule-set SR_SET_1 to zone untrust1
user@host# set source rule-set SR_SET_1 rule rule1 match source-address 10.11.0.0/16
user@host# set source rule-set SR_SET_1 rule rule1 match source-address 10.147.0.0/16
user@host# set source rule-set SR_SET_1 rule rule1 match destination-address 0.0.0.0/0
user@host# set source rule-set SR_SET_1 rule rule1 then source-nat pool pool_1
user@host# set source rule-set SR_SET_1 rule rule2 match source-address 10.148.1.0/27
user@host# set source rule-set SR_SET_1 rule rule2 match destination-address 0.0.0.0/0
user@host# set source rule-set SR_SET_1 rule rule2 then source-nat interface
user@host# set source rule-set SR_SET_2 from zone trust
user@host# set source rule-set SR_SET_2 to zone untrust2
user@host# set source rule-set SR_SET_2 rule rule3 match source-address 10.140.21.0/27
user@host# set source rule-set SR_SET_2 rule rule3 then source-nat pool pool_3
user@host# set source rule-set SR_SET_2 rule rule4 match source-address 10.150.45.0/24
user@host# set source rule-set SR_SET_2 rule rule4 then source-nat off
```

5. Configure destination NAT pools and rules.

```
[edit security nat]
user@host#set destination pool dppol_1 address 10.101.1.10/32
user@host#set destination pool dppol_1 address port 21
user@host#set destination pool dppol_2 address 10.101.1.11/32
user@host#set destination pool dppol_2 address port 2101
user@host#set destination pool dppol_3 address 10.103.12.251/32
user@host#set destination pool dppol_3 address port 23
user@host#set destination pool dppol_4 address 10.103.12.241/32
user@host#set destination pool dppol_4 address port 23
user@host#set destination pool dppol_5 address 10.103.1.11/32
user@host#set destination pool dppol_5 address port 22
```

```

user@host#set destination rule-set DR_SET1 from routing-instance isp1
user@host#set destination rule-set DR_SET1 rule rule1 match destination-address
192.168.0.10/32
user@host#set destination rule-set DR_SET1 rule rule1 match destination-port 7230
user@host#set destination rule-set DR_SET1 rule rule1 then destination-nat pool dppol_1
user@host#set destination rule-set DR_SET1 rule rule2 match destination-address 192.169.1.0/24
user@host#set destination rule-set DR_SET1 rule rule2 then destination-nat pool dppol_2
user@host#set destination rule-set DR_SET2 from routing-instance isp2
user@host#set destination rule-set DR_SET2 rule rule3 match destination-address 192.168.2.2/32
user@host#set destination rule-set DR_SET2 rule rule3 match destination-port 7351
user@host#set destination rule-set DR_SET2 rule rule3 then destination-nat pool dppol_3
user@host#set destination rule-set DR_SET2 rule rule4 match destination-address
192.168.4.171/32
user@host#set destination rule-set DR_SET2 rule rule4 match destination-port 3451
user@host#set destination rule-set DR_SET2 rule rule4 then destination-nat pool dppol_4

```

6. Configure static NAT rules.

```

[edit security nat]
user@host#set static rule-set ST_SET1 from zone trust
user@host#set static rule-set ST_SET1 rule rule1 match destination-address 10.0.10.0/24
user@host#set static rule-set ST_SET1 rule rule1 then static-nat prefix 192.168.5.0/24
user@host#set static rule-set ST_SET2 from routing-instance isp1
user@host#set static rule-set ST_SET2 rule rule2 match destination-address 192.168.6.0/24
user@host#set static rule-set ST_SET2 rule rule2 then static-nat prefix 10.107.30.0/24
user@host#set static rule-set ST_SET2 rule rule3 match destination-address 192.168.7.2/32
user@host#set static rule-set ST_SET2 rule rule3 then static-nat prefix 10.171.9.23/32

```

Results

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

```

user@host# show configuration routing-intances
routing-instances {
    isp1 {
        instance-type virtual-router;
    }
}

```

```

        interface ge-0/0/2.0;
        routing-options {
            static {
                route 10.0.0.0/8 next-table inet.0;
                route 0.0.0.0/0 next-hop 192.0.2.20;
            }
        }
    }
    isp2 {
        instance-type virtual-router;
        interface ge-0/0/3.0;
        routing-options {
            static {
                route 10.0.0.0/8 next-table inet.0;
                route 0.0.0.0/0 next-hop 198.51.100.251;
            }
        }
    }
}

```

```

user@host# show configuration routing-options
routing-options {
    interface-routes {
        rib-group inet isp;
    }
    static {
        route 10.0.0.0/8 next-hop 10.0.21.254;
    }
    rib-groups {
        isp {
            import-rib [ isp1.inet.0 isp2.inet.0 ];
        }
    }
}

```

```

user@host# show configuration policies
policies {
    from-zone trust to-zone untrust1 {
        policy tr-untr1-pol {
            match {

```

```

        source-address any;
        destination-address any;
        application any;
    }
    then {
        permit;
    }
}
}
from-zone trust to-zone untrust2 {
    policy tr-untr2-pol {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            permit;
        }
    }
}
from-zone untrust1 to-zone untrust2 {
    policy untr1-untr2-pol {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            reject;
        }
    }
}
from-zone untrust2 to-zone untrust1 {
    policy untr2-untr1-pol {
        match {
            source-address any;
            destination-address any;
            application any;
        }
        then {
            reject;
        }
    }
}

```

```

    }
}
from-zone untrust1 to-zone trust {
    policy untr1-tr-pol {
        match {
            source-address any;
            destination-address [ ftp-ser telnet-ser ];
            application [ junos-ftp junos-telnet ];
        }
        then {
            permit;
        }
    }
}
from-zone untrust2 to-zone trust {
    policy untr2-tr-pol {
        match {
            source-address any;
            destination-address [ 10.171.9.23/32 http-ser 10.103.12.0/24 ];
            application [ junos-http junos-icmp-all junos-dhcp-server ];
        }
        then {
            permit;
        }
    }
}
}
}

```

```

user@host# show configuration security nat
security {
    nat {
        source {
            pool pool_1 {
                address {
                    192.0.2.40/32 to 192.0.2.190/32;
                }
            }
            pool pool_2 {
                address {
                    192.0.2.250/32;
                }
            }
        }
    }
}

```

```

}
pool pool_3 {
    address {
        198.51.100.20/32 to 198.51.100.30/32;
    }
}
address-persistent;
pool-utilization-alarm raise-threshold 90 clear-threshold 80;
rule-set SR_SET_1 {
    from zone trust;
    to zone untrust1;
    rule rule1 {
        match {
            source-address [ 10.11.0.0/16 10.147.0.0/16 ];
            destination-address 0.0.0.0/0;
        }
        then {
            source-nat {
                pool {
                    pool_1;
                }
            }
        }
    }
}
rule rule2 {
    match {
        source-address 10.148.1.0/27;
        destination-address 0.0.0.0/0;
    }
    then {
        source-nat {
            interface;
        }
    }
}
}
rule-set SR_SET_2 {
    from zone trust;
    to zone untrust2;
    rule rule3 {
        match {
            source-address 10.140.21.0/27;
        }
    }
}

```

```

        then {
            source-nat {
                pool {
                    pool_3;
                }
            }
        }
    }
    rule rule4 {
        match {
            source-address 10.150.45.0/24;
        }
        then {
            source-nat {
                off;
            }
        }
    }
}
}
}

```

user@host# **show configuration security nat**

```

destination {
    pool dppol_1 {
        address 10.101.1.10/32 port 21;
    }
    pool dppol_2 {
        address 10.101.1.11/32 port 2101;
    }
    pool dppol_3 {
        address 10.103.12.251/32 port 23;
    }
    pool dppol_4 {
        address 10.103.12.241/32 port 23;
    }
    pool dppol_5 {
        address 10.103.1.11/32 port 22;
    }
    rule-set DR_SET1 {
        from routing-instance isp1;
        rule rule1 {

```



```

        match {
            destination-address 192.168.0.10/32;
            destination-port 7230;
        }
        then {
            destination-nat pool dppol_1;
        }
    }
    rule rule2 {
        match {
            destination-address 192.169.1.0/24;
        }
        then {
            destination-nat pool dppol_2;
        }
    }
}
rule-set DR_SET2 {
    from routing-instance isp2;
    rule rule3 {
        match {
            destination-address 192.168.2.2/32;
            destination-port 7351;
        }
        then {
            destination-nat pool dppol_3;
        }
    }
    rule rule4 {
        match {
            destination-address 192.168.4.171/32;
            destination-port 3451;
        }
        then {
            destination-nat pool dppol_4;
        }
    }
}

```

```

    }
}

```

```

user@host# show configuration static nat
static {
    rule-set ST_SET1 {
        from zone trust;
        rule rule1 {
            match {
                destination-address 10.0.10.0/24;
            }
            then {
                static-nat prefix 192.168.5.0/24;
            }
        }
    }
    rule-set ST_SET2 {
        from routing-instance isp1;
        rule rule2 {
            match {
                destination-address 192.168.6.0/24;
            }
            then {
                static-nat prefix 10.107.30.0/24;
            }
        }
        rule rule3 {
            match {
                destination-address 192.168.7.2/32;
            }
            then {
                static-nat prefix 10.171.9.23/32;
            }
        }
    }
}
}
}
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Interfaces | 24](#)

Verifying Interfaces

Purpose

Verify that the interfaces are configured correctly.

Action

From operational mode, enter the following commands:

- `show interfaces`
- `show zones`
- `show routing-instances`
- `show routing-options`
- `show policies`
- `show source nat`
- `show destination nat`
- `show static nat`

Configuring Proxy ARP for NAT (CLI Procedure)

You use NAT proxy ARP functionality to configure proxy ARP entries for IP addresses that require either source or destination NAT and that are in the same subnet as the ingress interface.

NOTE: On SRX Series Firewalls, you must explicitly configure NAT proxy ARP.

When configuring NAT proxy ARP, you must specify the logical interface on which to configure proxy ARP. Then you enter an address or address range.

The device performs proxy ARP for the following conditions:

- When addresses defined in the static NAT and source NAT pool are in the same subnet as that of the ingress interface
- When addresses in the original destination address entry in the destination NAT rules are in the same subnet as that of the ingress interface

```
user@host# set security nat proxy-arp interface fe-0/0/0.0 address 10.1.1.10 to 10.1.1.20
```

Configuring NAT trace options

IN THIS SECTION

- [Purpose | 25](#)
- [Action | 26](#)

Purpose

The NAT trace options hierarchy configures trace file and flags for verification purposes.

SRX Series Firewalls have two main components: the Routing Engine (RE) and the Packet Forwarding Engine (PFE). The PFE is divided into the ukernel portion and the real-time portion.

When a NAT configuration is committed, the configuration is first checked and validated on the RE. After validation, the configuration is pushed to the PFE. The configuration is installed on the ukernel PFE, then action is taken on each packet that matches NAT rules on the real-time PFE.

For verification, you can turn on flags individually to debug NAT functionality on the RE, ukernel PFE, or real-time PFE:

- The `nat-re` flag records the trace of the NAT configuration validation on the RE and the configuration push to the PFE.
- The `nat-pfe` flag records the trace of the NAT configuration installation on the ukernel PFE.
- The `nat-rt` flag records the trace of the NAT rule match, and subsequent action on the real-time PFE.

The trace data is written to `/var/log/security-trace` by default, and can be viewed using the command `show log security-trace`.

NOTE: If session logging has been enabled in the policy configurations on the device, the session logs will include specific NAT details for each session. See *Monitoring Security Policy Statistics* for information on how to enable session logging and [Information Provided in Session Log Entries for SRX Series Services Gateways](#) for a description of information provided in session logs.

Action

To verify that NAT configurations are correctly updated to the device upon commit, and that the NAT rule match and subsequent actions are correct, use the `security nat traceoptions` statement.

```
user@host# set security nat traceoptions flag all
user@host# set security nat traceoptions flag destination-nat-pfe
user@host# set security nat traceoptions flag destination-nat-re
user@host# set security nat traceoptions flag destination-nat-rti
user@host# set security nat traceoptions flag source-nat-pfe
user@host# set security nat traceoptions flag source-nat-re
user@host# set security nat traceoptions flag source-nat-rt
user@host# set security nat traceoptions flag static-nat-pfe
user@host# set security nat traceoptions flag static-nat-re
user@host# set security nat traceoptions flag static-nat-rt
```

To verify that NAT translations are being applied to the traffic, and to view individual traffic flow processing with NAT translations, use both the `security nat traceoptions` command and the `security flow traceoptions` command together. The commands are used together because the NAT trace, configured using the `security nat traceoptions` command, is not recorded unless the `flow traceoptions` command is also configured.

To filter a specific flow, you can define a packet filter and use it as a traceoption :

```
user@host# set security flow traceoptions packet-filter packet-filter
user@host# set security flow traceoptions packet-filter packet-filter apply-groups
user@host# set security flow traceoptions packet-filter packet-filter apply-groups-except
user@host# set security flow traceoptions packet-filter packet-filter destination-port
user@host# set security flow traceoptions packet-filter packet-filter destination-prefix
user@host# set security flow traceoptions packet-filter packet-filter interface
user@host# set security flow traceoptions packet-filter packet-filter protocol
user@host# set security flow traceoptions packet-filter packet-filter source-port
user@host# set security flow traceoptions packet-filter packet-filter source-prefix
```

To verify NAT traffic and to enable all traffic trace in data plane, use the traceoptions set security flow traceoptions flag basic-datapath command, as shown in the following example using a simple packet filter:

```
user@host# set security flow traceoptions file filename
user@host# set security flow traceoptions flag basic-datapath
user@host# set security flow traceoptions packet-filter client-traffic source-prefixprefix
user@host# set security flow traceoptions packet-filter client-traffic destination-prefixprefix
user@host# set security nat traceoptions flag all
```

Monitoring NAT Incoming Table Information

IN THIS SECTION

- Purpose | 27
- Action | 28

Purpose

View NAT table information.

Action

Select **Monitor>NAT>Incoming Table** in the J-Web user interface, or enter the following CLI command:

```
show security nat incoming-table
```

[Table 3 on page 28](#) summarizes key output fields in the incoming table display.

Table 3: Summary of Key Incoming Table Output Fields

Field	Values
Statistics	
In use	Number of entries in the NAT table.
Maximum	Maximum number of entries possible in the NAT table.
Entry allocation failed	Number of entries failed for allocation.
Incoming Table	
Clear	
Destination	Destination IP address and port number.
Host	Host IP address and port number that the destination IP address is mapped to.
References	Number of sessions referencing the entry.
Timeout	Timeout, in seconds, of the entry in the NAT table.
Source-pool	Name of source pool where translation is allocated.

Monitoring Interface NAT Port Information

IN THIS SECTION

- Purpose | 29
- Action | 29

Purpose

View port usage for an interface source pool information.

Action

To monitoring interface NAT port information, do one of the following:

- If you are using SRX5400, SRX5600, or SRX5800 platforms, select **Monitor>Firewall/NAT>Interface NAT** in the J-Web user interface or enter the CLI command `show security nat interface-nat-ports`.
- Select **Monitor>NAT>Interface NAT Ports** in the J-Web user interface.

[Table 4 on page 29](#) summarizes key output fields in the interface NAT display.

Table 4: Summary of Key Interface NAT Output Fields

Field	Values	Additional Information
Interface NAT Summary Table		
Pool Index	Port pool index.	–
Total Ports	Total number of ports in a port pool.	–

Table 4: Summary of Key Interface NAT Output Fields *(Continued)*

Field	Values	Additional Information
Single Ports Allocated	Number of ports allocated one at a time that are in use.	-
Single Ports Available	Number of ports allocated one at a time that are free for use.	-
Twin Ports Allocated	Number of ports allocated two at a time that are in use.	-
Twin Ports Available	Number of ports allocated two at a time that are free for use.	-

2

CHAPTER

Types of NAT

Source NAT | 32

Destination NAT | 125

Static NAT | 157

Source NAT

IN THIS SECTION

- [Understanding Source NAT | 33](#)
- [Understanding Central Point Architecture Enhancements for NAT | 34](#)
- [Port Overflow Burst Mode | 35](#)
- [Optimizing Source NAT Performance | 36](#)
- [Monitoring Source NAT Information | 39](#)
- [Source NAT Configuration Overview | 48](#)
- [Example: Configuring Source NAT for Egress Interface Translation | 48](#)
- [Example: Configuring Source NAT for Single Address Translation | 54](#)
- [Example: Configuring Source and Destination NAT Translations | 61](#)
- [Understanding Source NAT Rules | 71](#)
- [Example: Configuring Source NAT with Multiple Rules | 72](#)
- [Understanding Source NAT Pools | 81](#)
- [Understanding Source NAT Pool Capacities | 83](#)
- [Understanding Persistent Addresses for Source NAT Pools | 84](#)
- [Example: Configuring Capacity for Source NAT Pools with PAT | 85](#)
- [Understanding Source NAT Pools with Address Pooling | 88](#)
- [Understanding Source NAT Pools with Address Shifting | 88](#)
- [Example: Configuring Source NAT Pools with Address Shifting | 89](#)
- [Understanding Source NAT Pools with PAT | 96](#)
- [Example: Configuring Source NAT for Multiple Addresses with PAT | 97](#)
- [Understanding Source NAT Pools Without PAT | 104](#)
- [Example: Configuring a Single IP Address in a Source NAT Pool Without PAT | 105](#)
- [Example: Configuring Multiple Addresses in a Source NAT Pool Without PAT | 110](#)
- [Understanding Shared Addresses in Source NAT Pools without PAT | 117](#)
- [Understanding NAT Session Persistence | 118](#)
- [Configure Port Block Allocation Size | 120](#)
- [Configuring the NAT Session Hold Timeout and NAT Session Persistence Scan | 122](#)
- [Understanding NAT Configuration Check on Egress Interfaces after Reroute | 123](#)

Source NAT is most commonly used for translating private IP address to a public routable address to communicate with the host. Source NAT changes the source address of the packets that pass through the Router. A NAT pool is a set of addresses that are designed as a replacement for client IP addresses. For more information, see the following topics:

Understanding Source NAT

Source NAT is the translation of the source IP address of a packet leaving the Juniper Networks device. Source NAT is used to allow hosts with private IP addresses to access a public network.

Source NAT allows connections to be initiated only for outgoing network connections—for example, from a private network to the Internet. Source NAT is commonly used to perform the following translations:

- Translate a single IP address to another address (for example, to provide a single device in a private network with access to the Internet).
- Translate a contiguous block of addresses to another block of addresses of the same size.
- Translate a contiguous block of addresses to another block of addresses of smaller size.
- Translate a contiguous block of addresses to a single IP address or a smaller block of addresses using port translation.
- Translate a contiguous block of addresses to the address of the egress interface.

Translation to the address of the egress interface does not require an address pool; all other source NAT translations require configuration of an address pool. One-to-one and many-to-many translations for address blocks of the same size do not require port translation because there is an available address in the pool for every address that would be translated.

If the size of the address pool is smaller than the number of addresses that would be translated, either the total number of concurrent addresses that can be translated is limited by the size of the address pool or port translation must be used. For example, if a block of 253 addresses is translated to an address pool of 10 addresses, a maximum of 10 devices can be connected concurrently unless port translation is used.

The following types of source NAT are supported:

- Translation of the original source IP address to the egress interface's IP address (also called interface NAT). Port address translation is always performed.
- Translation of the original source IP address to an IP address from a user-defined address pool without port address translation. The association between the original source IP address to the

translated source IP address is dynamic. However, once there is an association, the same association is used for the same original source IP address for new traffic that matches the same NAT rule.

- Translation of the original source IP address to an IP address from a user-defined address pool with port address translation. The association between the original source IP address to the translated source IP address is dynamic. Even if an association exists, the same original source IP address may be translated to a different address for new traffic that matches the same NAT rule.
- Translation of the original source IP address to an IP address from a user-defined address pool by shifting the IP addresses. This type of translation is one-to-one, static, and without port address translation. If the original source IP address range is larger than the IP address range in the user-defined pool, untranslated packets are dropped.

Understanding Central Point Architecture Enhancements for NAT

System session capacity and session ramp-up rate are limited by central point memory capacity and CPU capacity. Starting in Junos OS Release 15.1X49-D30 and Junos OS Release 17.3R1, the central point architecture for NAT has been enhanced to handle higher system session capacity and session ramp-up rate for the SRX5000 line. Hence, the workload on the central point is reduced to increase the session capacity and to support more sessions to achieve higher connections per second (CPS). Starting in Junos OS Release 17.4R1, source NAT resources handled by the central point architecture have been offloaded to the SPUs when the SPC number is more than four, resulting in more efficient resource allocation. The following list describes the enhancements to NAT to improve performance:

- The central point architecture no longer supports central point sessions. Therefore, NAT needs to maintain a NAT tracker to track the IP address or port allocation and usage. NAT tracker is a global array for SPU session ID to NAT IP or port mapping that is used to manage NAT resources.
- By default, a NAT rule alarm and trap statistics counter update message is sent from the Services Processing Unit (SPU) to the central point at intervals of 1 second instead of updating the statistics based on each session trigger in the central point system.
- To support a specific NAT IP address or port allocated such that the 5-tuple hash after NAT is the same as the original 5-tuple hash before NAT, select a NAT port that results in the same hash as the original hash by the specific calculation. Hence, the forwarding session is reduced. When NAT is used, the reverse wing is hashed to a different SPU. A forward session has to be installed to forward reverse traffic to a session SPU. NAT tries to select a port that can be used by the hash algorithm to make the reverse wing be hashed to the same SPU as the initial wing. So, both NAT performance and throughput are improved with this approach.
- To improve NAT performance, IP shifting pool (non-PAT pool) management is moved from the central point to the SPU so that all local NAT resources for that pool are managed locally instead of sending

the NAT request to the central point. Hence, IP address-shifting NAT pool connections per second and throughput are improved.

Port Overflow Burst Mode

The port overflow burst mode allows you to use the ports beyond the allocated port blocks. You can configure a burst pool with a range of ports in an IP address to be reserved for bursting.

There are primary and burst pool types, the device uses the burst pool once the subscribers reach the limit configured in the primary pool.

Burst mode is supported on:

1. Deterministic NAT source NAT pool with PBA type burst pool.
2. Deterministic NAT source NAT pool with dynamic Network Address Port Translation (NAPT) type burst pool.
3. Regular PBA source NAT pool with PBA type burst pool.
4. Regular PBA source NAT pool with dynamic NAPT type burst pool.

Table 5: Burst Mode Operations

NAT Type	Before the Configured Port Block Limit not Exceed	After the Configured Port Block Limit not Exceed
Deterministic NAT source NAT pool with PBA type burst pool	Port blocks from the primary DetNAT pool are used.	Port block from the burst pool configured in PBA.
Deterministic NAT source NAT pool with dynamic Network Address Port Translation (NAPT) type burst pool	Port blocks from the primary DetNAT pool are used.	Port block from the burst pool configured in dynamic NAPT.
Regular PBA source NAT pool with PBA type burst pool	Port blocks from the primary PBA pool are used.	Port block from the burst pool configured in PBA.
Regular PBA source NAT pool with dynamic NAPT type burst pool	Port blocks from the primary PBA pool are used.	Port block from the burst pool configured in dynamic NAPT.

PBA Burst Type Method—PBA supports APP and non-APP mode of operations.

- **APP Mode**—Ports are allocated from the primary pool. When then subscriber limit exceeds from primary pool, if there are any available ports for the same IP address from the burst pool, then new sessions are created.
- **non-APP Mode**—Ports are allocated from the primary pool. When subscriber limit exceeds from primary pool, new sessions are created from the burst pool with any available IP address and port.

DetNAT Burst Type Method—Ports are allocated from the primary pool. If the same IP address from the burst pool or all the available ports are not available from same IP address, then new session are created with another IP address. If the burst pool is configured with a different IP from primary pool, uses another IP from the burst pool.

Optimizing Source NAT Performance

IN THIS SECTION

- [Port Randomization Mode \(Default\) | 36](#)
- [Round-Robin Mode | 37](#)
- [Session Affinity Mode | 38](#)

Source NAT can be optimized based on functionality and performance needs.

Port Randomization Mode (Default)

For pool-based source NAT and interface NAT, port randomization mode is enabled and used by default.

In this mode, the device selects IP addresses on a round-robin basis, and the port selection is random. That is, when the device performs NAT translation it first chooses the IP address by round robin, then chooses the port used for that IP address by randomization.

Although randomized port number allocation can provide protection from security threats such as DNS poison attacks, it can also affect performance and memory usage due to the computations and NAT table resources involved.

Round-Robin Mode

A less resource-intensive NAT translation method involves using only the round-robin allocation method. Whereas randomization requires computational work for each assigned port, the round robin method simply selects ports sequentially.

In this mode, the device selects both IP addresses and ports on a round-robin basis. That is, when the device performs NAT translation it first chooses the IP address by round robin, then chooses the port used for that IP address by round robin.

For example, if the source pool contains only one IP address:

- When the first packet of a flow arrives (creating a session), it is translated to IP1, port N. Subsequent packets in that flow are allocated to the same IP/port.
- When the first packet of a new flow arrives, it is translated to IP1, port N+1, and so on.

If the source pool contains two IP addresses:

- When the first packet of a flow arrives (creating a session), it is translated to IP1, port X. Subsequent packets in that flow are allocated to the same IP/port.
- When the first packet of a second flow arrives, it is translated to IP2, port X.
- When the first packet of a third flow arrives, it is translated to IP1, port X+1.
- With the first packets of a fourth flow arrives, it is translated to IP2, port X+1, and so on.

Configuration

Round-robin mode is enabled by default, however port randomization mode (also enabled) has higher priority. To use round-robin mode, disable the higher-priority port randomization mode, as follows:

```
user@host# set security nat source port-randomization disable
```

To disable round-robin mode (and re-enable port randomization), delete the configuration statement, as follows:

```
user@host# delete security nat source port-randomization disable
```


Session Affinity Mode

Starting in Junos OS Release 15.1X49-D30 and Junos OS Release 17.3R1, you can further improve NAT performance and throughput on SRX5000 line devices using “session affinity” mode.

With the modes noted above, a given session is processed by the inbound SPU based on a 5-tuple (source IP, dest IP, source port, dest port, protocol) hash. When NAT is involved, the 5-tuple hash will be different for the outbound part of the session vs. the return part of the session. Therefore, the outbound NAT session information may be located in one SPU, while the return (reverse) NAT session information may be located in another SPU. The goal of session affinity mode is to maintain the forwarding session information for both the outbound and return traffic on the same SPU.

In this mode, the device uses a “reverse NAT enhancement” translation algorithm for IP and port selection, to improve performance for NAT sessions and throughput. The NAT module attempts to select an IP address and port that can be used with the hash algorithm to ensure the selected SPU for the outbound and return flow elements can be identical.

Configuration

Session affinity mode is enabled by default, however both port randomization and round-robin modes (also enabled) have higher priority. To use session affinity mode, disable both port randomization and round-robin modes, as follows:

```
user@host# set security nat source port-randomization disable
user@host# set security nat source port-round-robin disable
```

To disable session affinity mode, and re-enable either round-robin or port randomization mode, delete one or both of the configuration statements, as follows:

```
user@host# delete security nat source port-round-robin disable
user@host# delete security nat source port-randomization disable
```

Usage Notes

Notes and guidelines for session affinity mode include:

- Use large NAT port pools whenever possible (see Security Considerations below)
- The algorithm chooses a port from within the configured port range. If no port is available, the NAT port will be allocated based on random selection.
- Static NAT and destination NAT cannot use affinity mode.

Security Considerations

Although session affinity improves performance by consolidating forwarding sessions, it decreases security to some degree since the algorithm selects the IP address and port based on a pre-defined algorithm with specific parameters, instead of pure randomization. That said, the fact there are typically multiple eligible ports for the algorithm to choose from and so there is still some degree of randomization.

The best way to mitigate the security risk is to ensure the source port number used is less predictable. That is, the larger the NAT pool resource range from which ephemeral ports are selected, the smaller the chances of an attacker guessing the selected port number. Given this, it is recommended to configure large NAT port pools whenever possible.

Monitoring Source NAT Information

IN THIS SECTION

- Purpose | 39
- Action | 39

Purpose

Display configured information about source Network Address Translation (NAT) rules, pools, persistent NAT, and paired addresses.

Action

Select **Monitor>NAT>Source NAT** in the J-Web user interface, or enter the following CLI commands:

- `show security nat source summary`
- `show security nat source pool pool-name`
- `show security nat source persistent-nat-table`
- `show security nat source paired-address`

[Table 6 on page 40](#) describes the available options for monitoring source NAT.

Table 6: Source NAT Monitoring Page

Field	Description	Action
Rules		
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.
Total rules	Number of rules configured.	–
ID	Rule ID number.	–
Name	Name of the rule .	–
From	Name of the routing instance/zone/interface from which the packet flows.	–
To	Name of the routing instance/zone/interface to which the packet flows.	–
Source address range	Source IP address range in the source pool.	–
Destination address range	Destination IP address range in the source pool.	–
Source ports	Source port numbers.	–
Ip protocol	IP protocol.	–
Action	Action taken for a packet that matches a rule.	–

Table 6: Source NAT Monitoring Page *(Continued)*

Field	Description	Action
Persistent NAT type	Persistent NAT type.	–
Inactivity timeout	Inactivity timeout interval for the persistent NAT binding.	–
Alarm threshold	Utilization alarm threshold.	
Max session number	The maximum number of sessions.	–
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. <ul style="list-style-type: none"> • Succ–Number of successful session installations after the NAT rule is matched. • Failed–Number of unsuccessful session installations after the NAT rule is matched. • Current–Number of sessions that reference the specified rule. 	–
Translation Hits	Number of times a translation in the translation table is used for a source NAT rule.	–
Pools		
Pool Name	The names of the pools.	Select all pools or a specific pool to display from the list.
Total Pools	Total pools added.	–
ID	ID of the pool.	–

Table 6: Source NAT Monitoring Page *(Continued)*

Field	Description	Action
Name	Name of the source pool.	–
Address range	IP address range in the source pool.	–
Single/ Twin ports	Number of allocated single and twin ports.	–
Port	Source port number in the pool.	–
Address assignme nt	Displays the type of address assignment.	–
Alarm threshold	Utilization alarm threshold.	–
Port overloadi ng factor	Port overloading capacity.	–
Routing instance	Name of the routing instance.	–
Total addresses	Total IP address, IP address set, or address book entry.	–
Host address base	Host base address of the original source IP address range.	–
Translatio n hits	Number of times a translation in the translation table is used for source NAT.	–

Table 6: Source NAT Monitoring Page *(Continued)*

Field	Description	Action
Top 10 Translation Hits		
Graph	Displays the graph of top 10 translation hits.	–
Persistent NAT		
Persistent NAT table statistics		
binding total	Displays the total number of persistent NAT bindings for the FPC.	–
binding in use	Number of persistent NAT bindings that are in use for the FPC.	–
enode total	Total number of persistent NAT enodes for the FPC.	–
enode in use	Number of persistent NAT enodes that are in use for the FPC.	–
Persistent NAT table		
Source NAT pool	Name of the pool.	Select all pools or a specific pool to display from the list.
Internal IP	Internal IP address.	Select all IP addresses or a specific IP address to display from the list.
Internal port	Displays the internal ports configured in the system.	Select the port to display from the list.
Internal protocol	Internal protocols .	Select all protocols or a specific protocol to display from the list.

Table 6: Source NAT Monitoring Page (Continued)

Field	Description	Action
Internal IP	Internal transport IP address of the outgoing session from internal to external.	-
Internal port	Internal transport port number of the outgoing session from internal to external.	-
Internal protocol	Internal protocol of the outgoing session from internal to external.	-
Reflective IP	Translated IP address of the source IP address.	-
Reflective port	Displays the translated number of the port.	-
Reflective protocol	Translated protocol.	-
Source NAT pool	Name of the source NAT pool where persistent NAT is used.	-
Type	Persistent NAT type.	-
Left time/ Conf time	Inactivity timeout period that remains and the configured timeout value.	-
Current session num/Max session num	Number of current sessions associated with the persistent NAT binding and the maximum number of sessions.	-
Source NAT rule	Name of the source NAT rule to which this persistent NAT binding applies.	-

Table 6: Source NAT Monitoring Page *(Continued)*

Field	Description	Action
External node table		
Internal IP	Internal transport IP address of the outgoing session from internal to external.	–
Internal port	Internal port number of the outgoing session from internal to external.	–
External IP	External IP address of the outgoing session from internal to external.	–
External port	External port of the outgoing session from internal to external.	–
Zone	External zone of the outgoing session from internal to external.	–
Paired Address		
Pool name	Name of the pool.	Select all pools or a specific pool to display from the list.
Specified Address	IP address.	Select all addresses, or select the internal or external IP address to display, and enter the IP address.
Pool name	Displays the selected pool or pools.	–
Internal address	Displays the internal IP address.	–
External address	Displays the external IP address.	–

Table 6: Source NAT Monitoring Page *(Continued)*

Field	Description	Action
Resource Usage		
Utilization for all source pools		
Pool name	Name of the pool.	To view additional usage information for Port Address Translation (PAT) pools, select a pool name. The information displays under Detail Port Utilization for Specified Pool.
Pool type	Pool type: PAT or Non-PAT.	–
Port overloading factor	Port overloading capacity for PAT pools.	–
Address	Addresses in the pool.	–
Used	<p>Number of used resources in the pool.</p> <p>For Non-PAT pools, the number of used IP addresses is displayed.</p> <p>For PAT pools, the number of used ports is displayed.</p>	–
Available	<p>Number of available resources in the pool.</p> <p>For Non-PAT pools, the number of available IP addresses is displayed.</p> <p>For PAT pools, the number of available ports is displayed.</p>	–

Table 6: Source NAT Monitoring Page *(Continued)*

Field	Description	Action
Total	<p>Number of used and available resources in the pool.</p> <p>For Non-PAT pools, the total number of used and available IP addresses is displayed.</p> <p>For PAT pools, the total number of used and available ports is displayed.</p>	–
Usage	<p>Percent of resources used.</p> <p>For Non-PAT pools, the percent of IP addresses used is displayed.</p> <p>For PAT pools, the percent of ports, including single and twin ports, is displayed.</p>	–
Peak usage	Percent of resources used during the peak date and time.	–
Detail Port Utilization for Specified Pool		
Address Name	IP addresses in the PAT pool.	Select the IP address for which you want to display detailed usage information.
Factor-Index	Index number.	–
Port-range	Displays the number of ports allocated at a time.	–
Used	Displays the number of used ports.	–
Available	Displays the number of available ports.	–
Total	Displays the number of used and available ports.	–

Table 6: Source NAT Monitoring Page *(Continued)*

Field	Description	Action
Usage	Displays the percentage of ports used during the peak date and time.	–

Source NAT Configuration Overview

The main configuration tasks for source NAT are as follows:

1. Configure an address pool or an interface NAT mapping of private addresses to the public address of an egress interface.
For an address pool, also do the following:
 - a. Specify the name of the pool, the addresses or address ranges, the routing instance, and whether to perform port address translation (PAT).
 - b. (Optional) Configure address pool options, such as overflow pool, IP address shifting, address sharing, address pooling, and pool utilization alarms.
 - c. Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.
2. (Optional) Configure the persistent address.
3. Configure source NAT rules that align with your network and security requirements.

Example: Configuring Source NAT for Egress Interface Translation

IN THIS SECTION

- [Requirements | 49](#)
- [Overview | 49](#)
- [Configuration | 51](#)
- [Verification | 53](#)

This example describes how to configure a source NAT mapping of private addresses to the public address of an egress interface.

Requirements

Before you begin:


1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 2 on page 50](#), devices with private addresses in the trust zone access a public network through the egress interface ge-0/0/0. For packets that enter the Juniper Networks security device from the trust zone with a destination address in the untrust zone, the source IP address is translated to the IP address of the egress interface.

NOTE: No source NAT pool is required for source NAT using an egress interface. Proxy ARP does not need to be configured for the egress interface.

Figure 2: Source NAT Egress Interface Translation



Original Source IP	Translated Source IP
0.0.0.0/0	203.0.113.63 (Interface IP)

This example describes the following configurations:

- Source NAT rule set `rs1` with a rule `r1` to match any packet from the trust zone to the untrust zone. For matching packets, the source address is translated to the IP address of the egress interface.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

IN THIS SECTION

- [Procedure | 51](#)

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat interface
set security policies from-zone trust to-zone untrust policy internet-access match source-
address any
set security policies from-zone trust to-zone untrust policy internet-access match destination-
address any
set security policies from-zone trust to-zone untrust policy internet-access match application
any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT translation to an egress interface:

1. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust
```

2. Configure a rule that matches packets and translates the source address to the address of the egress interface.

```
[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address 0.0.0.0/0
user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r1 then source-nat interface
```

3. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
    rule-set rs1 {
        from zone trust;
        to zone untrust;
        rule r1 {
            match {
                source-address 0.0.0.0/0;
                destination-address 0.0.0.0/0;
```

```
    }  
    then {  
        source-nat {  
            interface;  
        }  
    }  
}  
  
}
```

user@host# **show security policies**

```
from-zone trust to-zone untrust {  
    policy internet-access {  
        match {  
            source-address any;  
            destination-address any;  
            application any;  
        }  
        then {  
            permit;  
        }  
    }  
}
```


Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Example: Configuring Source NAT for Single Address Translation

IN THIS SECTION

- Requirements | 54
- Overview | 55
- Configuration | 57
- Verification | 60

This example describes how to configure a source NAT mapping of a single private address to a public address.

Requirements


Before you begin:

1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 3 on page 56](#), a device with the private address 192.168.1.200 in the trust zone accesses a public network. For packets sent by the device to a destination address in the untrust zone, the Juniper Networks security device translates the source IP address to the public IP address 203.0.113.200/32.

Figure 3: Source NAT Single Address Translation



Original Source IP	Translated Source IP
192.168.1.200/32	203.0.113.200/32

This example describes the following configurations:

- Source NAT pool `src-nat-pool-1` that contains the IP address `203.0.113.200/32`.
- Source NAT rule set `rs1` with rule `r1` to match packets from the trust zone to the untrust zone with the source IP address `192.168.1.200/32`. For matching packets, the source address is translated to the IP address in `src-nat-pool-1` pool.
- Proxy ARP for the address `203.0.113.200` on interface `ge-0/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

IN THIS SECTION

- [Procedure](#) | 57

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.200/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 192.168.1.200/32
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security policies from-zone trust to-zone untrust policy internet-access match source-
address any
set security policies from-zone trust to-zone untrust policy internet-access match destination-
address any
set security policies from-zone trust to-zone untrust policy internet-access match application
any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT translation for a single IP address:

1. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool src-nat-pool-1 address 203.0.113.200/32
```

2. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust
```

3. Configure a rule that matches packets and translates the source address to the address in the pool.

```
[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address 192.168.1.200/32
user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
```

4. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200
```

5. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
    pool src-nat-pool-1 {
        address {
            203.0.113.200/32;
        }
    }
    rule-set rs1 {
        from zone trust;
        to zone untrust;
        rule r1 {
            match {
                source-address 192.168.1.200/32;
            }
            then {
                source-nat {
                    pool {
                        src-nat-pool-1;
                    }
                }
            }
        }
    }
}
proxy-arp {
    interface ge-0/0/0.0 {
        address {
            203.0.113.200/32;
        }
    }
}

user@host# show security policies
from-zone trust to-zone untrust {
    policy internet-access {
        match {
```

```
        source-address any;  
        destination-address any;  
        application any;  
    }  
    then {  
        permit;  
    }  
}  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Source NAT Pool Usage | 60](#)
- [Verifying Source NAT Rule Usage | 60](#)
- [Verifying NAT Application to Traffic | 61](#)

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

Verifying Source NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the source NAT pool.

Action

From operational mode, enter the `show security nat source pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Example: Configuring Source and Destination NAT Translations

IN THIS SECTION

- [Requirements | 61](#)
- [Overview | 62](#)
- [Configuration | 64](#)
- [Verification | 69](#)

This example describes how to configure both source and destination NAT mappings.

Requirements

Before you begin:


1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 4 on page 63](#), the following translations are performed on the Juniper Networks security device:

- The source IP address in packets sent by the device with the private address 192.168.1.200 in the trust zone to any address in the untrust zone is translated to a public address in the range from 203.0.113.10 through 203.0.113.14.
- The destination IP address 203.0.113.100/32 in packets sent from the trust zone to the untrust zone is translated to the address 10.1.1.200/32.

Figure 4: Source and Destination NAT Translations



This example describes the following configurations:

- Source NAT pool `src-nat-pool-1` that contains the IP address range 203.0.113.10 through 203.0.113.14.
- Source NAT rule set `rs1` with rule `r1` to match any packets from the trust zone to the untrust zone. For matching packets, the source address is translated to an IP address in the `src-nat-pool-1` pool.
- Destination NAT pool `dst-nat-pool-1` that contains the IP address 10.1.1.200/32.

- Destination NAT rule set `rs1` with rule `r1` to match packets from the trust zone with the destination IP address `203.0.113.100`. For matching packets, the destination address is translated to the IP address in the `dst-nat-pool-1` pool.
- Proxy ARP for the addresses `203.0.113.10` through `203.0.113.14` and `203.0.113.100/32` on interface `ge-0/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policy to permit traffic from the trust zone to the untrust zone.
- Security policy to permit traffic from the untrust zone to the translated destination IP addresses in the trust zone.

Configuration

IN THIS SECTION

- [Procedure](#) | 64

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.10/32 to 203.0.113.14/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat destination pool dst-nat-pool-1 address 10.1.1.200/32
set security nat destination rule-set rs1 from zone untrust
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.100/32
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.10/32 to 203.0.113.24/32
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.100/32
set security policies from-zone trust to-zone untrust policy internet-access match source-
```

```

address any
set security policies from-zone trust to-zone untrust policy internet-access match destination-
address any
set security policies from-zone trust to-zone untrust policy internet-access match application
any
set security policies from-zone trust to-zone untrust policy internet-access then permit
set security address-book global address dst-nat-pool-1 10.1.1.200/32
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access match source-
address any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access match
destination-address dst-nat-pool-1
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access match
application any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-1-access then permit

```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the source and destination NAT translations:

1. Create a source NAT pool.

```

[edit security nat source]
user@host# set pool src-nat-pool-1 address 203.0.113.10 to 203.0.113.14

```

2. Create a source NAT rule set.

```

[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust

```

3. Configure a rule that matches packets and translates the source address to an address in the source NAT pool.

```

[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address 0.0.0.0/0

```

```
user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
```

4. Create a destination NAT pool.

```
[edit security nat destination]
user@host# set pool dst-nat-pool-1 address 10.1.1.200/32
```

5. Create a destination NAT rule set.

```
[edit security nat destination]
user@host# set rule-set rs1 from zone untrust
```

6. Configure a rule that matches packets and translates the destination address to the address in the destination NAT pool.

```
[edit security nat destination]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.100/32
user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
```

7. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.10 to 203.0.113.14
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.100
```

8. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

9. Configure an address in the global address book.

```
[edit security address-book global]
user@host# set address dst-nat-pool-1 10.1.1.200/32
```

10. Configure a security policy that allows traffic from the untrust zone to the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy dst-nat-pool-1-access match source-address any destination-address
dst-nat-pool-1 application any
user@host# set policy dst-nat-pool-1-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.10/32 to 203.0.113.14/32;
    }
  }
}
rule-set rs1 {
  to zone untrust;
  rule r1 {
    match {
      source-address 0.0.0.0/0;
      destination-address 0.0.0.0/0;
    }
    then {
      source-nat {
        pool {
          src-nat-pool-1;
        }
      }
    }
  }
}
```

```

    }
  }
}
destination {
  pool dst-nat-pool-1 {
    address 10.1.1.200/32;
  }
  rule-set rs1 {
    from zone untrust;
    rule r1 {
      match {
        destination-address 203.0.113.100/32;
      }
      then {
        destination-nat pool dst-nat-pool-1;
      }
    }
  }
}
proxy-arp {
  interface ge-0/0/0.0 {
    address {
      203.0.113.10/32 to 203.0.113.14/32;
      203.0.113.100/32;
    }
  }
}
user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
    }
  }
  policy internet-access {
    then {
      permit;
    }
  }
}

```

```

}
  from-zone untrust to-zone trust {
    policy dst-nat-pool-1-access {
      match {
        source-address any;
        destination-address dst-nat-pool-1;
        application any;
      }
      then {
        permit;
      }
    }
  }
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Source NAT Pool Usage | 69](#)
- [Verifying Source NAT Rule Usage | 70](#)
- [Verifying Destination NAT Pool Usage | 70](#)
- [Verifying Destination NAT Rule Usage | 70](#)
- [Verifying NAT Application to Traffic | 70](#)

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

Verifying Source NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the source NAT pool.

Action

From operational mode, enter the `show security nat source pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying Destination NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the destination NAT pool.

Action

From operational mode, enter the `show security nat destination pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose

Verify that there is traffic matching the destination NAT rule.

Action

From operational mode, enter the `show security nat destination rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Understanding Source NAT Rules

Source NAT rules specify two layers of match conditions:

- **Traffic direction**—Allows you to specify combinations of from interface, from zone, or from routing-instance and to interface, to zone, or to routing-instance. You cannot configure the same from and to contexts for different rule sets.
- **Packet information**—Can be source and destination IP addresses or subnets, source port numbers or port ranges, destination port numbers or port ranges, protocols, or applications.

For all ALG traffic, except FTP, we recommend that you not use the source-port rule option. Data session creation can fail if this option is used because the IP address and the source port value, which is a random value, might not match the rule.

In addition, we recommend that you not use the destination-port option or the application option as matching conditions for ALG traffic. If these options are used, translation may fail because the port value in the application payload might not match the port value in the IP address.

If multiple source NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 to zone 2 and rule B specifies traffic from zone 1 to interface ge-0/0/0, rule B is used to perform source NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match.

The actions you can specify for a source NAT rule are:

- **off**—Do not perform source NAT.
- **pool**—Use the specified user-defined address pool to perform source NAT.
- **interface**—Use the egress interface's IP address to perform source NAT.

Source NAT rules are applied to traffic in the first packet that is processed for the flow or in the fast path for the ALG. Source NAT rules are processed after static NAT rules, destination NAT rules, and reverse mapping of static NAT rules and after route and security policy lookup.

When zones are not configured under rule-set and when active source NAT is configured with missing mandatory statement "from" then, the following message is displayed when performing commit **"Missing mandatory statement: 'from' error: configuration check-out failed"** and the configuration check-out fails.

Example: Configuring Source NAT with Multiple Rules

IN THIS SECTION

- Requirements | 72
- Overview | 72
- Configuration | 75
- Verification | 80

This example describes how to configure source NAT mappings with multiple rules.

Requirements

Before you begin:


- Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
- Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 5 on page 73](#), the following translations are performed on the Juniper Networks security device for the source NAT mapping for traffic from the trust zone to the untrust zones:

- The source IP address in packets sent by the 10.1.1.0/24 and 10.1.2.0/24 subnets to any address in the untrust zone is translated to a public address in the range from 192.0.2.1 to 192.0.2.24 with port translation.
- The source IP address in packets sent by the 192.168.1.0/24 subnet to any address in the untrust zone is translated to a public address in the range from 192.0.2.100 to 192.0.2.249 with no port translation.
- The source IP address in packets sent by the 192.168.1.250/32 host device is not translated.

Figure 5: Source NAT with Multiple Translation Rules



Original Source IP	Translated Source IP
10.1.1.0/24, 10.1.2.0/24	192.0.2.1 – 192.0.2.24 (w/port translation)
192.168.1.0/24	192.0.2.100 - 192.0.2.249 (no port translation)
192.168.1.250/32	(no source NAT translation)

This example describes the following configurations:

- Source NAT pool `src-nat-pool-1` that contains the IP address range 192.0.2.1 through 192.0.2.24.
- Source NAT pool `src-nat-pool-2` that contains the IP address range 192.0.2.100 through 192.0.2.249, with port address translation disabled.

NOTE: When port address translation is disabled, the number of translations that the source NAT pool can support concurrently is limited to the number of addresses in the pool, unless the `address-shared` option is enabled. Packets are dropped if there are no addresses available in the source NAT pool. You can optionally specify an overflow pool from which IP addresses and port numbers are allocated when there are no addresses available in the original source NAT pool.

- Source NAT rule set `rs1` to match packets from the trust zone to the untrust zone. Rule set `rs1` contains multiple rules:
 - Rule `r1` to match packets with a source IP address in either the 10.1.1.0/24 or 10.1.2.0/24 subnets. For matching packets, the source address is translated to an IP address in the `src-nat-pool-1` pool.
 - Rule `r2` to match packets with a source IP address of 192.168.1.250/32. For matching packets, there is no NAT translation performed.
 - Rule `r3` to match packets with a source IP address in the 192.168.1.0/24 subnet. For matching packets, the source address is translated to an IP address in the `src-nat-pool-2` pool.

NOTE: The order of rules in a rule set is important, as the first rule in the rule set that matches the traffic is used. Therefore, rule `r2` to match a specific IP address must be placed before rule `r3` that matches the subnet on which the device is located.

- Proxy ARP for the addresses 192.0.2.1 through 192.0.2.24 and 192.0.2.100 through 192.0.2.249 on interface `ge-0/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policies to permit traffic from the trust zone to the untrust zone.

On SRX4600 devices, when you configure source NAT rule or pool with rule name or pool name as interface or service-set you will receive the following error message: **syntax error, expecting <data>**.

- If there is a source NAT rule named `interface`, the rule cannot be viewed using the `show security nat source rule interface` command.

- If there is a source NAT rule named `service-set`, the rule cannot be viewed using the `show security nat source rule service-set` command.
- If there is a source NAT pool named `interface`, the pool cannot be viewed using the `show security nat source pool interface` command.
- If there is a source NAT pool named `service-set`, the pool cannot be viewed using the `show security nat source pool service-set` command.
- If there is a source NAT pool named `interface`, the paired-address cannot be viewed using the `show security nat source paired-address pool-name interface` command.
- If there is a source NAT pool named `service-set`, the paired-address cannot be viewed using the `show security nat source paired-address pool-name service-set` command.

Configuration

IN THIS SECTION

- [Procedure | 75](#)

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-pool-1 address 192.0.2.1/32 to 192.0.2.24/32
set security nat source pool src-nat-pool-2 address 192.0.2.100/32 to 192.0.2.249/32
set security nat source pool src-nat-pool-2 port no-translation
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 10.1.1.0/24
set security nat source rule-set rs1 rule r1 match source-address 10.1.2.0/24
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat source rule-set rs1 rule r2 match source-address 192.168.1.250/32
set security nat source rule-set rs1 rule r2 match destination-address 0.0.0.0/0
```

```

set security nat source rule-set rs1 rule r2 then source-nat off
set security nat source rule-set rs1 rule r3 match source-address 192.168.1.0/24
set security nat source rule-set rs1 rule r3 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r3 then source-nat pool src-nat-pool-2
set security nat proxy-arp interface ge-0/0/0.0 address 192.0.2.1/32 to 192.0.2.24/32
set security nat proxy-arp interface ge-0/0/0.0 address 192.0.2.100/32 to 192.0.2.249/32
set security policies from-zone trust to-zone untrust policy internet-access match source-
address any
set security policies from-zone trust to-zone untrust policy internet-access match destination-
address any
set security policies from-zone trust to-zone untrust policy internet-access match application
any
set security policies from-zone trust to-zone untrust policy internet-access then permit

```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure multiple source NAT rules in a rule set:

1. Create a source NAT pool.

```

[edit security nat source]
user@host# set pool src-nat-pool-1 address 192.0.2.1 to 192.0.2.24

```

2. Create a source NAT pool with no port translation.

```

[edit security nat source]
user@host# set pool src-nat-pool-2 address 192.0.2.100 to 192.0.2.249
user@host# set pool src-nat-pool-2 port no-translation

```

NOTE: To configure an overflow pool for src-nat-pool-2 using the egress interface:

```

[edit security nat source]
user@host# set pool src-nat-pool-2 overflow-pool interface

```

3. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust
```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

```
[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address [10.1.1.0/24 10.1.2.0/24]
user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
```

5. Configure a rule to match packets for which the source address is not translated.

```
[edit security nat source]
user@host# set rule-set rs1 rule r2 match source-address 192.168.1.250/32
user@host# set rule-set rs1 rule r2 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r2 then source-nat off
```

6. Configure a rule to match packets and translate the source address to an address in the pool with no port translation.

```
[edit security nat source]
user@host# set rule-set rs1 rule r3 match source-address 192.168.1.0/24
user@host# set rule-set rs1 rule r3 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r3 then source-nat pool src-nat-pool-2
```

7. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 192.0.2.1 to 192.0.2.24
user@host# set proxy-arp interface ge-0/0/0.0 address 192.0.2.100 to 192.0.2.249
```


8. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
    pool src-nat-pool-1 {
        address {
            192.0.2.1/32 to 192.0.2.24/32;
        }
    }
    pool src-nat-pool-2 {
        address {
            192.0.2.100/32 to 192.0.2.249/32;
        }
        port no-translation;
    }
    rule-set rs1 {
        from zone trust;
        to zone untrust;
        rule r1 {
            match {
                source-address [ 10.1.1.0/24 10.1.2.0/24 ];
                destination-address 0.0.0.0/0;
            }
            then {
                source-nat {
                    pool {
                        src-nat-pool-1;
                    }
                }
            }
        }
    }
}
```



```
        application any;
    }
    then {
        permit;
    }
}
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Source NAT Pool Usage | 80](#)
- [Verifying Source NAT Rule Usage | 80](#)
- [Verifying NAT Application to Traffic | 81](#)

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

Verifying Source NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the source NAT pool.

Action

From operational mode, enter the `show security nat source pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Understanding Source NAT Pools

A NAT pool is a user-defined set of IP addresses that are used for translation. Unlike static NAT, where there is a one-to-one mapping that includes destination IP address translation in one direction and source IP address translation in the reverse direction, with source NAT, you translate the original source IP address to an IP address in the address pool.

For source Network Address Translation (NAT) address pools, specify the following:

- Name of the source NAT address pool.
- Up to eight address or address ranges.

NOTE: Do not overlap NAT addresses for source NAT, destination NAT, and static NAT within one routing instance.

- Routing instance—Routing instance to which the pool belongs (the default is the main **inet.0** routing instance).
- Port —The Port Address Translation (PAT) for a source pool. By default, PAT is performed with source NAT. If you specify the **no-translation** option, the number of hosts that the source NAT pool can support is limited to the number of addresses in the pool. If you specify **block-allocation**, a block of ports is allocated for translation, instead of individual ports being allocated. If you specify **deterministic**, an incoming (source) IP address and port always map to the specific destination address and port block, based on predefined, deterministic NAT algorithm. If you specify **port-overloading**, you

can configure the port overloading capacity in source NAT. If you specify `range`, you can provide the port number range attached to each address in the pool, and the twin port range for source NAT pools.

- **Overflow pool (optional)**—Packets are dropped if there are no addresses available in the designated source NAT pool. To prevent that from happening when the **port no-translation** option is configured, you can specify an overflow pool. Once addresses from the original source NAT pool are exhausted, IP addresses and port numbers are allocated from the overflow pool. A user-defined source NAT pool or an egress interface can be used as the overflow pool. (When the overflow pool is used, the pool ID is returned with the address.)
- **IP address shifting (optional)**—A range of original source IP addresses can be mapped to another range of IP addresses, or to a single IP address, by shifting the IP addresses. Specify the **host-address-base** option with the base address of the original source IP address range.
- **Address sharing (optional)**—Multiple internal IP addresses can be mapped to the same external IP address. This option can be used only when the source NAT pool is configured with no port translation. Specify the `address-shared` option when a source NAT pool has few external IP addresses available, or only one external IP address. With a many-to-one mapping, use of this option increases NAT resources and improves traffic.
- **Address pooling (optional)**— Address pooling can be configured as `paired` or `no-paired`. Specify `address-pooling paired` for applications that require all sessions associated with one internal IP address to be mapped to the same external IP address for the duration of a session. This differs from the `persistent-address` option, in which the same internal address is translated to the same external address every time. Specify `address-pooling no-paired` for applications that can be assigned IP addresses in a round-robin fashion. If either `address-pooling paired` or `address-pooling no-paired` is configured for a source NAT pool with PAT, the `persistent-address` option is disabled. If `address-shared` is configured on a source NAT pool without PAT, then the `persistent-address` option is enabled. Both `address-shared` and `address-pooling paired` can be configured on the same source NAT pool without PAT.
- **Pool utilization alarm (optional)**— When the **raise-threshold** option is configured for source NAT, an SNMP trap is triggered if the source NAT pool utilization rises above this threshold. If the optional **clear-threshold** option is configured, an SNMP trap is triggered if the source NAT pool utilization drops below this threshold. If **clear-threshold** is not configured, it is set by default to 80 percent of the **raise-threshold** value.

You can use the **show security nat resource usage source pool** command to view address use in a source NAT pool without PAT, and to view port use in a source NAT pool with PAT.

Understanding Source NAT Pool Capacities

Maximum capacities for source pools and IP addresses on SRX300, SRX320, SRX340, SRX345 and SRX650 devices are as follows:

Pool/PAT Maximum Address Capacity	SRX300 SRX320	SRX340 SRX345	SRX650
Source NAT pools	1024	2048	1024
IP addresses supporting port translation	1024	2048	1024
PAT port number	64M	64M	64M

Maximum capacities for source pools and IP addresses on SRX1400, SRX1500, SRX3400, SRX3600, SRX4100, SRX4200, SRX5400, SRX5600, and SRX5800 devices are as follows:

Pool/PAT Maximum Address Capacity	SRX1400 SRX1500	SRX3400 SRX3600	SRX4100 SRX4200	SRX5400 SRX5600 SRX5800
Source NAT pools	8192	10,240	10,240	12,288
IP addresses supporting port translation	8192	12,288	12,288	1M
PAT port number	256M	384M	384M	384M

NOTE: In Release 12.3X48-D40, and in Release 15.1X49-D60 and later releases, you can increase the source NAT port capacity to 2.4G on SRX5400, SRX5600, and SRX5800 devices with next-generation Services Processing Cards (SPCs) using the port-scaling-enlargement statement at the [edit security nat source] hierarchy level supported .

NOTE: Platform support depends on the Junos OS release in your installation.

Increasing the total number of IP addresses used for source NAT, either by increasing the number of pools in the configuration and/or by increasing the capacity or IP-addresses per pool, consumes memory needed for port allocation. When source NAT pool and IP address limits are reached, port ranges should be reassigned. That is, the number of ports for each IP address should be decreased when the number of IP addresses and source NAT pools is increased. This ensures NAT does not consume too much memory.

For example, in a source NAT pool for SRX5000 devices, when the number of IP addresses supporting port translation reaches the limit of 1M, the total number of PAT ports is 64G, which exceeds the 384M limitation. This is because, by default, each IP address supports 64,512 ports. To ensure that PAT port numbers are within capacity, the port range for each IP needs to be configured to decrease the total number of PAT ports.

Use the `range` and `range twin-port` options at the `[edit security nat source pool port]` hierarchy level to assign a new port range or twin port range for a specific pool. Use the `pool-default-port-range` and the `pool-default-twin-port-range` options at the `[edit security nat source]` hierarchy level to specify the global default port range or twin port range for all source NAT pools.

Configuring port overloading should also be done carefully when source NAT pools are increased.

For a source pool with PAT in range (63,488 through 65,535), two ports are allocated at one time for RTP/RTCP applications, such as SIP, H.323, and RTSP. In these scenarios, each IP address supports PAT, occupying 2048 ports (63,488 through 65,535) for ALG module use.

Understanding Persistent Addresses for Source NAT Pools

By default, port address translation is performed with source NAT. However, an original source address may not be translated to the same IP address for different traffic that originates from the same host. The source NAT `address-persistent` option ensures that the same IP address is assigned from the source NAT pool to a specific host for multiple concurrent sessions.

This option differs from the `address-pooling paired` option, where the internal address is mapped to an external address within the pool on a first-come, first-served basis, and might be mapped to a different external address for each session.

Example: Configuring Capacity for Source NAT Pools with PAT

IN THIS SECTION

- [Requirements | 85](#)
- [Overview | 85](#)
- [Configuration | 85](#)
- [Verification | 87](#)

This example describes how to configure the capacity of source NAT pools with Port Address Translation (PAT) if a default port range is not set or you want to override it. Translations are set for each IP address. When the source pool is increased, ports should be reassigned if the current port number exceeds limitations.

Requirements

Before you begin:

1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example shows how to configure a PAT pool of 2048 IP addresses with 32,000 ports for each IP address.

Configuration

IN THIS SECTION

- [Procedure | 86](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
[edit security nat source]
set pool src-nat-pat-addr address 192.168.0.0/32 to 192.168.3.255/32
set pool src-nat-pat-addr address 192.168.4.0/32 to 192.168.7.255/32
set pool-default-port-range 2001
set pool-default-port-range to 32720
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure capacity for a source NAT pool with PAT:

1. Specify a source NAT pool with PAT and an IP address range.

```
[edit security nat source]
user@host# set pool src-nat-pat-addr address 192.168.0.0/32 to 192.168.3.255/32
user@host# set pool src-nat-pat-addr address 192.168.4.0/32 to 192.168.7.255/32
```

2. Specify a default port range for the source pool.

```
[edit security nat source]
user@host# set pool-default-port-range 2001
user@host# set pool-default-port-range to 32720
```

Results

From configuration mode, confirm your configuration by entering the `show security nat-source-summary` command. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@host> run show security nat source summary
Total port number usage for port translation pool: 16515072
Maximum port number for port translation pool: 134217728
Total pools: 1

Pool Address Routing PAT Total Name Range Instance Address pool2 203.0.113.1 - 203.0.113.3
default yes 2048
Name Range Instance Address
pool1 198.51.100.0 - 198.51.100.255 default yes 256

Total rules: 1
Rule name Rule set From To Action
rule 1 ruleset1 ge-2/2/2.0 ge-2/2/3.0 pool1
rule 1 ge-2/2/4.0 ge-2/2/5.0
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Capacity of Source NAT Pools | 87](#)

Verifying Capacity of Source NAT Pools

Purpose

View port and pool information. Port limitations are automatically checked, so the configuration will not be committed if port limitations are exceeded.

Action

From operational mode, enter the `show security nat source summary` command to view port and pool details.

Understanding Source NAT Pools with Address Pooling

When a host initiates several sessions that match a policy that requires NAT, and is assigned an IP address from a source pool that has port address translation enabled, a different source IP address is used for each session.

Because some applications require the same source IP address for each session, you can use the address-pooling paired feature to enable all sessions associated with one internal IP address to map to the same external IP address for the duration of the sessions. When the sessions end, the mapping between the internal IP address and the external IP address ceases. The next time the host initiates a session, a different IP address from the pool might be assigned to it.

This differs from the source NAT address-persistent feature, which keeps the mapping static; the same internal IP address is mapped to the same external IP address every time. It also differs from the address-persistent feature in that address-pooling paired is configured for a specific pool. The address-persistent feature is a global configuration that applies to all source pools.

Understanding Source NAT Pools with Address Shifting

The match conditions for a source NAT rule set do not allow you to specify an address range; only address prefixes may be specified in a rule. When configuring a source NAT pool, you can specify the host-base-address option; this option specifies the IP address where the original source IP address range begins.

The range of original source IP addresses that are translated is determined by the number of addresses in the source NAT pool. For example, if the source NAT pool contains a range of ten IP addresses, then up to ten original source IP addresses can be translated, starting with a specified base address. This type of translation is one-to-one, static, and without port address translation.

The match condition in a source NAT rule may define a larger address range than that specified in the source NAT pool. For example, a match condition might specify an address prefix that contains 256 addresses, but the source NAT pool might contain a range of only a few IP addresses, or only one IP address. A packet's source IP address can match a source NAT rule, but if the source IP address is not within the address range specified in the source NAT pool, the source IP address is not translated.

Example: Configuring Source NAT Pools with Address Shifting

IN THIS SECTION

- Requirements | 89
- Overview | 90
- Configuration | 92
- Verification | 95

This example describes how to configure a source NAT mapping of a private address range to public addresses, with optional address shifting. This mapping is one-to-one between the original source IP addresses and translated IP addresses.

NOTE: The match conditions for a source NAT rule set do not allow you to specify an address range; only address prefixes may be specified in a rule. When configuring a source NAT pool, you can specify the `host-base-address` option; this option specifies the IP address where the original source IP address range begins, and disables port translation.

The range of original source IP addresses that are translated is determined by the number of addresses in the source NAT pool. For example, if the source NAT pool contains a range of ten IP addresses, then up to ten original source IP addresses can be translated, starting with a specified base address.

The match condition in a source NAT rule may define a larger address range than that specified in the source NAT pool. For example, a match condition might specify an address prefix that contains 256 addresses, but the source NAT pool contains a range of only ten IP addresses. A packet's source IP address can match a source NAT rule, but if the source IP address is not within the address range specified in the source NAT pool, the source IP address is not translated.

Requirements


Before you begin:

- Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
- Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 6 on page 91](#), a range of private addresses in the trust zone is mapped to a range of public addresses in the untrust zone. For packets sent from the trust zone to the untrust zone, a source IP address in the range of 192.168.1.10/32 through 192.168.1.20/32 is translated to a public address in the range of 203.0.113.30/32 through 203.0.113.40/32.

Figure 6: Source NAT with Address Shifting



Original Source IP	Translated Source IP
192.168.1.10/32 - 192.168.1.20/32	203.0.113.30/32 - 203.0.113.40/32

This example describes the following configurations:

- Source NAT pool `src-nat-pool-1` that contains the IP address range 203.0.113.30/32 through 203.0.113.40/32. For this pool, the beginning of the original source IP address range is 192.168.1.10/32 and is specified with the `host-address-base` option.
- Source NAT rule set `rs1` with rule `r1` to match packets from the trust zone to the untrust zone with a source IP address in the 192.168.1.0/24 subnet. For matching packets that fall within the source IP address range specified by the `src-nat-pool-1` configuration, the source address is translated to the IP address in `src-nat-pool-1` pool.
- Proxy ARP for the addresses 203.0.113.30/32 through 203.0.113.40/32 on interface `ge-0/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

IN THIS SECTION

- [Procedure | 92](#)

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.30/32 to 203.0.113.40/32
set security nat source pool src-nat-pool-1 host-address-base 192.168.1.10/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 192.168.1.0/24
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.30/32 to 203.0.113.40/32
set security policies from-zone trust to-zone untrust policy internet-access match source-address any
```

```
set security policies from-zone trust to-zone untrust policy internet-access match destination-
address any
set security policies from-zone trust to-zone untrust policy internet-access match application
any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping with address shifting:

1. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool src-nat-pool-1 address 203.0.113.30/32 to 203.0.113.40/32
```

2. Specify the beginning of the original source IP address range.

```
[edit security nat source]
user@host# set pool src-nat-pool-1 host-address-base 192.168.1.10/32
```

3. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust
```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

```
[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address 192.168.1.0/24
user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
```


5. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.30/32 to 203.0.113.40/32
```

6. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.30/32 to 203.0.113.40/32;
    }
    host-address-base 192.168.1.10/32;
  }
  rule-set rs1 {
    from zone trust;
    to zone untrust;
    rule r1 {
      match {
        source-address 192.168.1.0/24;
      }
      then {
        source-nat {
          pool {
            src-nat-pool-1;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

proxy-arp {
  interface ge-0/0/0.0 {
    address {
      203.0.113.30/32 to 203.0.113.40/32;
    }
  }
}

user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
}
}
```

Verifying Source NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the source NAT pool.

Action

From operational mode, enter the `show security nat source pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Understanding Source NAT Pools with PAT

Using the source pool with Port Address Translation (PAT), Junos OS translates both the source IP address and the port number of the packets. When PAT is used, multiple hosts can share the same IP address.

Junos OS maintains a list of assigned port numbers to distinguish what session belongs to which host. When PAT is enabled, up to 63,488 hosts can share a single IP address. Each source pool can contain

multiple IP addresses, multiple IP address ranges, or both. For a source pool with PAT, Junos OS may assign different addresses to a single host for different concurrent sessions, unless the source pool or Junos OS has the persistent address feature or the paired address pooling feature enabled.

For interface source pool and source pool with PAT, range (1024, 65535) is available for port number mapping per IP address. Within range (1024, 63487) one port is allocated at a time, for a total of 62,464 ports. In range (63488, 65535), two ports are allocated at a time for RTP/RTCP applications such as SIP, H.323, and RTSP, for a total of 2,048 ports.

When a host initiates several sessions that match a policy that requires network address translation and is assigned an address from a source pool that has PAT enabled, the device assigns a different source IP address for each session. Such random address assignment can be problematic for services that create multiple sessions that require the same source IP address for each session. For example, it is important to have the same IP address for multiple sessions when using the AOL Instant Message (AIM) client.

To ensure that the router assigns the same IP address from a source pool to a host for multiple concurrent sessions, you can enable a persistent IP address per router. To ensure that the device assigns the same IP address from a source pool to a host for the duration of a single session, you can enable paired address pooling.

Example: Configuring Source NAT for Multiple Addresses with PAT

IN THIS SECTION

- [Requirements | 97](#)
- [Overview | 98](#)
- [Configuration | 100](#)
- [Verification | 103](#)

This example describes how to configure a source NAT mapping of a private address block to a smaller public address block using port address translation.

Requirements

Before you begin:


1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 7 on page 99](#), the source IP address in packets sent from the trust zone to the untrust zone is mapped to a smaller block of public addresses in the range from 203.0.113.1/32 through 203.0.113.24/32. Because the size of the source NAT address pool is smaller than the number of potential addresses that might need to be translated, port address translation is used.

NOTE: Port address translation includes a source port number with the source IP address mapping. This allows multiple addresses on a private network to map to a smaller number of public IP addresses. Port address translation is enabled by default for source NAT pools.

Figure 7: Source NAT Multiple Addresses with PAT



Original Source IP	Translated Source IP
10.1.1.0/24	203.0.113.1 (with port address translation)
10.1.2.0/24	
192.168.1.0/24	

This example describes the following configurations:

- Source NAT pool `src-nat-pool-1` that contains the IP address range 203.0.113.1/32 through 203.0.113.24/32.
- Source NAT rule set `rs1` to match all packets from the trust zone to the untrust zone. For matching packets, the source IP address is translated to an IP address in the `src-nat-pool-1` pool.
- Proxy ARP for the addresses 203.0.113.1/32 through 203.0.113.24/32 on interface `ge-0/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

IN THIS SECTION

- [Procedure](#) | 100

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.1/32 to 203.0.113.24/32
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 10.1.1.0/24
set security nat source rule-set rs1 rule r1 match source-address 10.1.2.0/24
set security nat source rule-set rs1 rule r1 match source-address 192.168.1.0/24
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.24/32
set security policies from-zone trust to-zone untrust policy internet-access match source-address any
set security policies from-zone trust to-zone untrust policy internet-access match destination-
```

```

address any
set security policies from-zone trust to-zone untrust policy internet-access match application
any
set security policies from-zone trust to-zone untrust policy internet-access then permit

```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping from a private address block to a smaller public address block using PAT:

1. Create a source NAT pool.

```

[edit security nat source]
user@host# set pool src-nat-pool-1 address 203.0.113.1 to 203.0.113.24

```

2. Create a source NAT rule set.

```

[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust

```

3. Configure a rule that matches packets and translates the source address to an address in the pool.

```

[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address [10.1.1.0/24 10.1.2.0/24
192.168.1.0/24]
user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1

```

4. Configure proxy ARP.

```

[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1 to 203.0.113.24

```


5. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
    pool src-nat-pool-1 {
        address {
            203.0.113.1/32 to 203.0.113.24/32;
        }
    }
    rule-set rs1 {
        from zone trust;
        to zone untrust;
        rule r1 {
            match {
                source-address [10.1.1.0/24 10.1.2.0/24 192.168.1.0/24];
                destination-address 0.0.0.0/0;
            }
            then {
                source-nat {
                    pool {
                        src-nat-pool-1;
                    }
                }
            }
        }
    }
}
proxy-arp {
```

```
interface ge-0/0/0.0 {  
  address {  
    203.0.113.1/32 to 203.0.113.24/32;  
  }  
}  
}  
user@host# show security policies  
from-zone trust to-zone untrust {  
  policy internet-access {  
    match {  
      source-address any;  
      destination-address any;  
      application any;  
    }  
    then {  
      permit;  
    }  
  }  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Source NAT Pool Usage | 103](#)
- [Verifying Source NAT Rule Usage | 104](#)
- [Verifying NAT Application to Traffic | 104](#)

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

Verifying Source NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the source NAT pool.

Action

From operational mode, enter the `show security nat source pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage**Purpose**

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic**Purpose**

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Understanding Source NAT Pools Without PAT

When you define a source pool, Junos OS enables PAT by default. To disable PAT, you must specify no port translation when you are defining a source pool.

When using a source pool without PAT, Junos OS performs source Network Address Translation for the IP address without performing PAT for the source port number. For applications that require that a particular source port number remain fixed, you must use source pool without PAT.

The source pool can contain multiple IP addresses, multiple IP address ranges, or both. For source pool without PAT, Junos OS assigns one translated source address to the same host for all its concurrent sessions unless the address-pooling no-paired option is enabled.

The number of hosts that a source NAT pool without PAT can support is limited to the number of addresses in the pool. When you have a pool with a single IP address, only one host can be supported,

and traffic from other hosts is blocked because there are no resources available. If a single IP address is configured for a source NAT pool without PAT when NAT resource assignment is not in active-backup mode in a chassis cluster, traffic through node 1 will be blocked.

Pool utilization for each source pool without PAT is computed. You can turn on pool utilization alarm by configuring alarm thresholds. An SNMP trap is triggered every time pool utilization rises above a threshold and goes below a threshold.

NOTE: If a static NAT rule is for one-to-one IP translation, avoid dividing the rule into a destination rule and a source rule when source no-pat pool without address sharing is used. If you choose to divide the rule, you will then have to use source pat-pool with single IP or source no-pat pool with multiple IP.

Example: Configuring a Single IP Address in a Source NAT Pool Without PAT

IN THIS SECTION

- [Requirements | 106](#)
- [Overview | 106](#)
- [Configuration | 106](#)
- [Verification | 109](#)

This example describes how to configure a private address block to a single public address in a source NAT pool without Port Address Translation.

NOTE: PAT is enabled by default for source NAT pools. When PAT is disabled, the number of translations that the source NAT pool can concurrently support is limited to the number of addresses in the pool. Packets are dropped if there are no addresses available in the source NAT pool. However, using the address-shared option, you can map more than one private IP address to a single public IP address as long as the traffic is from different source ports.

Requirements

Before you begin:

1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. The source IP address of packets sent from the trust zone to the untrust zone are mapped to a single public address.

This example describes the following configurations:

- Source NAT pool `src-nat-pool-1` that contains the IP address `203.0.113.1/30`. The port `no-translation` option and the address `shared` option are specified for the pool.
- Source NAT rule set `rs1` to match all packets from the trust zone to the untrust zone. For matching packets, the source IP address is translated to an IP address in the `src-nat-pool-1` pool.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

IN THIS SECTION

- [Procedure | 106](#)

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.1/30
set security nat source pool src-nat-pool-1 port no-translation
set security nat source pool src-nat-pool-1 address-shared
set security nat source rule-set rs1 from zone trust
```

```
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule1 match source address 192.0.2.0/24
set security nat source rule-set rs1 rule r1 then source src-nat-pool-1
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping from a private address block to a single public address without PAT:

1. Create a source NAT pool with a single IP address for the shared address.

```
[edit security nat source]
user@host# set pool src-nat-pool-1 address 203.0.113.1/30
```

Specify the port no-translation option.

```
[edit security nat source]
user@host# set pool src-nat-pool-1 port no-translation
```

2. Specify the address-shared option.

```
[edit security nat source]
user@host# set pool pool-src-nat-pool-1 address-shared
```

3. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust
```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

```
[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address 192.0.2.0/24
user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
```

5. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat source pool` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
    pool src-nat-pool-1 {
        address {
            203.0.113.1/30
        }
        port no-translation;
    }
    address-shared;
    rule-set rs1 {
        from zone trust;
        to zone untrust;
        rule r1 {
            match {
                source-address [192.0.2.0/24]
            }
            then {
                source-nat {
                    pool {
                        src-nat-pool-1;
                    }
                }
            }
        }
    }
}
```

```
user@host# show security policies
from-zone trust to-zone untrust {
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Shared Address | 109](#)
- [Verifying Shared Address Application to Traffic | 110](#)

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

Verifying Shared Address

Purpose

Verify that two internal IP addresses, with different source ports, share one external IP address.

Action

From operational mode, enter the `show security nat source pool` command. View the **Address assignment** field to verify that it is shared.

Verifying Shared Address Application to Traffic

Purpose

Verify that two sessions are using the same IP address.

Action

From operational mode, enter the `show security flow session` command.

Example: Configuring Multiple Addresses in a Source NAT Pool Without PAT

IN THIS SECTION

- [Requirements | 110](#)
- [Overview | 111](#)
- [Configuration | 113](#)
- [Verification | 116](#)

This example describes how to configure a source NAT mapping of a private address block to a smaller public address block without port address translation.

NOTE: Port address translation is enabled by default for source NAT pools. When port address translation is disabled, the number of translations that the source NAT pool can concurrently support is limited to the number of addresses in the pool. Packets are dropped if there are no addresses available in the source NAT pool. You can optionally specify an overflow pool from which IP addresses and port numbers are allocated when there are no addresses available in the original source NAT pool.

Requirements

Before you begin:


1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).

2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 8 on page 112](#), the source IP address in packets sent from the trust zone to the untrust zone is mapped to a smaller block of public addresses in the range from 203.0.113.1/32 through 203.0.113.24/32.

Figure 8: Source NAT Multiple Addresses Without PAT



Original Source IP	Translated Source IP
10.1.1.0/24	203.0.113.1 (no port address translation)
10.1.2.0/24	
192.168.1.0/24	

This example describes the following configurations:

- Source NAT pool `src-nat-pool-1` that contains the IP address range `203.0.113.1/32` through `203.0.113.24/32`. The port `no-translation` option is specified for the pool.
- Source NAT rule set `rs1` to match all packets from the trust zone to the untrust zone. For matching packets, the source IP address is translated to an IP address in the `src-nat-pool-1` pool.
- Proxy ARP for the addresses `203.0.113.1/32` through `203.0.113.24/32` on interface `ge-0/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policies to permit traffic from the trust zone to the untrust zone.

Configuration

IN THIS SECTION

- [Procedure](#) | 113

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-pool-1 address 203.0.113.1/32 to 203.0.113.24/32
set security nat source pool src-nat-pool-1 port no-translation
set security nat source rule-set rs1 from zone trust
set security nat source rule-set rs1 to zone untrust
set security nat source rule-set rs1 rule r1 match source-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 match destination-address 0.0.0.0/0
set security nat source rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.24/32
set security policies from-zone trust to-zone untrust policy internet-access match source-
address any
set security policies from-zone trust to-zone untrust policy internet-access match destination-
address any
```

```
set security policies from-zone trust to-zone untrust policy internet-access match application
any
set security policies from-zone trust to-zone untrust policy internet-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a source NAT mapping from a private address block to a smaller public address block without PAT:

1. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool src-nat-pool-1 address 203.0.113.1 to 203.0.113.24
```

2. Specify the port no-translation option.

```
[edit security nat source]
user@host# set pool src-nat-pool-1 port no-translation
```

3. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs1 from zone trust
user@host# set rule-set rs1 to zone untrust
```

4. Configure a rule that matches packets and translates the source address to an address in the pool.

```
[edit security nat source]
user@host# set rule-set rs1 rule r1 match source-address 0.0.0.0/0
user@host# set rule-set rs1 rule r1 match destination-address 0.0.0.0/0
user@host# set rule-set rs1 rule r1 then source-nat pool src-nat-pool-1
```

5. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1 to 203.0.113.24
```

6. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-pool-1 {
    address {
      203.0.113.1/32 to 203.0.113.24/32;
    }
    port no-translation;
  }
  rule-set rs1 {
    from zone trust;
    to zone untrust;
    rule r1 {
      match {
        source-address 0.0.0.0/0;
        destination-address 0.0.0.0/0;
      }
      then {
        source-nat {
          pool {
            src-nat-pool-1;
          }
        }
      }
    }
  }
}
```


Verifying Source NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the source NAT pool.

Action

From operational mode, enter the `show security nat source pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Understanding Shared Addresses in Source NAT Pools without PAT

Source NAT pools with no port address translation perform static, one-to-one mappings from one source IP address to one external IP address. When there is only one external IP address, or very few available in a source no-pat pool, the address-shared option enables you to map many source IP addresses to one external IP address as long as the traffic comes from different source ports.

For example, if there is a source NAT pool with no port translation containing only two IP addresses, IP 1 and IP 2, when a packet arrives from

1. Source IP 1, port 1, it is translated to IP 1, port 1.
2. Source IP 2, port 2, it is translated to IP 2, port 2.
3. Source IP 3, port 1, it is translated to IP 2, port 1. (It cannot be translated to IP 1 port 1 because that port is already used.

However, if another packet arrives from Source IP 3, port 1 for a different destination IP and port, it cannot be translated to IP 1, port 1 or IP 2, port 1 because port 1 is already used for both available IP addresses. The session will fail.

This option increases NAT resources and improves the possibility of setting up successful translated traffic. It cannot be used on source NAT pools with port address translation because address sharing is already their default behavior.

Understanding NAT Session Persistence

IN THIS SECTION

- [Limitations of NAT Session Persistence | 119](#)

Network Address Translation (NAT) session persistence provides a means to retain existing sessions, instead of clearing them, when there changes in the NAT configuration. If session persistence is enabled, the retained sessions continue to process and forward packets as time and resources are optimally used to rebuild the impacted sessions. Thus, packet forwarding does not stop even if the NAT configuration is changed for some or all sessions.

From Junos OS Release 18.3R1 onward, with the support for NAT session persistence, the Packet Forwarding Engine scans the sessions and decides whether to keep the sessions or clear the sessions. In releases before Junos OS Release 18.3R1, the NAT sessions are cleared if there is a change in the NAT configuration.

The Packet Forwarding Engine performs the following two types of scans to decide whether to retain or drop sessions:

- **Source NAT pool session persistence scan**—The Packet Forwarding Engine compares the existing session IP address with source pool address range. If the existing session IP address is in the specified source pool address range, the session is kept alive, otherwise the session is cleared.
- **Source NAT rule session persistence scan**—The Packet Forwarding Engine uses the rule ID to compare the source IP address, source port, destination IP address, and destination port between the old and new configurations. If the new and old configurations are the same, then the session is kept alive, otherwise the session is cleared.

NOTE:

- NAT session persistence is not supported for static NAT and destination NAT.
- NAT session persistence is not supported if the PAT pool is configured with the address persistent, address pooling paired, source address-persistent, port block allocation, port deterministic, persistent nat, and port overloading factor fields.

NAT session persistence is supported only for source NAT in the following scenarios:

- **Source pool**—Change in an address range in a Port Address Translation (PAT) pool.
- **Source rule**—Change in match conditions for the address book, application, destination IP address, destination port, source IP address, and destination port information.

To enable the NAT session persistence scanning, include the `session-persistence-scan` statement at the `[edit security nat source]` hierarchy level.

You can also configure a timeout value to retain the sessions for the specified time period by using the `set security nat source session-drop-hold-down` CLI command. The value of the `session-drop-hold-down` option ranges from 30 through 28,800 seconds (eight hours). The session expires after the configured timeout period.

Limitations of NAT Session Persistence

- When there is a change in IP addresses in the NAT source pool, the newly configured IP addresses are appended to the NAT source pool. After the NAT source pool is rebuilt, the new IP addresses are not the same as the existing IP addresses. The differences in the IP addresses in the NAT source pool impacts the round-robin mode of picking IP addresses from the NAT source pool.
- If the scan types identify sessions that will never be timed out (that is, the sessions for which the `session-drop-hold-down` value is not configured or is configured as 8 hours), then the Packet Forwarding Engine ignores those sessions, and the sessions are retained.

Configure Port Block Allocation Size

Before you begin:

- Understand the guidelines for configuring port block allocation. Read [Guidelines for Configuring Secured Port Block Allocation](#).

You can configure secured port block allocation, which allocates blocks of ports to a NAT subscriber. With port block allocation, we generate one syslog log per set of ports allocated for a subscriber. Use this procedure to configure the port block allocation size.

1. Configure the IPv4 addresses.

```
user@host# set security nat source pool root_src_v4_pat address 200.0.0.1/32 to 200.16.0.0/32
```

2. Configure the starting and ending port value.

```
user@host# set security nat source pool root_src_v4_pat port range 61044
user@host# set security nat source pool root_src_v4_pat port range to 63500
```

3. Configure the port block allocation size.

```
user@host# set security nat source pool root_src_v4_pat port block-allocation block-size 8
```

If you configure the port block allocation size lesser than 8 on SRX5400, SRX5600, and SRX5800, the system displays the warning message warning: To save system memory, the block size is recommended to be no less than 8.

Starting in Junos OS Release 20.3R1, you can configure the port block allocation size on SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600. To save system memory, the recommended port block allocation size is 64. If you configure the port block allocation size lesser than 64, the system displays the warning message warning: To save system memory, the block size is recommended to be no less than 64.

4. Configure the interim log interval time.

```
user@host#set security nat source pool root_src_v4_pat port block-allocation interim-logging-
interval 1800
```

5. Configure the last port block timeout value.

```
user@host#set security nat source pool root_src_v4_pat port block-allocation last-block-
recycle-timeout 120
```

6. Commit the configuration

```
user@host#commit
2020-05-14 19:56:33.758167 CST: Running FIPS Self-tests
Veriexec is not enforced, FIPS mode not available
2020-05-14 19:56:33.771303 CST: FIPS Self-tests Skipped
commit complete
```

7. Verify the output value for configured block-size.

```
user@host#run show security nat source pool all
```

```
Total pools: 1
Pool name      : root_src_v4_pat
Pool id       : 4
Routing instance : default
Port          : [61044, 63500]
Port overloading : 1
Address assignment : no-paired
Total addresses  : 1048576
Translation hits : 0
Port block size  : 8
Max blocks per host : 8
Active block timeout : 0
Last block recycle timeout : 0
Interim logging interval : 0
PBA block log    : Enable
Used/total port blocks: 0/321912832
```

Address range	Single Ports	Twin Ports
200.0.0.1 - 200.16.0.0	0	0
Total used ports :	0	0

Configuring the NAT Session Hold Timeout and NAT Session Persistence Scan

This configuration shows how to configure the NAT session hold timeout and NAT session persistence.

Configuring NAT Session Hold Timeout

The following configuration shows how to configure the NAT session hold timeout.

- To set the NAT session hold timeout period:

```
[edit security nat source]
user@host# set session-drop-hold-down time;
```

The value of the *time* variable ranges from 30 through 28,800 seconds (eight hours). The session expires after the configured timeout period.

Results

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

```
[edit]
user@host# show security
  nat {
    source {
      session-drop-hold-down 28800;
    }
  }
```

Configuring NAT Session Persistence Scan

The following configuration shows how to configure the NAT session persistence scan.

- To enable the NAT session persistence scan:

```
[edit security nat source]
user@host# set session-persistence-scan
```

Results

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

```
[edit]
user@host# show security
  nat {
    source {
      session-persistence-scan;
    }
  }
```

Understanding NAT Configuration Check on Egress Interfaces after Reroute

The Network Address Translation (NAT) configuration often changes to accommodate more users and to enhance shortest route to transfer the traffic. If there is a change in egress interface because of rerouting of traffic, you can use the `set security flow enable-reroute-uniform-link-check nat` command to retain the existing NAT configuration and rule.

When the `enable-reroute-uniform-link-check nat` command is enabled:

- The session is retained with the existing NAT rule, if the new egress interface and the previous egress interface are in the same security zone, and there is no change in the matched NAT rule or if no rule is applied before and after rerouting.
- The session expires if the new egress interface and the previous egress interface are in the same security zone and the matched NAT rule is changed.

When the `enable-reroute-uniform-link-check nat` command is disabled:

- The traffic is forwarded to the new egress interface if the new egress interface and the previous egress interface are in the same security zone.

Configuration

To enable the NAT configuration for an existing session when there is a change in egress interface because of rerouting, use the following command:

```
[edit]
```

```
user@host# set security flow enable-reroute-uniform-link-check nat
```

The new configuration is applied when you commit the configuration changes.

The `enable-reroute-uniform-link-check nat` command is disabled by default.

Limitations

Retaining the NAT configuration using the `set security flow enable-reroute-uniform-link-check nat` command has the following limitations:

- The TCP synchronization does not allow the new session to transfer the traffic. You must disable the TCP synchronization to allow the transfer of traffic in new sessions.
- The packet information might be lost if reroute is initiated after a three-way handshake to initialize communication. You must disable the Junos OS Services Framework (JSF) like Application Layer Gateway (ALG) to allow the transfer of traffic in new sessions.

Release History Table

Release	Description
17.4R1	Starting in Junos OS Release 17.4R1, source NAT resources handled by the central point architecture have been offloaded to the SPUs when the SPC number is more than four, resulting in more efficient resource allocation.
15.1X49-D30	Starting in Junos OS Release 15.1X49-D30 and Junos OS Release 17.3R1, the central point architecture for NAT has been enhanced to handle higher system session capacity and session ramp-up rate for the SRX5000 line.
12.3X48-D40	In Release 12.3X48-D40, and in Release 15.1X49-D60 and later releases, you can increase the source NAT port capacity to 2.4G on SRX5400, SRX5600, and SRX5800 devices with next-generation Services Processing Cards (SPCs) using the <code>port-scaling-enlargement</code> statement at the <code>[edit security nat source]</code> hierarchy level supported

Destination NAT

IN THIS SECTION

- [Understanding Destination NAT | 125](#)
- [Understanding Destination NAT Address Pools | 126](#)
- [Understanding Destination NAT Rules | 127](#)
- [Destination NAT Configuration Overview | 127](#)
- [Example: Configuring Destination NAT for Single Address Translation | 128](#)
- [Example: Configuring Destination NAT for IP Address and Port Translation | 139](#)
- [Example: Configuring Destination NAT for Subnet Translation | 146](#)
- [Monitoring Destination NAT Information | 153](#)

Destination NAT changes the destination address of packets passing through the Router. It also offers the option to perform the port translation in the TCP/UDP headers. Destination NAT mainly used to redirect incoming packets with an external address or port destination to an internal IP address or port inside the network.

Understanding Destination NAT

Destination NAT is the translation of the destination IP address of a packet entering the Juniper Networks device. Destination NAT is used to redirect traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

NOTE: When destination NAT is performed, the destination IP address is translated according to configured destination NAT rules and then security policies are applied.

Destination NAT allows connections to be initiated only for incoming network connections—for example, from the Internet to a private network. Destination NAT is commonly used to perform the following actions:

- Translate a single IP address to another address (for example, to allow a device on the Internet to connect to a host on a private network).

- Translate a contiguous block of addresses to another block of addresses of the same size (for example, to allow access to a group of servers).
- Translate a destination IP address and port to another destination IP address and port (for example, to allow access to multiple services using the same IP address but different ports).

The following types of destination NAT are supported:

- Translation of the original destination IP address to an IP address from a user-defined pool. This type of translation does not include Port Address Translation (PAT). If the original destination IP address range is larger than the address range in the user-defined address pool, any untranslated packets are dropped.
- Translation of the original destination IP address (and optional port number) to one specific IP address (and port number) from a user-defined pool.

Understanding Destination NAT Address Pools

A NAT pool is a user-defined set of IP addresses that are used for translation. Unlike static NAT, where there is a one-to-one mapping that includes destination IP address translation in one direction and source IP address translation in the reverse direction, with destination NAT, you translate the original destination address to an IP address in the address pool.

For destination NAT address pools, specify the following:

- Name of the destination NAT address pool
- Destination address or address range

NOTE: Do not overlap NAT addresses for source NAT, destination NAT, and static NAT within one routing instance.

- Destination port that is used for port forwarding
- Routing instance to which the pool belongs—A destination NAT pool that does not specify a specific routing instance will default to the routing instance of the ingress zone.

NOTE: You can configure a NAT pool to exist in the default routing instance. Configuration option to specify that a NAT pool exists in the default routing-instance is available. As a

result, the NAT pool is reachable from zones in the default routing instance, and from zones in other routing instances.

Understanding Destination NAT Rules

Destination NAT rules specify two layers of match conditions:

- Traffic direction—Allows you to specify from interface, from zone, or from routing-instance.
- Packet information—Can be source IP addresses, destination IP address or subnet, destination port numbers or port ranges, protocols, or applications.

For ALG traffic, we recommend that you not use the destination-port option or the application option as matching conditions. If these options are used, translation may fail because the port value in the application payload might not match the port value in the IP address.

If multiple destination NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 and rule B specifies traffic from interface `ge-0/0/0`, rule B is used to perform destination NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match.

The actions you can specify for a destination NAT rule are:

- off—Do not perform destination NAT.
- pool—Use the specified user-defined address pool to perform destination NAT.

Destination NAT rules are applied to traffic in the first packet that is processed for the flow or in the fast path for the ALG. Destination NAT rules are processed after static NAT rules but before source NAT rules.

Destination NAT Configuration Overview

The main configuration tasks for destination NAT are as follows:

1. Configure a destination NAT address pool that aligns with your network and security requirements.
2. Configure destination NAT rules that align with your network and security requirements.
3. Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

Example: Configuring Destination NAT for Single Address Translation

IN THIS SECTION

- [Requirements | 128](#)
- [Overview | 128](#)
- [Configuration | 131](#)
- [Verification | 135](#)

This example describes how to configure a destination NAT mapping of a single public address to a private address.

NOTE: Mapping one destination IP address to another can also be accomplished with static NAT. Static NAT mapping allows connections to be established from either side of the gateway device, whereas destination NAT only allows connections to be established from one side. However, static NAT only allows translations from one address to another or between blocks of addresses of the same size.

Requirements

This example uses the following hardware and software components:

- SRX Series Firewalls
- Server

Before you begin:

- Configure network interfaces on the device. See the [Interfaces User Guide for Security Devices](#).
- Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

IN THIS SECTION

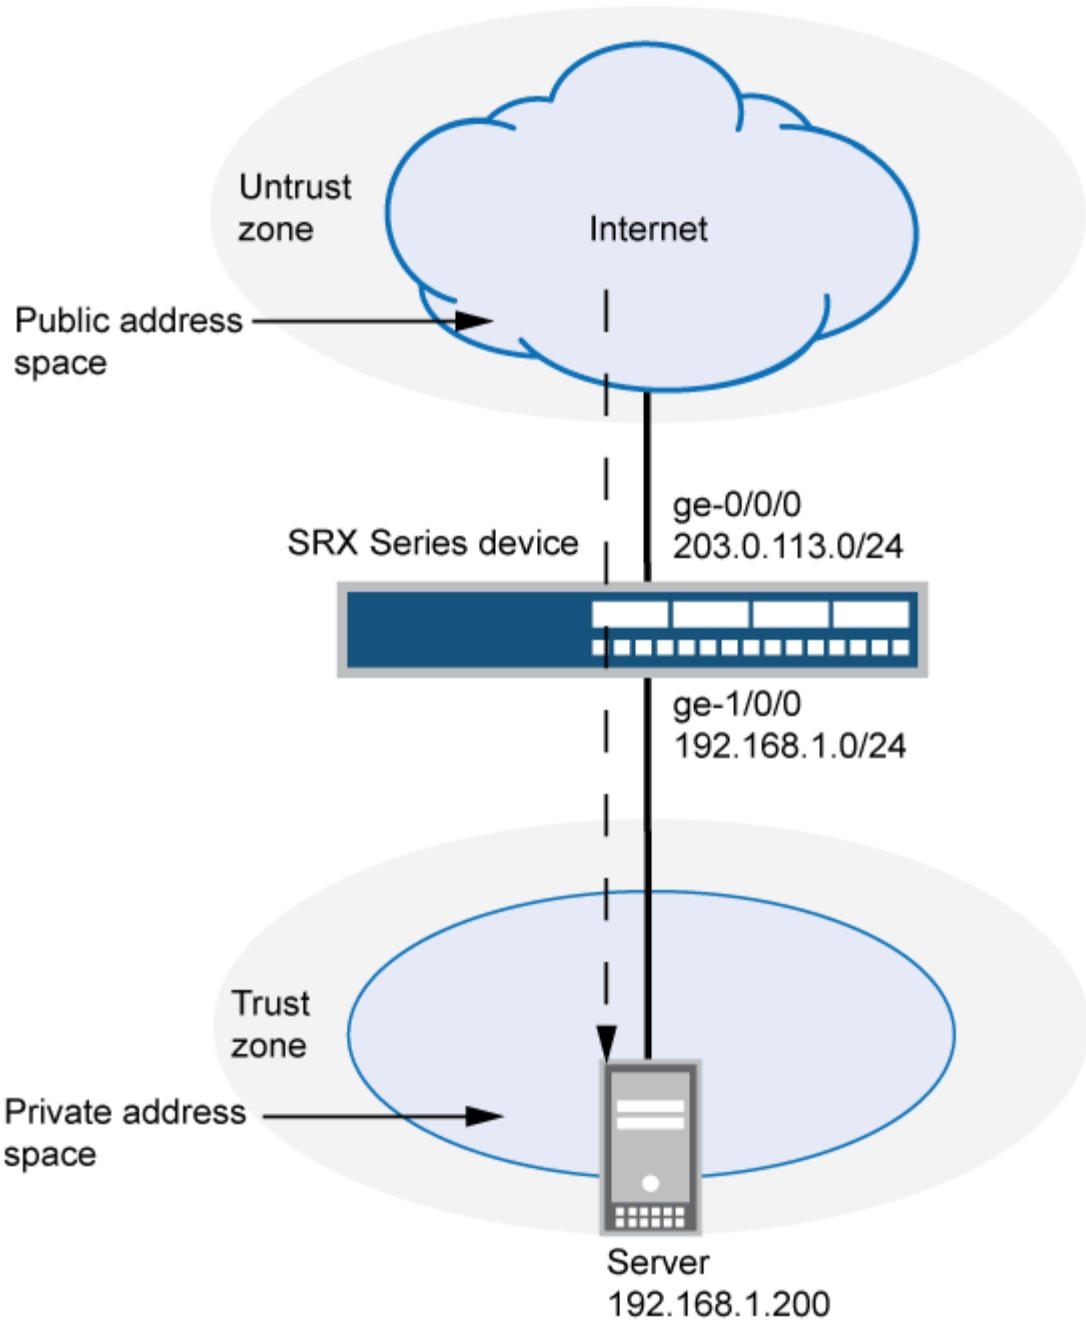
- [Topology | 130](#)

Destination NAT is commonly used to distribute a service located in a private network with a publicly accessible IP address. This allows users to use the private service with the public IP address. Destination NAT address pool and destination NAT rules configurations are used to align your network and improve security requirements.

In this example, first you configure the trust security zone for the private address space and then you configure the untrust security zone for the public address space. In [Figure 9 on page 130](#), devices in the untrust zone access a server in the trust zone by way of public address 203.0.113.200/32. For packets that enter the Juniper Networks security device from the untrust zone with the destination IP address 203.0.113.200/32, the destination IP address is translated to the private address 192.168.1.200/32.

Topology

Figure 9: Destination NAT Single Address Translation



Original Destination IP	Translated Destination IP
203.0.113.200/32	192.168.1.200/32

Table 7 on page 131 shows the parameters configured in this example.

Table 7: Interfaces, Zones, Server, and IP Address Information

Parameter	Description
Trust Zone	Security zone for the private address space.
Untrust Zone	Security zone for the public address space.
192.168.1.200/32	Translated destination NAT IP address.
192.168.1.0/24	Private subnet in private zone.
203.0.113.200/32	Public address of the server.
Server	Server address of the private address space.
ge-0/0/0 and ge-1/0/0	NAT interfaces for traffic direction.

This example describes the following configurations:

- Destination NAT pool `dst-nat-pool-1` that contains the IP address `192.168.1.200/32`.
- Destination NAT rule set `rs1` with rule `r1` to match packets received from the `ge-0/0/0.0` interface with the destination IP address `203.0.113.200/32`. For matching packets, the destination address is translated to the address in the `dst-nat-pool-1` pool.
- Proxy ARP for the address `203.0.113.200/32` on interface `ge-0/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the untrust zone to the translated destination IP address in the trust zone.

Configuration

IN THIS SECTION

Procedure | 132

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security nat destination pool dst-nat-pool-1 address 192.168.1.200/32
set security nat destination rule-set rs1 from interface ge-0/0/0.0
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.200/32
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security address-book global address server-1 192.168.1.200/32
set security policies from-zone untrust to-zone trust policy server-access match source-address any
set security policies from-zone untrust to-zone trust policy server-access match destination-address server-1
set security policies from-zone untrust to-zone trust policy server-access match application any
set security policies from-zone untrust to-zone trust policy server-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure a destination NAT mapping from a public address to a private address:

1. Create the destination NAT pool.

```
[edit security nat destination]
user@host# set pool dst-nat-pool-1 address 192.168.1.200/32
```

2. Create a destination NAT rule set.

```
[edit security nat destination]
user@host# set rule-set rs1 from interface ge-0/0/0.0
```

3. Configure a rule that matches packets and translates the destination address to the address in the pool.

```
[edit security nat destination]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.200/32
user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
```

4. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
```

5. Configure an address in the global address book.

```
[edit security address-book global]
user@host# set address server-1 192.168.1.200/32
```

6. Configure a security policy that allows traffic from the untrust zone to the server in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy server-access match source-address any
user@host# set policy server-access match destination-address server-1
user@host# set policy server-access match application any
user@host# set policy server-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show interfaces`, `show security zones`, and `show bridge-domains` command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security nat
destination {
  pool dst-nat-pool-1 {
    address 192.168.1.200/32;
  }
}
```



```

rule-set rs1 {
    from interface ge-0/0/0.0;
    rule r1 {
        match {
            destination-address 203.0.113.200/32;
        }
        then {
            destination-nat pool dst-nat-pool-1;
        }
    }
}

}

proxy-arp {
    interface ge-0/0/0.0 {
        address {
            203.0.113.200/32;
        }
    }
}

[edit]
user@host# show security address-book
global {
    address server-1 192.168.1.200/32;
}

user@host# show security policies
from-zone untrust to-zone trust {
    policy server-access {
        match {
            source-address any;
            destination-address server-1;
            application any;
        }
        then {
            permit;
        }
    }
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Destination NAT Pool Usage | 135](#)
- [Verifying Destination NAT Rule Usage | 136](#)
- [Verifying Destination NAT for a Single Address Translation | 136](#)
- [Verifying NAT Application to Traffic | 137](#)

Confirm that the configuration is working properly.

Verifying Destination NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the destination NAT pool.

Action

From operational mode, enter the `show security nat destination pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

```
user@host>show security nat destination pool all
Total destination-nat pools: 1

Pool name      : dst-nat-pool-1
Pool id       : 1
Total address  : 1
Translation hits: 71
Address range          Port
192.168.1.200 - 192.168.1.200      0
```

Meaning

The `show security nat destination pool all` command displays the pool of translated addresses. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose

Verify that there is traffic matching the destination NAT rule.

Action

From operational mode, enter the `show security nat destination rule all` command.

```
user@host>show security nat destination rule all
Total destination-nat rules: 1
Total referenced IPv4/IPv6 ip-prefixes: 1/0

Destination NAT rule: r1                Rule-set: rs1
Rule-Id                        : 1
Rule position                  : 1
From interface                  : ge-0/0/0.0
  Destination addresses        : 203.0.113.200 - 203.0.113.200
Action                         : dst-nat-pool-1
Translation hits                : 75
  Successful sessions          : 75
  Failed sessions              : 0
Number of sessions              : 4
```

Meaning

The `show security nat destination rule all` command displays the destination NAT rule. View the Translation hits field to check for traffic that matches the destination rule.

Verifying Destination NAT for a Single Address Translation

Purpose

Verify the configuration of destination NAT for a single address translation.

Action

From operational mode, enter the `show security nat destination summary` command.

```
user@host>show security nat destination summary
Total pools: 1
Pool name          Address Range          Routing
                   192.168.1.200 - 192.168.1.200 Instance
                   Port    Total
                   Address
dst-nat-pool-1     0      1

Total rules: 1
Rule name          Rule set    From          Action
r1                 rs1         ge-0/0/0.0    dst-nat-pool-1
```

Meaning

The `show security nat destination summary` command displays information about destination NAT configuration. You can verify the following information:

- Rule sets
- Rules
- Address range
- NAT pool
- Port details

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

```
user@host>show security flow session
Session ID: 26415, Policy name: server-access/11, Timeout: 2, Valid
In: 203.0.113.219/30 --> 203.0.113.200/54850;icmp, If: ge-0/0/0.0, Pkts: 1, Bytes: 84
```

```

Out: 192.168.1.200/54850 --> 203.0.113.219/30;icmp, If: ge-0/0/1.0, Pkts: 1, Bytes: 84

Session ID: 26420, Policy name: server-access/11, Timeout: 2, Valid
In: 203.0.113.219/31 --> 203.0.113.200/54850;icmp, If: ge-0/0/0.0, Pkts: 1, Bytes: 84
Out: 192.168.1.200/54850 --> 203.0.113.219/31;icmp, If: ge-0/0/1.0, Pkts: 1, Bytes: 84

Session ID: 26425, Policy name: server-access/11, Timeout: 4, Valid
In: 203.0.113.219/32 --> 203.0.113.200/54850;icmp, If: ge-0/0/0.0, Pkts: 1, Bytes: 84
Out: 192.168.1.200/54850 --> 203.0.113.219/32;icmp, If: ge-0/0/1.0, Pkts: 1, Bytes: 84

Session ID: 26431, Policy name: server-access/11, Timeout: 4, Valid
In: 203.0.113.219/33 --> 203.0.113.200/54850
;icmp, If: ge-0/0/0.0, Pkts: 1, Bytes: 84
Out: 192.168.1.200/54850 --> 203.0.113.219/33;icmp, If: ge-0/0/1.0, Pkts: 1, Bytes: 84
Total sessions: 9

```

Meaning

The `show security flow session` command displays active sessions on the device and each session's associated security policy. The output shows traffic entering the device destined to a public host at 203.0.113.200 that gets translated to private destination IP address 192.168.1.200.

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name or number of packets in and out.
- **server-access**—Policy name that permitted the traffic from the untrust zone to the translated destination IP address in the trust zone.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers, the session is ICMP, and the source interface for this session is ge-0/0/0.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers, the session is ICMP, and the destination interface for this session is ge-0/0/1.0).

Example: Configuring Destination NAT for IP Address and Port Translation

IN THIS SECTION

- [Requirements | 139](#)
- [Overview | 139](#)
- [Configuration | 141](#)
- [Verification | 145](#)

This example describes how to configure destination NAT mappings of a public address to private addresses, depending on the port number.

Requirements

Before you begin:


1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 10 on page 140](#), devices in the untrust zone access servers in the trust zone by way of public address 203.0.113.200 on port 80 or 8000. Packets entering the Juniper Networks security device from the untrust zone are mapped to the private addresses of the servers as follows:

- The destination IP address 203.0.113.200 and port 80 is translated to the private address 192.168.1.200 and port 80.
- The destination IP address 203.0.113.200 and port 8000 is translated to the private address 192.168.1.220 and port 8000.

Figure 10: Destination NAT Address and Port Translation



Original Destination IP	Translated Destination IP
203.0.113.200 port 80	192.168.1.200 port 80
203.0.113.200 port 8000	192.168.1.220 port 8000

This example describes the following configurations:

- Destination NAT pool `dst-nat-pool-1` that contains the IP address `192.168.1.200` port `80`.
- Destination NAT pool `dst-nat-pool-2` that contains the IP address `192.168.1.220` and port `8000`.
- Destination NAT rule set `rs1` with rule `r1` to match packets received from the untrust zone with the destination IP address `203.0.113.200` and destination port `80`. For matching packets, the destination address is translated to the address in the `dst-nat-pool-1` pool.
- Destination NAT rule set `rs1` with rule `r2` to match packets received from the untrust zone with the destination IP address `203.0.113.200` and destination port `8000`. For matching packets, the destination IP address and port are translated to the address and port in the `dst-nat-pool-2` pool.
- Proxy ARP for the address `203.0.113.200/32`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic from the untrust zone to the translated destination IP addresses in the trust zone.

Configuration

IN THIS SECTION

- [Procedure | 141](#)

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat destination pool dst-nat-pool-1 address 192.168.1.200/32
set security nat destination pool dst-nat-pool-1 address port 80
set security nat destination pool dst-nat-pool-2 address 192.168.1.220/32
set security nat destination pool dst-nat-pool-2 address port 8000
set security nat destination rule-set rs1 from zone untrust
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.200/32
set security nat destination rule-set rs1 rule r1 match destination-port 80
```



```

set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat destination rule-set rs1 rule r2 match destination-address 203.0.113.200/32
set security nat destination rule-set rs1 rule r2 match destination-port 8000
set security nat destination rule-set rs1 rule r2 then destination-nat pool dst-nat-pool-2
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security address-book global address server-2 192.168.1.220/32
set security address-book global address server-1 192.168.1.200/32
set security policies from-zone untrust to-zone trust policy server-access match source-address
any
set security policies from-zone untrust to-zone trust policy server-access match destination-
address server-1
set security policies from-zone untrust to-zone trust policy server-access match destination-
address server-2
set security policies from-zone untrust to-zone trust policy server-access match application any
set security policies from-zone untrust to-zone trust policy server-access then permit

```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a destination NAT mapping from a public address to a private address:

1. Create destination NAT pools.

```

[edit security nat destination]
user@host# set pool dst-nat-pool-1 address 192.168.1.200 port 80
user@host# set pool dst-nat-pool-2 address 192.168.1.220 port 8000

```

2. Create a destination NAT rule set.

```

[edit security nat destination]
user@host# set rule-set rs1 from zone untrust

```

3. Configure a rule that matches packets and translates the destination address to the address in the pool.

```

[edit security nat destination]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.200

```

```
user@host# set rule-set rs1 rule r1 match destination-port 80
user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
```

4. Configure a rule that matches packets and translates the destination address to the address in the pool.

```
[edit security nat destination]
user@host# set rule-set rs1 rule r2 match destination-address 203.0.113.200
user@host# set rule-set rs1 rule r2 match destination-port 8000
user@host# set rule-set rs1 rule r2 then destination-nat pool dst-nat-pool-2
```

5. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
```

6. Configure addresses in the global address book.

```
[edit security address-book global]
user@host# set address server-2 192.168.1.220/32
user@host# set address server-1 192.168.1.200/32
```

7. Configure a security policy that allows traffic from the untrust zone to the servers in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy server-access match source-address any destination-address [server-1
server-2] application any
user@host# set policy server-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
```

```

destination {
    pool dst-nat-pool-1 {
        address 192.168.1.200/32 port 80;
    }
    pool dst-nat-pool-2 {
        address 192.168.1.220/32 port 8000;
    }
    rule-set rs1 {
        from zone untrust;
        rule r1 {
            match {
                destination-address 203.0.113.200/32;
                destination-port 80;
            }
            then {
                destination-nat pool dst-nat-pool-1;
            }
        }
        rule r2 {
            match {
                destination-address 203.0.113.200/32;
                destination-port 8000;
            }
            then {
                destination-nat pool dst-nat-pool-2;
            }
        }
    }
}

proxy-arp {
    interface ge-0/0/0.0 {
        address {
            203.0.113.200/32;
        }
    }
}

user@host# show security policies
from-zone untrust to-zone trust {
    policy server-access {
        match {
            source-address any;
            destination-address [ server-1 server-2 ];
            application any;

```

```

    }
    then {
        permit;
    }
}
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Destination NAT Pool Usage | 145](#)
- [Verifying Destination NAT Rule Usage | 145](#)
- [Verifying NAT Application to Traffic | 146](#)

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

Verifying Destination NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the destination NAT pool.

Action

From operational mode, enter the `show security nat destination pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose

Verify that there is traffic matching the destination NAT rule.

Action

From operational mode, enter the `show security nat destination rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Example: Configuring Destination NAT for Subnet Translation

IN THIS SECTION

- [Requirements | 146](#)
- [Overview | 147](#)
- [Configuration | 149](#)
- [Verification | 152](#)

This example describes how to configure a destination NAT mapping of a public subnet address to a private subnet address.

NOTE: Mapping addresses from one subnet to another can also be accomplished with static NAT. Static NAT mapping allows connections to be established from either side of the gateway device, whereas destination NAT allows connections to be established from only one side. However, static NAT only allows translations between blocks of addresses of the same size.

Requirements


Before you begin:

1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 11 on page 148](#), devices in the untrust zone access devices in the trust zone by way of public subnet address 203.0.113.0/24. For packets that enter the Juniper Networks security device from the untrust zone with a destination IP address in the 203.0.113.0/24 subnet, the destination IP address is translated to a private address on the 192.168.1.0/24 subnet.

Figure 11: Destination NAT Subnet Translation



This example describes the following configurations:

- Destination NAT pool `dst-nat-pool-1` that contains the IP address `192.168.1.0/24`.
- Destination NAT rule set `rs1` with rule `r1` to match packets received from the `ge-0/0/0.0` interface with the destination IP address on the `203.0.113.0/24` subnet. For matching packets, the destination address is translated to the address in the `dst-nat-pool-1` pool.
- Proxy ARP for the addresses `203.0.113.1/32` through `203.0.113.62/32` on the interface `ge-0/0/0.0`; these are the IP addresses of the hosts that should be translated from the `203.0.113.0/24` subnet. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses. The address `203.0.113.0/24` is assigned to the interface itself, so this address is not included in the proxy ARP configuration. The addresses that are not in the `203.0.113.1/32` through `203.0.113.62/32` range are not expected to be present on the network and would not be translated.
- Security policies to permit traffic from the untrust zone to the translated destination IP addresses in the trust zone.

Configuration

IN THIS SECTION

- [Procedure](#) | 149

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat destination pool dst-nat-pool-1 address 192.168.1.0/24
set security nat destination rule-set rs1 from interface ge-0/0/0.0
set security nat destination rule-set rs1 rule r1 match destination-address 203.0.113.0/24
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.62/32
set security address-book global address internal-net 192.168.1.0/24
set security policies from-zone untrust to-zone trust policy internal-access match source-
```



```

address any
set security policies from-zone untrust to-zone trust policy internal-access match destination-
address internal-net
set security policies from-zone untrust to-zone trust policy internal-access match application
any
set security policies from-zone untrust to-zone trust policy internal-access then permit

```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a destination NAT mapping from a public subnet address to a private subnet address:

1. Create the destination NAT pool.

```

[edit security nat destination]
user@host# set pool dst-nat-pool-1 address 192.168.1.0/24

```

2. Create a destination NAT rule set.

```

[edit security nat destination]
user@host# set rule-set rs1 from interface ge-0/0/0.0

```

3. Configure a rule that matches packets and translates the destination address to an address in the pool.

```

[edit security nat destination]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.0/24
user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool-1

```

4. Configure proxy ARP.

```

[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.62/32

```

5. Configure an address in the global address book.

```
[edit security address-book global]
user@host# set address internal-net 192.168.1.0/24
```

6. Configure a security policy that allows traffic from the untrust zone to the devices in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy internal-access match source-address any destination-address internal-
net application any
user@host# set policy internal-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
destination {
  pool dst-nat-pool-1 {
    address 192.168.1.0/24;
  }
  rule-set rs1 {
    from interface ge-0/0/0.0;
    rule r1 {
      match {
        destination-address 203.0.113.0/24;
      }
      then {
        destination-nat pool dst-nat-pool-1;
      }
    }
  }
}
proxy-arp {
  interface ge-0/0/0.0 {
    address {
```

```

        203.0.113.1/32 to 203.0.113.62/32;
    }
}
}
user@host# show security policies
from-zone untrust to-zone trust {
    policy internal-access {
        match {
            source-address any;
            destination-address internal-net;
            application any;
        }
        then {
            permit;
        }
    }
}
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Destination NAT Pool Usage | 152](#)
- [Verifying Destination NAT Rule Usage | 153](#)
- [Verifying NAT Application to Traffic | 153](#)

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

Verifying Destination NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the destination NAT pool.

Action

From operational mode, enter the `show security nat destination pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose

Verify that there is traffic matching the destination NAT rule.

Action

From operational mode, enter the `show security nat destination rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Monitoring Destination NAT Information

IN THIS SECTION

- [Purpose | 154](#)
- [Action | 154](#)

Purpose

View the destination Network Address Translation (NAT) summary table and the details of the specified NAT destination address pool information.

Action

Select **Monitor>NAT> Destination NAT** in the J-Web user interface, or enter the following CLI commands:

- show security nat destination summary
- show security nat destination pool *pool-name*

[Table 8 on page 154](#) summarizes key output fields in the destination NAT display.

Table 8: Summary of Key Destination NAT Output Fields

Field	Values	Action
Rules		
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.
Total rules	Number of rules configured.	–
ID	Rule ID number.	–
Name	Name of the rule .	–
Ruleset Name	Name of the rule set.	–
From	Name of the routing instance/zone/interface from which the packet flows.	–

Table 8: Summary of Key Destination NAT Output Fields (Continued)

Field	Values	Action
Source address range	Source IP address range in the source pool.	–
Destination address range	Destination IP address range in the source pool.	–
Destination port	Destination port in the destination pool.	–
IP protocol	IP protocol.	–
Action	Action taken for a packet that matches a rule.	–
Alarm threshold	Utilization alarm threshold.	–
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. <ul style="list-style-type: none"> • Succ–Number of successful session installations after the NAT rule is matched. • Failed–Number of unsuccessful session installations after the NAT rule is matched. • Current–Number of sessions that reference the specified rule. 	–
Translation hits	Number of times a translation in the translation table is used for a destination NAT rule.	–
Pools		

Table 8: Summary of Key Destination NAT Output Fields *(Continued)*

Field	Values	Action
Pool Name	The names of the pools.	Select all pools or a specific pool to display from the list.
Total Pools	Total pools added.	–
ID	ID of the pool.	–
Name	Name of the destination pool.	–
Address range	IP address range in the destination pool.	–
Port	Destination port number in the pool.	–
Routing instance	Name of the routing instance.	–
Total addresses	Total IP address, IP address set, or address book entry.	–
Translation hits	Number of times a translation in the translation table is used for destination NAT.	–
Top 10 Translation Hits		
Graph	Displays the graph of top 10 translation hits.	–

Static NAT

IN THIS SECTION

- [Understanding Static NAT | 157](#)
- [Understanding Static NAT Rules | 158](#)
- [Static NAT Configuration Overview | 159](#)
- [Example: Configuring Static NAT for Single Address Translation | 159](#)
- [Example: Configuring Static NAT for Subnet Translation | 165](#)
- [Example: Configuring Static NAT for Port Mapping | 172](#)
- [Monitoring Static NAT Information | 181](#)

Static NAT maps network traffic from a static external IP address to an internal IP address or network. It creates a static translation of real addresses to mapped addresses. Static NAT provides internet connectivity to networking devices through a private LAN with an unregistered private IP address.

Understanding Static NAT

Static NAT defines a one-to-one mapping from one IP subnet to another IP subnet. The mapping includes destination IP address translation in one direction and source IP address translation in the reverse direction. From the NAT device, the original destination address is the virtual host IP address while the mapped-to address is the real host IP address.

Static NAT allows connections to be originated from either side of the network, but translation is limited to one-to-one or between blocks of addresses of the same size. For each private address, a public address must be allocated. No address pools are necessary.

Static NAT also supports the following types of translation:

- To map multiple IP addresses and specified ranges of ports to a same IP address and different range of ports
- To map a specific IP address and port to a different IP address and port

The port address translation (PAT) is also supported by giving static mapping between destination-port (range) and mapped-port (range).

NOTE: The original destination address, along with other addresses in source and destination NAT pools, must not overlap within the same routing instance.

In NAT rule lookup, static NAT rules take precedence over destination NAT rules and reverse mapping of static NAT rules take precedence over source NAT rules.

Understanding Static NAT Rules

Static Network Address Translation (NAT) rules specify two layers of match conditions:

- Traffic direction—Allows you to specify **from interface**, **from zone**, or **from routing-instance**.
- Packet information—Can be source addresses and ports, and destination addresses and ports.

For all ALG traffic, except FTP, we recommend that you not use the static NAT rule options `source-address` or `source-port`. Data session creation can fail if these options are used because the IP address and the source port value, which is a random value, might not match the static NAT rule. For FTP ALG traffic, the `source-address` option can be used because an IP address can be provided to match the source address of a static NAT rule.

When both source and destination addresses are configured as match conditions for a rule, traffic is matched to both the source address and destination address. Because static NAT is bidirectional, traffic in the opposite direction reverse matches the rule, and the destination address of the traffic is matched to the configured source address.

If multiple static NAT rules overlap in the match conditions, the most specific rule is chosen. For example, if rules A and B specify the same source and destination IP addresses, but rule A specifies traffic from zone 1 and rule B specifies traffic from interface `ge-0/0/0`, rule B is used to perform static NAT. An interface match is considered to be more specific than a zone match, which is more specific than a routing instance match.

Because static NAT rules do not support overlapping addresses and ports, they should not be used to map one external IP address to multiple internal IP addresses for ALG traffic. For example, if different sites want to access two different FTP servers, the internal FTP servers should be mapped to two different external IP addresses.

For the static NAT rule action, specify the translated address and (optionally) the routing instance.

In NAT lookup, static NAT rules take precedence over destination NAT rules and reverse mapping of static NAT rules takes precedence over source NAT rules.

Static NAT Configuration Overview

The main configuration tasks for static NAT are as follows:

1. Configure static NAT rules that align with your network and security requirements.
2. Configure NAT proxy ARP entries for IP addresses in the same subnet of the ingress interface.

Example: Configuring Static NAT for Single Address Translation

IN THIS SECTION

- [Requirements | 159](#)
- [Overview | 159](#)
- [Configuration | 161](#)
- [Verification | 164](#)

This example describes how to configure a static NAT mapping of a single private address to a public address.

Requirements

Before you begin:


1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space.

In [Figure 12 on page 160](#), devices in the untrust zone access a server in the trust zone by way of public address 203.0.113.200/32. For packets that enter the Juniper Networks security device from the untrust zone with the destination IP address 203.0.113.200/32, the destination IP address is translated to the private address 192.168.1.200/32. For a new session originating from the server, the source IP address in the outgoing packet is translated to the public address 203.0.113.200/32.

Figure 12: Static NAT Single Address Translation



Original Destination IP	Translated Destination IP
203.0.113.200/32	192.168.1.200/32

This example describes the following configurations:

- Static NAT rule set rs1 with rule r1 to match packets from the untrust zone with the destination address 203.0.113.200/32. For matching packets, the destination IP address is translated to the private address 192.168.1.200/32.
- Proxy ARP for the address 203.0.113.200 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for that address.
- Security policies to permit traffic to and from the 192.168.1.200 server.

Configuration

IN THIS SECTION

- [Procedure | 161](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set rs1 from zone untrust
set security nat static rule-set rs1 rule r1 match destination-address 203.0.113.200/32
set security nat static rule-set rs1 rule r1 then static-nat prefix 192.168.1.200/32
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.200/32
set security address-book global address server-1 192.168.1.200/32
set security policies from-zone trust to-zone untrust policy permit-all match source-address server-1
set security policies from-zone trust to-zone untrust policy permit-all match destination-address any
set security policies from-zone trust to-zone untrust policy permit-all match application any
set security policies from-zone trust to-zone untrust policy permit-all then permit
set security policies from-zone untrust to-zone trust policy server-access match source-address any
set security policies from-zone untrust to-zone trust policy server-access match destination-address server-1
```

```
set security policies from-zone untrust to-zone trust policy server-access match application any
set security policies from-zone untrust to-zone trust policy server-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a static NAT mapping from a private address to a public address:

1. Create a static NAT rule set.

```
[edit security nat static]
user@host# set rule-set rs1 from zone untrust
```

2. Configure a rule that matches packets and translates the destination address in the packets to a private address.

```
[edit security nat static]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.200/32
user@host# set rule-set rs1 rule r1 then static-nat prefix 192.168.1.200/32
```

3. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.200
```

4. Configure an address in the global address book.

```
[edit security address-book global]
user@host# set address server-1 192.168.1.200/32
```

5. Configure a security policy that allows traffic from the untrust zone to the server in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy server-access match source-address any destination-address server-1
```

```

application any
user@host# set policy server-access then permit

```

6. Configure a security policy that allows all traffic from the server in the trust zone to the untrust zone.

```

[edit security policies from-zone trust to-zone untrust]
user@host# set policy permit-all match source-address server-1 destination-address any
application any
user@host# set policy permit-all then permit

```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@host# show security nat
static {
    rule-set rs1 {
        from zone untrust;
        rule r1 {
            match {
                destination-address 203.0.113.200/32;
            }
            then {
                static-nat prefix 192.168.1.200/32;
            }
        }
    }
}
proxy-arp {
    interface ge-0/0/0.0 {
        address {
            203.0.113.200/32;
        }
    }
}
user@host# show security policies
from-zone trust to-zone untrust {

```

```
policy permit-all {  
    match {  
        source-address server-1;  
        destination-address any;  
        application any;  
    }  
    then {  
        permit;  
    }  
}  
}  
from-zone untrust to-zone trust {  
    policy server-access {  
        match {  
            source-address any;  
            destination-address server-1;  
            application any;  
        }  
        then {  
            permit;  
        }  
    }  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Static NAT Configuration | 165](#)
- [Verifying NAT Application to Traffic | 165](#)

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

Verifying Static NAT Configuration

Purpose

Verify that there is traffic matching the static NAT rule set.

Action

From operational mode, enter the `show security nat static rule` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Example: Configuring Static NAT for Subnet Translation

IN THIS SECTION

- [Requirements | 166](#)
- [Overview | 166](#)
- [Configuration | 168](#)
- [Verification | 171](#)

This example describes how to configure a static NAT mapping of a private subnet address to a public subnet address.

NOTE: Address blocks for static NAT mapping must be of the same size.

Requirements


Before you begin:

1. Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. In [Figure 13 on page 167](#), devices in the untrust zone access devices in the trust zone by way of public subnet address 203.0.113.0/24. For packets that enter the Juniper Networks security device from the untrust zone with a destination IP address in the 203.0.113.0/24 subnet, the destination IP address is translated to a private address on the 192.168.1.0/24 subnet. For new sessions originating from the 192.168.1.0/24 subnet, the source IP address in outgoing packets is translated to an address on the public 203.0.113.0/24 subnet.

Figure 13: Static NAT Subnet Translation



Original Destination IP	Translated Destination IP
203.0.113.200/32	192.168.1.0/24

This example describes the following configurations:

- Static NAT rule set rs1 with rule r1 to match packets received on interface ge-0/0/0.0 with a destination IP address in the 203.0.113.0/24 subnet. For matching packets, the destination address is translated to an address on the 192.168.1.0/24 subnet.
- Proxy ARP for the address ranges 203.0.113.1/32 through 203.0.113.249/32 on interface ge-0/0/0.0. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses. The address 203.0.113.250/32 is assigned to the interface itself, so this address is not included in the proxy ARP configuration.
- Security policies to permit traffic to and from the 192.168.1.0/24 subnet.

Configuration

IN THIS SECTION

- [Procedure | 168](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set rs1 from interface ge-0/0/0.0
set security nat static rule-set rs1 rule r1 match destination-address 203.0.113.0/24
set security nat static rule-set rs1 rule r1 then static-nat prefix 192.168.1.0/24
set security nat proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.249/32
set security address-book global address server-group 192.168.1.0/24
set security policies from-zone trust to-zone untrust policy permit-all match source-address server-group
set security policies from-zone trust to-zone untrust policy permit-all match destination-address any
set security policies from-zone trust to-zone untrust policy permit-all match application any
set security policies from-zone trust to-zone untrust policy permit-all then permit
set security policies from-zone untrust to-zone trust policy server-access match source-address any
```

```
set security policies from-zone untrust to-zone trust policy server-access match destination-
address server-group
set security policies from-zone untrust to-zone trust policy server-access match application any
set security policies from-zone untrust to-zone trust policy server-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure a static NAT mapping from a private subnet address to a public subnet address:

1. Create a static NAT rule set.

```
[edit security nat static]
user@host# set rule-set rs1 from interface ge-0/0/0.0
```

2. Configure a rule that matches packets and translates the destination address in the packets to an address in a private subnet.

```
[edit security nat static]
user@host# set rule-set rs1 rule r1 match destination-address 203.0.113.0/24
user@host# set rule-set rs1 rule r1 then static-nat prefix 192.168.1.0/24
```

3. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface ge-0/0/0.0 address 203.0.113.1/32 to 203.0.113.249/32
```

4. Configure an address in the global address book.

```
[edit security address-book global]
user@host# set address server-group 192.168.1.0/24
```

5. Configure a security policy that allows traffic from the untrust zone to the subnet in the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy server-access match source-address any destination-address server-group
```

```
application any
user@host# set policy server-access then permit
```

6. Configure a security policy that allows all traffic from the subnet in the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy permit-all match source-address server-group destination-address any
application any
user@host# set policy permit-all then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
static {
    rule-set rs1 {
        from interface ge-0/0/0.0;
        rule r1 {
            match {
                destination-address 203.0.113.0/24;
            }
            then {
                static-nat prefix 192.168.1.0/24;
            }
        }
    }
}
proxy-arp {
    interface ge-0/0/0.0 {
        address {
            203.0.113.1/32 to 203.0.113.249/32;
        }
    }
}
user@host# show security policies
```

```
from-zone trust to-zone untrust {  
  policy permit-all {  
    match {  
      source-address server-group;  
      destination-address any;  
      application any;  
    }  
    then {  
      permit;  
    }  
  }  
}  
  
from-zone untrust to-zone trust {  
  policy server-access {  
    match {  
      source-address any;  
      destination-address server-group;  
      application any;  
    }  
    then {  
      permit;  
    }  
  }  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Static NAT Configuration | 172](#)
- [Verifying NAT Application to Traffic | 172](#)

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

Verifying Static NAT Configuration

Purpose

Verify that there is traffic matching the static NAT rule set.

Action

From operational mode, enter the `show security nat static rule` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

Example: Configuring Static NAT for Port Mapping

IN THIS SECTION

- [Requirements | 173](#)
- [Overview | 173](#)
- [Configuration | 175](#)
- [Verification | 179](#)
- [Troubleshooting | 180](#)

This example describes how to configure static NAT mappings of a public address to private addresses on a specified range of ports.

This topic includes the following sections:

Requirements

Before you begin:


- Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).
- Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

This example uses the trust security zone for the private address space and the untrust security zone for the public address space.

In [Figure 14 on page 174](#), devices in the untrust zone access a server in the trust zone by way of public addresses 203.0.113.1/32, 203.0.113.1/32, and 203.0.113.3/32. For packets that enter the Juniper Networks security device from the untrust zone with the destination IP addresses 203.0.113.1/32, 203.0.113.1/32, and 203.0.113.3/32, the destination IP address is translated to the private addresses 10.1.1.1/32, 10.1.1.2/32, and 10.1.1.2/32.

Figure 14: Static NAT for Port Mapping



Original Source IP	Translated Source IP
203.0.113.1/32 (port 100 to 200)	10.1.1.1/32 (port 300 to 400)
203.0.113.1/32 (port 300 to 400)	10.1.1.2/32 (port 300 to 400)
203.0.113.3/32 (port 300)	10.1.1.2/32 (port 200)

g034403

NOTE:

- To configure the destination port, you must use an IP address for the destination address field instead of an IP address prefix.
- You must configure the destination port to configure the mapped port and vice versa.
- Use the same number range for the ports while configuring the destination port and the mapped port.
- If you do not configure the destination port and the mapped port, the IP mapping will be the one-to-one mapping.
- Any address overlapping or any address and port overlapping is not allowed.

This example describes the following configurations:

- Static NAT rule set rs1 with rule r1 to match packets from the untrust zone with the destination address 203.0.113.1/32 and destination port 100 to 200. For matching packets, the destination IP address is translated to the private address 10.1.1.1/32 and mapped to port 300 to 400.
- Static NAT rule set rs1 with rule r2 to match packets from the untrust zone with the destination address 203.0.113.1/32 and destination port 300 to 400. For matching packets, the destination IP address is translated to the private address 10.1.1.2/32 and mapped to port 300 to 400.
- Static NAT rule set rs1 with rule r3 to match packets from the untrust zone with the destination address 203.0.113.3/32 and destination port 300. For matching packets, the destination IP address is translated to the private address 10.1.1.2/32 and mapped to port 200.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 176](#)
- [Procedure | 176](#)
- [Results | 177](#)

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set rs from zone untrust
set security nat static rule-set rs rule r1 match destination-address 203.0.113.1/32
set security nat static rule-set rs rule r1 match destination-port 100 to 200
set security nat static rule-set rs rule r1 then static-nat prefix 10.1.1.1/32
set security nat static rule-set rs rule r1 then static-nat prefix mapped-port 300 to 400
set security nat static rule-set rs rule r2 match destination-address 203.0.113.1/32
set security nat static rule-set rs rule r2 match destination-port 300 to 400
set security nat static rule-set rs rule r2 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs rule r2 then static-nat prefix mapped-port 300 to 400
set security nat static rule-set rs rule r3 match destination-address 203.0.113.3/32
set security nat static rule-set rs rule r3 match destination-port 300
set security nat static rule-set rs rule r3 then static-nat prefix 10.1.1.2/32
set security nat static rule-set rs rule r3 then static-nat prefix mapped-port 200
```

Procedure

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see [Using the CLI Editor in Configuration Mode](#).

To configure a static NAT mapping from a private subnet address to a public subnet address:

1. Create a static NAT rule set.

```
[edit security nat static]
user@host# set rule-set rs from zone untrust
```

2. Configure a rule that matches packets and translates the destination address in the packets to a private address.

```
[edit security nat static]
user@host# set rule-set rs rule r1 match destination-address 203.0.113.1/32
user@host# set rule-set rs rule r1 match destination-port 100 to 200
```

```

user@host# set rule-set rs rule r1 then static-nat prefix 10.1.1.1/32
user@host# set rule-set rs rule r1 then static-nat prefix mapped-port 300 to 400

```

3. Configure a rule that matches packets and translates the destination address in the packets to a private address.

```

[edit security nat static]
user@host# set rule-set rs rule r2 match destination-address 203.0.113.1/32
user@host# set rule-set rs rule r2 match destination-port 300 to 400
user@host# set rule-set rs rule r2 then static-nat prefix 10.1.1.2/32
user@host# set rule-set rs rule r2 then static-nat prefix mapped-port 300 to 400

```

4. Configure a rule that matches packets and translates the destination address in the packets to a private address.

```

[edit security nat static]
user@host# set rule-set rs rule r3 match destination-address 203.0.113.3/32
user@host# set rule-set rs rule r3 match destination-port 300
user@host# set rule-set rs rule r3 then static-nat prefix 10.1.1.2/32
user@host# set rule-set rs rule r3 then static-nat prefix mapped-port 200

```

Results

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

```
[edit]
```

```
user@host# show security nat
```

```

security {
  nat {
    static {
      rule-set rs {
        from zone untrust;
        rule r1 {
          match {
            destination-address 203.0.113.1/32;
            destination-port 100 to 200;

```


If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

Verifying Static NAT Configuration | 179

Verifying Static NAT Configuration

Purpose

Verify that there is traffic matching the static NAT rule set.

Action

From operational mode, enter the `show security nat static rule` command. View the Translation hits field to check for traffic that matches the rule.

```
user@host> show security nat static rule all
Total static-nat rules: 3

Static NAT rule: r2                Rule-set: rs
Rule-Id                           : 3
Rule position                      : 2
From zone                         : untrust
Destination addresses              : 203.0.113.1
Destination ports                  : 300 - 400
Host addresses                    : 10.1.1.2
Host ports                        : 300 - 400
Netmask                           : 32
Host routing-instance             : N/A
Translation hits                   : 0

Static NAT rule: r3                Rule-set: rs
Rule-Id                           : 4
Rule position                      : 3
From zone                         : untrust
Destination addresses              : 203.0.113.3
```

```

Destination ports      : 300 - 300
Host addresses         : 10.1.1.2
Host ports             : 200 - 200
Netmask                : 32
Host routing-instance : N/A
Translation hits       : 0

```

```

Static NAT rule: r1          Rule-set: rs
Rule-Id                     : 9
Rule position                : 1
From zone                    : untrust
Destination addresses        : 203.0.113.1
Destination ports            : 100 - 200
Host addresses               : 10.1.1.1
Host ports                   : 300 - 400
Netmask                      : 32
Host routing-instance        : N/A
Translation hits              : 0

```

Troubleshooting

IN THIS SECTION

- [Troubleshooting Static NAT Port Configuration | 180](#)

Troubleshooting Static NAT Port Configuration

Problem

Static NAT port mapping configuration failures occur during a commit.

Invalid configurations with overlapped IP addresses and ports result in commit failure.

The following example shows invalid configurations with overlapped addresses and ports:

- **set security nat static rule-set rs rule r1 match destination-address 203.0.113.1**
set security nat static rule-set rs rule r1 then static-nat prefix 10.1.1.1
- **set security nat static rule-set rs rule r2 match destination-address 203.0.113.1**

```
set security nat static rule-set rs rule r2 match destination-port 300 to 400
```

```
set security nat static rule-set rs rule r2 then static-nat prefix 10.1.1.2
```

```
set security nat static rule-set rs rule r2 then static-nat prefix mapped-port 300 to 400
```

- ```
set security nat static rule-set rs rule r1 match destination-address 203.0.113.1
```

```
set security nat static rule-set rs rule r1 match destination-port 100 to 200
```

```
set security nat static rule-set rs rule r1 then static-nat prefix 10.1.1.1
```

```
set security nat static rule-set rs rule r1 then static-nat prefix mapped-port 300 to 400
```

- ```
set security nat static rule-set rs rule r2 match destination-address 203.0.113.2
```

```
set security nat static rule-set rs rule r2 match destination-port 300 to 400
```

```
set security nat static rule-set rs rule r2 then static-nat prefix 10.1.1.1
```

```
set security nat static rule-set rs rule r2 then static-nat prefix mapped-port 390 to 490
```

The following error message was displayed when the aforementioned configuration was submitted for commit:

```
error: 'prefix/mapped-port' of static nat rule r2 overlaps with 'prefix/mapped-port' of static
nat rule r1
error: configuration check-out failed
```

Solution

To configure the destination port, you must avoid any address overlapping or any address and port overlapping. For an example of valid configuration, see ["Configuration" on page 175](#)

Monitoring Static NAT Information

IN THIS SECTION

● [Purpose | 182](#)

● [Action | 182](#)

Purpose

View static NAT rule information.

Action

Select **Monitor>NAT>Static NAT** in the J-Web user interface, or enter the following CLI command:

```
show security nat static rule
```

[Table 9 on page 182](#) summarizes key output fields in the static NAT display.

Table 9: Summary of Key Static NAT Output Fields

Field	Values	Action
Rule-set Name	Name of the rule set.	Select all rule sets or a specific rule set to display from the list.
Total rules	Number of rules configured.	–
ID	Rule ID number.	–
Position	Position of the rule that indicates the order in which it applies to traffic.	–
Name	Name of the rule.	–
Ruleset Name	Name of the rule set.	–
From	Name of the routing instance/interface/zone from which the packet comes	–
Source addresses	Source IP addresses.	–
Source ports	Source port numbers.	–

Table 9: Summary of Key Static NAT Output Fields (Continued)

Field	Values	Action
Destination addresses	Destination IP address and subnet mask.	–
Destination ports	Destination port numbers .	–
Host addresses	Name of the host addresses.	–
Host ports	Host port numbers.	
Netmask	Subnet IP address.	–
Host routing instance	Name of the routing instance from which the packet comes.	–
Alarm threshold	Utilization alarm threshold.	–
Sessions (Succ/ Failed/ Current)	Successful, failed, and current sessions. <ul style="list-style-type: none"> • Succ–Number of successful session installations after the NAT rule is matched. • Failed–Number of unsuccessful session installations after the NAT rule is matched. • Current–Number of sessions that reference the specified rule. 	–
Translation hits	Number of times a translation in the translation table is used for a static NAT rule.	–

Table 9: Summary of Key Static NAT Output Fields *(Continued)*

Field	Values	Action
Top 10 Translation Hits Graph	Displays the graph of top 10 translation hits.	-

3

CHAPTER

NAT Configuration Options

Persistent NAT and NAT64 | 186

NAT for Multicast Flows | 228

IPv6 NAT | 242

IPv6 Dual-Stack Lite | 268

NAT for VRF Routing Instance | 274

NAT for VRF group | 297

Persistent NAT and NAT64

IN THIS SECTION

- [Understanding Persistent NAT and NAT64 | 186](#)
- [Understanding Session Traversal Utilities for NAT \(STUN\) Protocol | 188](#)
- [Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation | 189](#)
- [Persistent NAT and NAT64 Configuration Overview | 191](#)
- [Example: Configuring Address Persistent NAT64 Pools | 193](#)
- [Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT | 196](#)
- [Example: Configuring Address-Dependent Filtering for IPv6 Clients | 205](#)
- [Example: Configuring Endpoint-Independent Filtering for IPv6 Clients | 210](#)
- [Example: Setting Maximum Persistent NAT Bindings | 216](#)
- [Persistent NAT Hairpinning Overview | 218](#)
- [Example: Configuring Persistent NAT Hairpinning with Source NAT Pool with Address Shifting | 220](#)

Network Address Translators (NATs) are well known to cause very significant problems with applications that carry IP addresses in the payload. Applications that suffer from this problem include Voice Over IP and Multimedia Over IP. Persistent NAT improves NATs behavior and defines a set of NAT requirement behavior which is useful for VOIP applications working. NAT64 is a translating mechanism used to translate IPv6 packets to IPv4 packets and vice versa by translating the packet headers according to IP/ICMP Translation Algorithm.

Understanding Persistent NAT and NAT64

Persistent NAT allows applications to use the Session Traversal Utilities for NAT (STUN) protocol when passing through NAT firewalls. Persistent NAT ensures that all requests from the same internal transport address (internal IP address and port) are mapped to the same reflexive transport address (the public IP address and port created by the NAT device closest to the STUN server).

NAT64 is a mechanism for translating IPv6 packets to IPv4 packets and vice versa that allows IPv6 clients to contact IPv4 servers using unicast UDP, TCP, or ICMP. It is an enhancement of Network Address Translation-Protocol Translation (NAT-PT).

NAT64 supports the following:

- Endpoint-independent mappings
- Endpoint-independent filtering and address-dependent filtering

NOTE: The mapping and filtering behaviors of NAT64 and persistent NAT are identical.

The following types of persistent NAT can be configured on the Juniper Networks device:

- Any remote host—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. Any external host can send a packet to the internal host by sending the packet to the reflexive transport address.
- Target host—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address.
- Target host port—All requests from a specific internal IP address and port are mapped to the same reflexive transport address. An external host can send a packet to an internal host by sending the packet to the reflexive transport address. The internal host must have previously sent a packet to the external host's IP address and port.

NOTE: The target-host-port configuration is not supported for NAT64 when configured with IPv6 address.

You configure any of the persistent NAT types with source NAT rules. The source NAT rule action can use a source NAT pool (with or without port translation) or an egress interface. Persistent NAT is not applicable for destination NAT, because persistent NAT bindings are based on outgoing sessions from internal to external.

NOTE: Port overloading is used in Junos OS only for normal interface NAT traffic. Persistent NAT does not support port overloading, and you must explicitly disable port overloading with one of the following options at the [edit security nat source] hierarchy level:

- port-overloading off
- port-overloading-factor 1

To configure security policies to permit or deny persistent NAT traffic, you can use two new predefined services—`junos-stun` and `junos-persistent-nat`.

NOTE: Persistent NAT is different from the persistent address feature (see "[Understanding Persistent Addresses for Source NAT Pools](#)" on page 84). The persistent address feature applies to address mappings for source NAT pools configured on the device. The persistent NAT feature applies to address mappings on an external NAT device, and is configured for a specific source NAT pool or egress interface. Also, persistent NAT is intended for use with STUN client/server applications.

[Table 10 on page 188](#) list the persistent NAT binding information for SRX4100, SRX4200, and vSRX.

Table 10: Persistent NAT Binding Support

Device	Persistent NAT Binding Support
SRX4200	1M (1048576)
SRX4100	512K (524288)
vSRX with 64GB	512K (524288)

Understanding Session Traversal Utilities for NAT (STUN) Protocol

Many video and voice applications do not work properly in a NAT environment. For example, Session Initiation Protocol (SIP), used with VoIP, encodes IP addresses and port numbers within application data. If a NAT firewall exists between the requestor and receiver, the translation of the IP address and port number in the data invalidates the information.

Also, a NAT firewall does not maintain a pinhole for incoming SIP messages. This forces the SIP application to either constantly refresh the pinhole with SIP messages or use an ALG to track registration, a function that may or may not be supported by the gateway device.

The Session Traversal Utilities for NAT (STUN) protocol, first defined in *RFC 3489, Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs)* and then later in *RFC 5389, Session Traversal Utilities for NAT*, is a simple client/server protocol. A STUN client sends requests to a STUN server, which returns responses to the client. A STUN client is usually part of an

application that requires a public IP address and/or port. STUN clients can reside in an end system such as a PC or in a network server whereas STUN servers are usually attached to the public Internet.

NOTE: Both the STUN client and STUN server must be provided by the application. Juniper Networks does not provide a STUN client or server.

The STUN protocol allows a client to:

- Discover whether the application is behind a NAT firewall.
- Determine the type of NAT binding being used.
- Learn the reflexive transport address, which is the IP address and port binding allocated by NAT device closest to the STUN server. (There may be multiple levels of NAT between the STUN client and the STUN server.)


The client application can use the IP address binding information within protocols such as SIP and H.323.

Understanding NAT64 IPv6 Prefix to IPv4 Address-Persistent Translation

The NAT64 mechanism enables IPv6 clients to contact IPv4 servers by translating IPv6 addresses to IPv4 addresses (and vice versa). However, some IPv4 applications and services cannot work correctly over IPv6-only networks with standard NAT64 in a dual-translation scenario, such as *464XLAT*. In those scenarios, address-persistent translation is required.

[Figure 15 on page 190](#) illustrates the 464XLAT architecture, whereby IPv4 packets are translated to IPv6 packets on the customer-side translator (CLAT), then go across the IPv6-only network, and are translated back to IPv4 packets on the provider-side translator (PLAT) to access global IPv4-only content in the core network. This architecture uses a combination of stateless translation on the CLAT and stateful translation on the PLAT.

Figure 15: 464XLAT Architecture



When a device functions as a PLAT, it is responsible for keeping the sticky mapping relationship between one specific IPv6 prefix and one translated IPv4 address. The device treats the IPv6 prefix as a single user. This mapping is accomplished by configuring the specific IPv6 prefix length in an IPv4 source NAT pool using the address-persistent feature.

Figure 16 on page 190 illustrates a NAT rule configured in the CLAT, which translates an IPv4 address to an IPv6 address with an address-persistent prefix. With stateless NAT46 translation on the CLAT and stateful NAT64 translation on the PLAT, the traffic from IPv4 host 192.168.1.2 reaches the global server 198.51.100.1 over an IPv6-only network.

Figure 16: NAT64 Translation on the PLAT




Table 11 on page 191 lists other NAT features and their compatibility with the address-persistent feature.

Table 11: NAT Feature Compatibility with the Address Persistent Feature

Feature			Compatible
PAT pools	IPv4	NAT IPv4 to IPv6	No
		NAT IPv6 to IPv4	Yes
	IPv6	NAT IPv4 to IPv6	No
		NAT IPv6 to IPv4	No
Non-PAT pools			No
Port-overloading			Yes
Persistent NAT in PAT pool			Yes
Port block allocation			Yes
Deterministic NAT			No
Address pooling paired			No
ALG (Existing ALG NAT translations , such as FTP/PPTP/RTSP/DNS/SIP from native IPv6 clients.)			Yes

Persistent NAT and NAT64 Configuration Overview

To configure persistent NAT, specify the following options with the source NAT rule action (for either a source NAT pool or an egress interface):

- The type of persistent NAT—One of the following: any remote host, target host, or target host port.

- (Optional) Address mapping—This option allows requests from a specific internal IP address to be mapped to the same reflexive IP address; internal and reflexive ports can be any ports. An external host using any port can send a packet to the internal host by sending the packet to the reflexive IP address (with a configured incoming policy that allows external to internal traffic). If this option is not configured, the persistent NAT binding is for specific internal and reflexive transport addresses.

You can only specify the address-mapping option when the persistent NAT type is any remote host and the source NAT rule action is one of the following actions:

- Source NAT pool with IP address shifting
- Source NAT pool with no port translation and no overflow pool
- (Optional) Inactivity timeout—Time, in seconds, that the persistent NAT binding remains in the device's memory when all the sessions of the binding entry have expired. When the configured timeout is reached, the binding is removed from memory. The default value is 300 seconds. Configure a value from 60 through 7200 seconds.

When all sessions of a persistent NAT binding have expired, the binding remains in a query state in the device's memory for the specified inactivity timeout period. The query binding is automatically removed from memory when the inactivity timeout period expires (the default is 300 seconds). You can explicitly remove all or specific persistent NAT query bindings with the `clear security nat source persistent-nat-table` command.

- (Optional) Maximum session number—Maximum number of sessions with which a persistent NAT binding can be associated. The default is 30 sessions. Configure a value from 8 through 100.

For interface NAT, you need to explicitly disable port overloading with one of the following options at the `[edit security nat source]` hierarchy level:

- `port-overloading off`
- `port-overloading-factor 1`

Finally, there are two predefined services that you can use in security policies to permit or deny STUN and persistent NAT traffic:

- `junos-stun`—STUN protocol traffic.
- `junos-persistent-nat`—Persistent NAT traffic.

For the any remote host persistent NAT type, the direction of the security policy is from external to internal. For target host or target host port persistent NAT types, the direction of the security policy is from internal to external.

Example: Configuring Address Persistent NAT64 Pools

IN THIS SECTION

- [Requirements | 193](#)
- [Overview | 193](#)
- [Configuration | 193](#)
- [Verification | 196](#)

This example shows how to configure address persistent NAT64 pools to ensure a sticky mapping relationship between one specific IPv6 prefix, which is calculated by the configured IPv6 prefix length, and one translated IPv4 address.

Requirements

Before you begin, be sure the existing NAT rules and pool configuration do not conflict with the new one.

Overview

In this example, you configure an IPv6 prefix length of /64 in an IPv4 source NAT pool for NAT IPv6 to IPv4 translations. Traffic matching the NAT rule and NAT pool perform address persistent translation between the IPv6 prefix and the IPv4 translated address. This configuration can be used on the provider-side translator (PLAT) in a dual-translation scenario, *464XLAT*, to enable IPv4 services to work over IPv6-only networks.

Configuration

IN THIS SECTION

- [CLI Quick Configuration | 194](#)
- [Procedure | 194](#)

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool NAT64 address 198.51.100.240/32 to 198.51.100.254/32
set security nat source pool NAT64 address-persistent subscriber ipv6-prefix-length 64
set security nat source rule-set RS1 from zone trust
set security nat source rule-set RS1 to zone untrust
set security nat source rule-set RS1 rule R1 match source-address 2001:db8::/32
set security nat source rule-set RS1 rule R1 match destination-address 198.51.100.198/32
set security nat source rule-set RS1 rule R1 then source-nat pool NAT64
```

Procedure

Step-by-Step Procedure

The following example requires you to navigate throughout various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

1. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool NAT64 address 198.51.100.240/32 to 198.51.100.254/32
```

2. Specify the IPv6 prefix length for the source NAT pool.

```
[edit security nat source]
user@host# set pool NAT64 address-persistent subscriber ipv6-prefix-length 64
```

3. Create a rule set.

```
[edit security nat source]
user@host# set rule-set RS1 from zone trust
user@host# set rule-set RS1 to zone untrust
```

4. Match the rule.

```
[edit security nat source]
user@host# set rule-set RS1 rule R1 match source-address 2001:db8::/32
user@host# set rule-set RS1 rule R1 match destination-address 198.51.100.198/32
```

5. Provide the action to be performed when the rule matches.

```
[edit security nat source]
user@host# set security nat source rule-set RS1 rule R1 then source-nat pool NAT64
```

Results

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

```
[edit]
user@host# show security nat
source {
  pool NAT64 {
    address {
      198.51.100.240/32 to 198.51.100.254/32;
    }
    address-persistent subscriber ipv6-prefix-length 64;
  }
  rule-set RS1 {
    from zone trust;
    to zone untrust;
    rule R1 {
      match {
        source-address 2001:db8::/32;
        destination-address 198.51.100.198/32;
      }
      then {
        source-nat {
          pool {
            NAT64;
          }
        }
      }
    }
  }
}
```

```

    }
  }
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying NAT Application to Traffic | 196](#)

Verifying NAT Application to Traffic

Purpose

Verify that the same IPv6 prefix is translated to the persistent IPv4 address.

Action

From operational mode, enter the `show security flow session` command.

Example: Supporting Network Configuration By Configuring Persistent NAT with Interface NAT

IN THIS SECTION

- [Requirements | 197](#)
- [Overview | 197](#)
- [Configuration | 199](#)
- [Verification | 203](#)

You can configure any of the persistent *NAT* types with source NAT rules. This example illustrates how to apply persistent NAT with an interface IP address and how to use an interface IP address as a NAT IP address to perform persistent NAT for a specific internal host. It also shows how to maintain persistent address *port mapping* behavior and persistent NAT filter behavior for the host. You must disable port overloading for interface NAT.

Requirements

This example uses the following hardware and software components:

- 1 SRX Series Firewall
- 4 PCs

Before you begin:

- Understand the concepts of persistent NAT. See ["Persistent NAT and NAT64 Configuration Overview" on page 191](#).

Overview

IN THIS SECTION

- [Topology | 198](#)

In a Carrier Grade NAT (CGN) network deployment, you can configure the interface IP address as a NAT address to perform persistent network address translation. In this way, the internal host can create one source NAT mapping relationship by the outgoing traffic initiated from internal to external. Then the external host sends traffic back to this internal host by sending the traffic to this interface NAT address through the shared NAT mapping relationship.

In this example, you first configure the interface NAT rule set int1 to match traffic from interface ge-0/0/1 to interface ge-0/0/2, and then you configure the NAT rule in1 to match the specific source and destination addresses to perform persistent NAT. You configure the any remote host persistent NAT type when interface NAT is performed.

For packets with source address 192.0.2.0/24 (internal phones) and destination address 198.51.100.0/24 (including *STUN* server, *SIP* proxy server, and external phones), you configure interface NAT with the any remote host persistent NAT type. Then you disable port overloading for interface NAT.

Next, you configure a security policy to allow persistent NAT traffic from the external network (external zone) to the internal network (internal zone) for any of the remote host persistent NAT types.

Topology

Figure 17 on page 198 shows an interface persistent NAT topology.

Figure 17: Interface Persistent NAT Topology

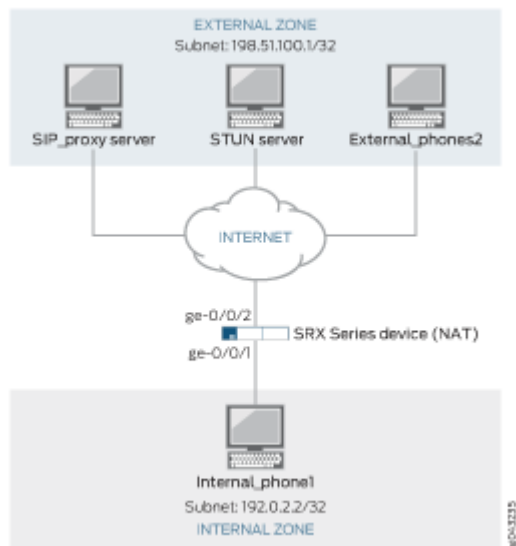


Table 12 on page 198 shows the parameters configured in this example.

Table 12: Interfaces, Zones, Servers, and IP Address Information

Parameter	Description
External Zone	External network
Internal Zone	Internal network
External_phones2	Phone2 address of external network
Internal_phone1	Phone1 address of internal network
SIP_proxy server	SIP proxy server address of external network
STUN server	STUN server address of external network

Table 12: Interfaces, Zones, Servers, and IP Address Information *(Continued)*

Parameter	Description
Subnet 198.51.100.1/32	Destination IP address
Subnet 192.0.2.2/32	Source IP address
ge-0/0/1 and ge-0/0/2	NAT interfaces for traffic direction

Configuration

IN THIS SECTION

- [Procedure](#) | 199

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```

set security nat source rule-set int1 from interface ge-0/0/1.0
set security nat source rule-set int1 to interface ge-0/0/2.0
set security nat source rule-set int1 rule in1 match source-address 192.0.2.0/24
set security nat source rule-set int1 rule in1 match destination-address 198.51.100.0/24
set security nat source rule-set int1 rule in1 then source-nat interface persistent-nat permit
any-remote-host
set security nat source interface port-overloading off
set security policies from-zone internal to-zone external policy stun_traffic match source-
address internal_phones destination-address stun_server application junos-stun
set security policies from-zone internal to-zone external policy sip_proxy_traffic match source-
address internal_phones destination-address sip_proxy_server application junos-sip
set security policies from-zone internal to-zone external policy sip_traffic match source-

```

```

address internal_phones destination-address external_phones application junos-persistent-nat
set security policies from-zone internal to-zone external policy sip_traffic then permit
set security policies from-zone internal to-zone external policy stun_traffic then permit
set security policies from-zone internal to-zone external policy sip_proxy_traffic then permit

```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure an interface NAT rule set:

1. Create a persistent NAT rule for an interface NAT.

```

[edit security nat source rule-set int1]
user@host# set from interface ge-0/0/1.0
user@host# set to interface ge-0/0/2.0
user@host# set rule in1 match source-address 192.0.2.0/24
user@host# set rule in1 match destination-address 198.51.100.0/24
user@host# set rule in1 then source-nat interface persistent-nat permit any-remote-host

```

2. Disable port overloading for interface NAT.

```

[edit security]
user@host# set nat source interface port-overloading off

```

3. Configure a security policy to allow STUN traffic from internal SIP phones to an external STUN server.

```

[edit security policies]
user@host# set from-zone internal to-zone external policy stun_traffic match source-address
internal_phones destination-address stun_server application junos-stun

```

4. Configure a security policy to allow SIP proxy traffic from internal SIP phones to an external SIP proxy server.

```
[edit security policies]
user@host# set from-zone internal to-zone external policy sip_proxy_traffic match source-
address internal_phones destination-address sip_proxy_server application junos-sip
```

5. Configure a security policy to allow SIP traffic from external SIP phones to internal SIP phones.

```
[edit security policies]
user@host# set from-zone internal to-zone external policy sip_traffic match source-address
internal_phones destination-address external_phones application junos-persistent-nat
user@host# set from-zone internal to-zone external policy sip_traffic then permit
user@host#set from-zone internal to-zone external policy stun_traffic then permit
user@host#set from-zone internal to-zone external policy sip_proxy_traffic then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@host# show security nat
source {
    interface {
        port-overloading off;
    }
    rule-set int1 {
        from interface ge-0/0/1.0;
        to interface ge-0/0/2.0;
        rule in1 {
            match {
                source-address 192.0.2.0/24;
                destination-address 198.51.100.0/24;
            }
            then {
                source-nat {
                    interface {
```

```

        persistent-nat {
            permit any-remote-host;
        }
    }
}

[edit]
user@host# show security policies
from-zone internal to-zone external {
    policy stun_traffic {
        match {
            source-address internal_phones;
            destination-address stun_server;
            application junos-stun;
        }
        then {
            permit;
        }
    }
    policy sip_proxy_traffic {
        match {
            source-address internal_phones;
            destination-address sip_proxy_server;
            application junos-sip;
        }
        then {
            permit;
        }
    }
    policy sip_traffic {
        match {
            source-address internal_phones;
            destination-address external_phones;
            application junos-persistent-nat;
        }
        then {
            permit;
        }
    }
}

```

```
    }  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That Rules Are Matched and Used | 203](#)
- [Verifying That NAT Traffic Sessions Are Established | 204](#)

Confirm that the configuration is working properly.

Verifying That Rules Are Matched and Used

Purpose

Verify that all the rules are matched and used.

Action

From operational mode, enter the `show security nat source persistent-nat-table all` command.

```
user@host>show security nat source persistent-nat-table all  
Internal  Reflective      Source      Type      Left_time/Curr_Sess_Num/  Source  
  In_IP    In_Port I_Proto Ref_IP    Ref_Port R_Proto NAT Pool  Conf_time Max_Sess_Num  
NAT Rule  
  192.0.2.12 17012  udp   198.51.100.1 28153    udp   interface  any-remote-host  
3528/3600    -/-      in1  
  192.0.2.12 7078   udp   198.51.100.1 6133     udp   interface  any-remote-host  -/  
300         1/30      in1
```

Meaning

The output displays a summary of persistent NAT information.

Verifying That NAT Traffic Sessions Are Established

Purpose

Verify that the sessions are established on the device.

Action

From operational mode, enter the `show security flow session` command.

```
user@host>show security flow session
```

```
Session ID: 6992, Policy name: sip_proxy_traffic/5, Timeout: 16, Valid
```

```
In: 192.0.2.12/17012 --> 198.51.100.45/5060;udp, If: ge-0/0/1.0, Pkts: 4, Bytes: 1850
```

```
Out: 198.51.100.45/5060 --> 198.51.100.1/28153;udp, If: ge-0/0/2.0, Pkts: 5, Bytes: 2258
```

```
Session ID: 7382, Policy name: stun_traffic/4, Timeout: 16, Valid
```

```
In: 192.0.2.12/7078 --> 198.51.100.49/3478;udp, If: ge-0/0/1.0, Pkts: 20, Bytes: 1040
```

```
Out: 198.51.100.49/3478 --> 198.51.100.1/6133;udp, If: ge-0/0/2.0, Pkts: 0, Bytes: 0
```

Meaning

The `show security flow session` command displays active sessions on the device and each session's associated security policy. The output shows traffic entering the device using the private source address 192.0.2.12 destined to a public host at 198.51.100.45. The return traffic from this flow travels to the translated public address 198.51.100.1.

- **Session ID**—Number that identifies the session. Use this ID to get more information about the session such as policy name or number of packets in and out.
- **sip_proxy_traffic**— Policy name that permitted the SIP traffic from the internal SIP phones to the external SIP proxy server.
- **In**—Incoming flow (source and destination IP addresses with their respective source and destination port numbers. The session is UDP, and the source interface for this session is ge-0/0/1.0).
- **Out**—Reverse flow (source and destination IP addresses with their respective source and destination port numbers. The session is UDP, and the destination interface for this session is ge-0/0/2.0).
- **stun_traffic**—Policy name that permitted the STUN traffic from the internal SIP phones to the external STUN server.

Example: Configuring Address-Dependent Filtering for IPv6 Clients

IN THIS SECTION

- [Requirements | 205](#)
- [Overview | 205](#)
- [Configuration | 206](#)
- [Verification | 209](#)

This example shows how to configure address-dependent filtering for IPv6 clients using NAT64.

Requirements

Before you begin:

- Ensure that IPv6 is enabled on the device.
- Ensure that the existing NAT rule and pool configuration do not conflict with the new ones.

Overview

IN THIS SECTION

- [Topology | 205](#)

In this example you use NAT64 to send packets from the IPv6 internal host to the IPv4 external host and from the IPv4 external host to the IPv4 internal host.

Topology

Configuration

IN THIS SECTION

- [Procedure](#) | 206

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address 2001:db8::/128
set security nat static rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
set security nat source pool myipv4 address 203.0.113.2
set security nat source rule-set myipv4_rs from interface ge-0/0/1
set security nat source rule-set myipv4_rs to interface ge-0/0/2
set security nat source rule-set myipv4_rs rule ipv4_rule match source-address 2001:db8::/96
set security nat source rule-set myipv4_rs rule ipv4_rule match destination-address 10.2.2.15
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool persistent-nat
permit target-host
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure address-dependent filtering for IPv6 clients:

1. Create a set of rules for NAT64.

```
[edit security nat static]
user@host# set rule-set test_rs from interface ge-0/0/1
```

2. Match the rule.

```
[edit security nat static]
user@host# set rule-set test_rs rule test_rule match destination-address 2001:db8::/128
```

3. Provide the action to be performed when the rule matches.

```
[edit security nat static]
user@host# set rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
```

4. Define a source address pool and add the address to the pool.

```
[edit security nat]
user@host# set source pool myipv4 address 203.0.113.2
```

5. Create another set of rules for NAT64.

```
[edit security nat]
user@host# set source rule-set myipv4_rs from interface ge-0/0/1
```

6. Match the rule with the source address.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule match source-address 2001:db8::/96
```

7. Match the rule with the destination address.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule match destination-address 10.2.2.15
```

8. Provide the action to be performed when the rules match.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
```

9. Configure persistent NAT.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool persistent-nat
permit target-host
```

Results

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

```
[edit security]
user@host#show nat source
pool myipv4 {
    address {
        203.0.113.2/32;
    }
}
rule-set test_rs {
    rule test_rule {
        match {
            destination-address 2001:db8::/128;
        }
    }
}

rule-set myipv4_rs {
    from interface ge-0/0/1.0;
    to interface ge-0/0/2.0;
    rule ipv4_rule {
        match {
            source-address 2001:db8::/96;
            destination-address 10.2.2.15/32;
        }
        then {
            source-nat {
                pool {
                    myipv4;
                    persistent-nat {
```

```

    permit target-host;
  }
}
}
}
}
}
}
}
}
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That the Configuration Is Enabled and Working | 209](#)
- [Verifying That Rules Are Matched and Used | 210](#)

Confirm that the configuration is working properly:

Verifying That the Configuration Is Enabled and Working

Purpose

Verify that the configuration is enabled and working.

Action

From operational mode, enter the following commands:

- `show security nat static rule test_rule`
- `show security nat source rule ipv4_rule`
- `show security nat source pool myipv4`

Verifying That Rules Are Matched and Used

Purpose

Verify that all the rules are matched and used.

Action

From operational mode, enter the `show security nat source persistent-nat-table all` command.

Example: Configuring Endpoint-Independent Filtering for IPv6 Clients

IN THIS SECTION

- [Requirements | 210](#)
- [Overview | 210](#)
- [Configuration | 211](#)
- [Verification | 215](#)

This example shows how to configure endpoint-independent filtering for IPv6 clients using NAT64.

Requirements

Before you begin:

- Ensure that IPv6 is enabled on the device
- Ensure that the existing NAT rules and pool configuration do not conflict with the new ones.

Overview

IN THIS SECTION

- [Topology | 211](#)

In this example you use NAT64 to send packets from the IPv6 internal host to the IPv4 external host and from the IPv4 external host to the IPv4 internal host.

Topology

Configuration

IN THIS SECTION

- [Procedure | 211](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address 2001:db8::/128
set security nat static rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
set security nat source pool myipv4 address 203.0.113.2
set security nat source rule-set myipv4_rs from interface ge-0/0/1
set security nat source rule-set myipv4_rs to interface ge-0/0/2
set security nat source rule-set myipv4_rs rule ipv4_rule match source-address 2001:db8::/96
set security nat source rule-set myipv4_rs rule ipv4_rule match destination-address 10.2.2.15
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool persistent-nat
permit any-remote-host
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure endpoint-independent filtering for IPv6 clients:

1. Create a set of rules for NAT64.

```
[edit security nat static]
user@host# set rule-set test_rs from interface ge-0/0/1
```

2. Match the rule.

```
[edit security nat static]
user@host# set rule-set test_rs rule test_rule match destination-address 2001:db8::/128
```

3. Provide the action to be performed when the rule matches.

```
[edit security nat static]
user@host# set rule-set test_rs rule test_rule then static-nat prefix 10.2.2.15/32
```

4. Define a source address pool and add the address to the pool.

```
[edit security nat]
user@host# set source pool myipv4 address 203.0.113.2
```

5. Create another set of rules for NAT64.

```
[edit security nat]
user@host# set source rule-set myipv4_rs from interface ge-0/0/1
```

6. Match the rule with the source address.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule match source-address 2001:db8::/96
```

7. Match the rule with the destination address.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule match destination-address 10.2.2.15
```

8. Provide the action to be performed when the rules match.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
```

9. Configure persistent NAT.

```
[edit security nat]
user@host# set source rule-set myipv4_rs rule ipv4_rule then source-nat pool persistent-nat
permit any-remote-host
```

Results

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

```
[edit security]
user@host#show security nat
source {
  pool myipv6_prefix {
    address {
      2001:db8::/64;
    }
  }
  pool myipv4 {
    address {
      203.0.113.2/32;
    }
  }
  rule-set myipv6_rs {
    from interface ge-0/0/1.0;
    to interface ge-0/0/2.0;
    rule ipv6_rule {
      match {
        source-address 10.1.1.0/30;
        destination-address 2001:db8::2/96;
      }
      then {
        source-nat {
```



```

        pool {
            myipv6_prefix;
        }
    }
}

rule-set myipv4_rs {
    from interface ge-0/0/1.0;
    to interface ge-0/0/2.0;
    rule ipv4_rule {
        match {
            source-address 2001:db8::/96;
            destination-address 10.2.2.15/32;
        }
        then {
            source-nat {
                pool {
                    myipv4;
                    persistent-nat {
                        permit target-host;
                    }
                }
            }
        }
    }
}

static {
    rule-set test_rs {
        from interface ge-0/0/1.0;
        rule test_rule {
            match {
                destination-address 2001:db8::/128;
            }
            then {
                static-nat {
                    prefix {
                        10.2.2.15/32;
                    }
                }
            }
        }
    }
}

```

```
}
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That the Configuration is Enabled and Working | 215](#)
- [Verifying That Rules Are Matched and Used | 215](#)

Confirm that the configuration is working properly:

Verifying That the Configuration is Enabled and Working

Purpose

Verify that the configuration is enabled and working.

Action

From operational mode, enter the following commands.

- `show security nat static rule test_rule`
- `show security nat source rule ipv4_rule`
- `show security nat source pool myipv4`

Verifying That Rules Are Matched and Used

Purpose

Verify that all the rules are matched and used.

Action

From operational mode, enter the `show security nat source persistent-nat-table all` command.

Example: Setting Maximum Persistent NAT Bindings

IN THIS SECTION

- Requirements | 216
- Overview | 216
- Configuration | 217
- Verification | 218

This example shows how to increase the persistent NAT capacity.

Requirements

Before you begin, see ["Understanding Persistent NAT and NAT64" on page 186](#).

Overview

In this example, you enable the maximize persistent NAT capacity option. This option is supported only on Services Processing Cards (SPCs) for SRX1400 devices with SRX1K-NPC-SPC-1-10-40, SRX3000 line with SRX3K-SPC-1-10-40, and SRX5000 line devices with SRX5K-SPC-2-10-40SPC and SRX5K-SPC3. Note that for the SRX5000 line devices with SRX5K-SPC-2-10-40SPC and SPC3, the persistent NAT binding number is maximized at the cost of reducing the maximum session number.

To enable this option, the supported central point maximum binding capacity can be approximately increased to 1/8 of the central point session capacity up to 2M and the supported SPU maximum binding capacity can be approximately increased to 1/4 of each SPU session capacity. Accordingly, the flow session capacity will decrease by 1/4 on both the CP and each of the SPU.

By default, the persistent NAT binding capacity on both the central point and the SPU of an SRX5400, SRX5600, or SRX5800 device is 64,000. In this example, you enable the session capacity to maximum 20,000,000 on the central point and maximum 1,100,000 on each of the SPUs with maximum session configuration. If you enable the `maximize-persistent-nat-capacity` option, an SRX5400, SRX5600, or SRX5800 device with 4 GB of memory can support maximum 2M persistent NAT bindings on the central point and 275,000 bindings on each of the SPUs.

Configuration

IN THIS SECTION

- [Procedure](#) | 217

Procedure

Step-by-Step Procedure

To increase the persistent NAT capacity:

1. Set maximize persistent NAT capacity option.

```
[edit]
user@host# set security forwarding-process application-services maximize-persistent-nat-
capacity
```

2. If you are done configuring the device, commit the configuration.

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

3. Restart the system from operational mode.

```
[edit]
user@host# request system reboot
```

NOTE: When switching to maximize persistent NAT capacity mode or back to regular mode, you must restart the device.

4. If you want to switch the device back to regular mode, delete the maximize persistent NAT capacity mode configuration.

```
[edit]
user@host# delete security forwarding-process application-services maximize-persistent-nat-
capacity
```

Verification

IN THIS SECTION

- [Verifying Increased Persistent NAT Capacity | 218](#)

Verifying Increased Persistent NAT Capacity

Purpose

Verify that you have increased the persistent NAT capacity.

Action

From operational mode, enter the `show security forwarding-process application-services` command.

Persistent NAT Hairpinning Overview

When traffic is sent between two hosts, the source host of the traffic may only know the destination host by its public IP address. In reality, the destination host may be in the same private address space as the source host. Hairpinning is the process of returning the traffic in the direction from where it came from as a way to get it to its destination host in a private subnetwork.

Generally, a source host in a subnetwork may not recognize that the traffic is intended for a destination host within the same subnetwork, because it identifies the destination host only by its public IP address. The NAT analyzes the IP packets and routes the packet back to the correct host.


NAT hairpinning support is required if two hosts on the internal network want to communicate with each other by using a binding on the NAT device. In this case, the NAT device receives a packet from the

internal network and forwards it back to the internal network. If hairpinning is not supported, forwarding the packet will fail and it will be dropped.

Hairpinning enables two endpoints (Host 1 and Host 2) on the private network to communicate even if they only use each other's external IP addresses and ports. When Host 1 sends traffic to Host 3, a NAT binding between Host 1's internal source IP address and port is associated in the NAT table with its external IP address and port. The same thing happens when Host 2 sends traffic to Host 3. In this way, when Host 1 and Host 2 want to communicate, they can identify each other's external IP addresses.

For example, if Host 1 communicates with Host 2, NAT (with hairpinning support) is used to route the packets, which contain Host 2's external address, back to Host 2's internal address.

Figure 18: Persistent NAT Hairpinning



In [Figure 18 on page 219](#), the following parameters are used:

- Host 1 IP address - `10.10.10.2/24`
- Host 2 IP address - `10.10.10.10/24`
- Intra-zone IP address - `10.10.10.254/24`
- Host 3 IP address - `198.51.100.2/24`

- Inter-zone IP address - 198.51.100.254/24
- Host 1 and Host 2 are in zone reht0z, and Host 3 is in reth1z zone

Table 13 on page 220 shows the binding table used in this example.

Table 13: Persistent NAT Binding Table

Original Source IP Address	Translated Source IP Address
10.10.10.2/24 to 10.10.10.11/24	192.0.2.1/32 to 192.0.2.10/32

Persistent NAT hairpinning applies only to any remote host persistent NAT type. To allow hairpinning, you must configure a security policy to allow traffic between endpoints in the same zone. Actually the two endpoints can be located in two different zones as well as long as either of the two hosts can only see the public address of the peer. NAT hairpinning behavior is not supported by target host persistent NAT and target host port persistent NAT. Only any remote host persistent NAT supports hairpinning behavior.

Example: Configuring Persistent NAT Hairpinning with Source NAT Pool with Address Shifting

IN THIS SECTION

- Requirements | 220
- Overview | 221
- Configuration | 223
- Verification | 226

This example shows how to configure persistent NAT hairpinning.

Requirements

Before you begin:

- Configure network interfaces on the device. See [Interfaces User Guide for Security Devices](#).

- Create security zones and assign interfaces to them. See *Understanding Security Zones*.

Overview

IN THIS SECTION

- [Topology](#) | **221**

Hairpinning allows packets from the private network to be translated and then looped back to the private network rather than being passed through to the public network. Hairpinning feature enables using a corresponding record in the NAT table to recognize that a packet is addressed to a host in the local network. Then it translates the destination IP address and sends the packet back to the local network (as well as in case of port mapping). This ensures that traffic between the two hosts work properly.


Topology

Hairpinning enables two endpoints (Host 1 and Host 2) on the private network to communicate even if they only use each other's external IP addresses and ports. This is explained in [Figure 19 on page 222](#).

When Host 1 sends traffic to Host 3, a NAT binding between Host 1's internal source IP address and port is associated in the NAT table with its external IP address and port. The same thing happens when Host 2 sends traffic to Host 3. In this way, when Host 1 and Host 2 want to communicate, they can identify each other's external IP addresses.

For example, if Host 1 communicates with Host 2, NAT (with hairpinning support) is used to route the packets, which contain Host 2's external address, back to Host 2's internal address.

Figure 19: Persistent NAT Hairpinning



In [Figure 19 on page 222](#), the following parameters are used:

- Host 1 IP address - 10.10.10.2/24
- Host 2 IP address - 10.10.10.10/24
- Intra-zone IP address - 10.10.10.254/24
- Host 3 IP address - 198.51.100.2/24
- Inter-zone IP address - 198.51.100.254/24
- Host 1 and Host 2 are in zone reth0z, and Host 3 is in reth1z zone

[Table 14 on page 223](#) shows the binding table used in this example.

Table 14: Persistent NAT Binding Table

Original Source IP Address	Translated Source IP Address
10.10.10.2/24 to 10.10.10.11/24	192.0.2.1/32 to 192.0.2.10/32

Configuration

IN THIS SECTION

Procedure | 223

Procedure

Step-by-Step Procedure

To configure persistent NAT hairpinning:

1. Configure interfaces.

```
[edit]
user@host# set interfaces ge-11/0/0 unit 0 family inet address 10.10.10.254/24
user@host# set interfaces ge-11/0/1 unit 0 family inet address 198.51.100.254/24
```

2. Create zones (reth0z and reth1z).

```
[edit]
user@host# set security zones security-zone reth0z host-inbound-traffic system-services all
user@host# set security zones security-zone reth0z host-inbound-traffic protocols all
user@host# set security zones security-zone reth0z interfaces ge-11/0/0.0
user@host# set security zones security-zone reth1z host-inbound-traffic system-services all
user@host# set security zones security-zone reth1z host-inbound-traffic protocols all
user@host# set security zones security-zone reth1z interfaces ge-11/0/1.0
```

3. Create policies for zones reth0z and reth1z.

```
[edit]
user@host# set security address-book global address subnet10 10.10.10.0/24
user@host# set security address-book global address subnet20 198.51.100.0/24
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 match source-
address subnet10
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 match destination-
address subnet20
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 match application
any
user@host# set security policies from-zone reth0z to-zone reth1z policy p1 then permit
user@host# set security policies default-policy deny-all
```

4. Add same zone policy to do persistent NAT hairpinning.

```
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 match source-
address subnet10
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 match destination-
address subnet10
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 match application
any
user@host# set security policies from-zone reth0z to-zone reth0z policy p2 then permit
```

5. Create a source NAT pool for Host 1 and Host 2 (src1).

```
[edit]
user@host# set security nat source pool src1 address 192.0.2.1/32 to 192.0.2.10/32
```

6. Specify the beginning of the original source IP address range for Host 1 and Host 2 (src1).

```
[edit]
user@host# set security nat source pool src1 host-address-base 10.10.10.2/24
```

7. Configure the source NAT rule set r1.

```
[edit]
user@host# set security nat source rule-set r1 from zone reth0z
```

```

user@host# set security nat source rule-set r1 to zone reth1z
user@host# set security nat source rule-set r1 to zone reth0z
user@host# set security nat source rule-set r1 rule rule1 match source-address 10.10.10.0/24
user@host# set security nat source rule-set r1 rule rule1 match destination-address
10.10.10.0/24
user@host# set security nat source rule-set r1 rule rule1 match destination-address
198.51.100.0/24
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool src1
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool persistent-nat
permit any-remote-host
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool persistent-nat
inactivity-timeout 900
user@host# set security nat source rule-set r1 rule rule1 then source-nat pool persistent-nat
max-session-number 20

```

Results

From configuration mode, enter the `show security nat` command to confirm your configuration. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```

[edit]
user@host# show security nat
source {
    pool src1 {
        address {
            192.0.2.1/32 to 192.0.2.10/32;
        }
        host-address-base 10.10.10.2/24;
    }
    rule-set r1 {
        from zone reth0z;
        to zone [ reth0z reth1z ];
        rule rule1 {
            match {
                source-address 10.10.10.0/24;
                destination-address [10.10.10.0/24 198.51.100.0/24];
            }
            then {
                source-nat {
                    pool {

```



```
user@host>show security nat source persistent-nat-table all
```

Internal	Reflective	Source		Type	Left_time/	Curr_Sess_Num/	Source
In_IP	In_Port	Ref_IP	Ref_Port	NAT Pool	Conf_time	Max_Sess_Num	NAT Rule
10.10.10.2	69	192.0.2.1	69	src1	any-remote-host	-/900	1/20 rule1

Traffic Sent Between the Hosts Creating Binding 2

Purpose

Verify traffic sent from between the hosts (Host 2 and Host 3) creating binding 2.

Action

```
sendip -d r28 -p ipv4 -iv 4 -is 10.10.10.10 -id 198.51.100.2 -p udp -us 69 -ud 69 198.51.100.2
Source-IP: 10.10.10.10
Source-port: 69
Dst-IP: 198.51.100.2
Dst-port: 69
Binding2 is below:
```

```
user@host>show security nat source persistent-nat-table all
```

Internal	Reflective	Source		Type	Left_time/	Curr_Sess_Num/	
Source							
In_IP	In_Port	Ref_IP	Ref_Port	NAT Pool	Conf_time	Max_Sess_Num	NAT Rule
10.10.10.2	69	192.0.2.1	69	src1	any-remote-host	-/900	1/20 rule1
10.10.10.10	69	192.0.2.9	69	src1	any-remote-host	-/900	1/20 rule1

Traffic Sent Between Two Hosts

Purpose

Verify the traffic sent from Host 1 to Host 2:

Action

```
user@host>show security flow session
sendip -d r28 -p ipv4 -iv 4 -is 10.10.10.2 -id 192.0.2.9 -p udp -us 69 -ud 69 192.0.2.9

Session ID: 100007628, Policy name: default-policy/2, Timeout: 52, Valid
In: 10.10.10.2/69 --> 192.0.2.9/69;udp, If: ge-0/0/0.0, Pkts: 2, Bytes: 112
Out: 10.10.10.10/69 --> 192.0.2.1/69;udp, If: ge-0/0/0.0, Pkts: 0, Bytes: 0
Total sessions: 1
```

NAT for Multicast Flows

IN THIS SECTION

- [Understanding NAT for Multicast Flows | 228](#)
- [Example: Configuring NAT for Multicast Flows | 229](#)

To implement multicast group address translation, either static NAT or destination NAT is used. With the help of NAT, source addresses in IPv4 are translated to IPv4 multicast group destination addresses.

Understanding NAT for Multicast Flows

Network Address Translation (NAT) can be used to translate source addresses in IPv4 multicast flows and to translate IPv4 multicast group destination addresses.

Either static NAT or destination NAT can be used to perform multicast group address translation. Static NAT allows connections to be originated from either side of the network, but translation is limited to one-to-one addresses or between blocks of addresses of the same size. No address pools are necessary. Use the static *configuration statement* at the [edit security nat] hierarchy level to configure static NAT rule sets for multicast traffic. Destination NAT allows connections to be initiated only for incoming network connections—for example, from the Internet to a private network. Use the destination *configuration statement* at the [edit security nat] hierarchy level to configure destination NAT pools and rule sets.

Source NAT for multicast traffic is supported only by using IP address shifting to translate the original source IP address to an IP address from a user-defined address pool. This type of translation is one-to-one, static, and without port address translation. If the original source IP address range is larger than the IP address range in the user-defined pool, untranslated packets are dropped. The mapping does not provide bidirectional mapping, which static NAT provides. Use the source configuration statement at the [edit security nat] hierarchy level to configure source NAT pools and rule sets. When you define the source NAT pool for this type of source NAT, use the host-address-base option to specify the start of the original source IP address range.

SEE ALSO

[Source NAT | 32](#)

[Static NAT | 157](#)

[Destination NAT | 125](#)

Example: Configuring NAT for Multicast Flows

IN THIS SECTION

- [Requirements | 229](#)
- [Overview | 230](#)
- [Configuration | 232](#)
- [Verification | 240](#)

This example shows how to configure a Juniper Networks device for address translation of multicast flows.

Requirements

Before you begin:

1. Configure network interfaces on the device. See the [Interfaces User Guide for Security Devices](#).
2. Create security zones and assign interfaces to them. See *Understanding Security Zones*.
3. Configure the device for multicast forwarding. See the [Multicast Overview](#).

Overview

IN THIS SECTION

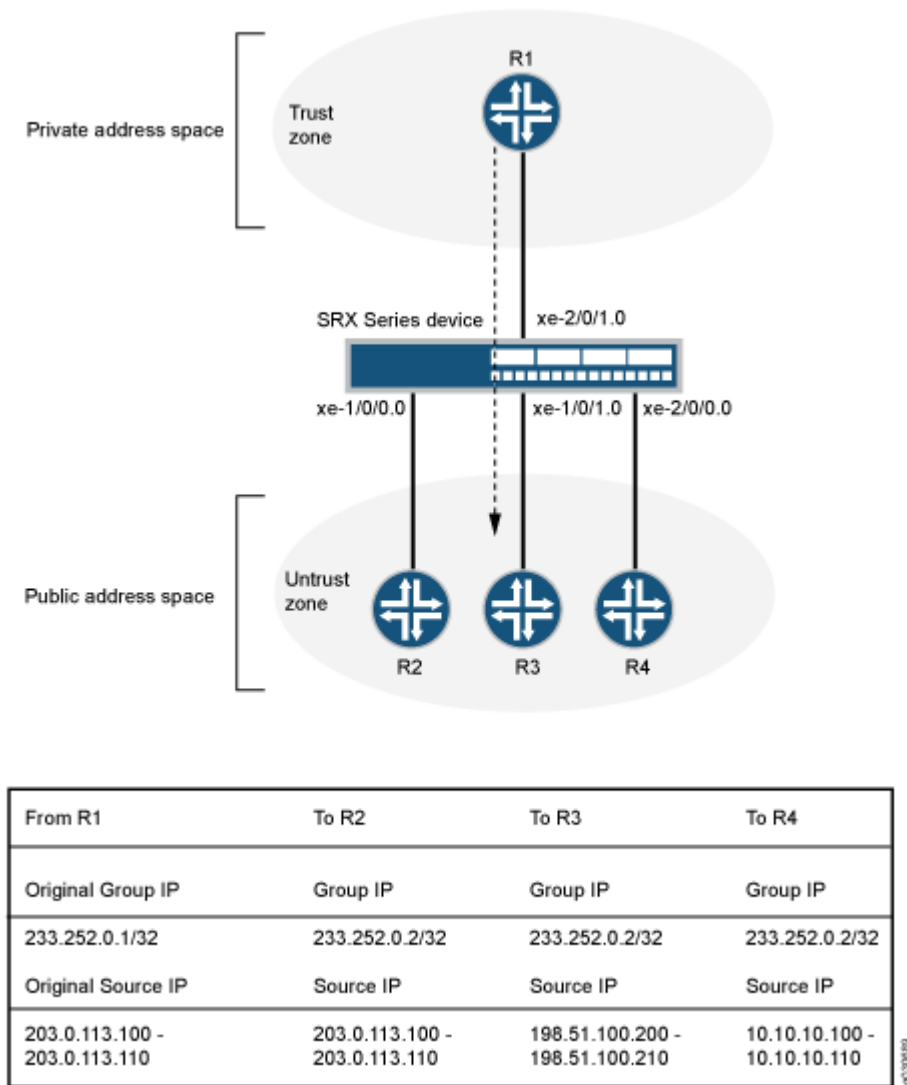
- [Topology | 232](#)

This example uses the trust security zone for the private address space and the untrust security zone for the public address space. [Figure 20 on page 231](#) depicts a typical deployment of the Juniper Networks device for multicast forwarding. The source router R1 sends multicast packets with source addresses in the range 203.0.113.100 through 203.0.113.110 and the group address 233.252.0.1/32 toward the Juniper Networks device. The source router R1 is in the private network (trust zone) upstream of the Juniper Networks device. There are several receivers in the public network (untrust zone) downstream of the device.

The Juniper Networks device translates incoming multicast packets from R1 before forwarding them out on the downstream interfaces. The following translations are applied:

- For the interface to R2, the source address is untranslated, and the group address is translated to 233.252.0.2/32.
- For the interface to R3, the source address is translated to an address in the range 198.51.100.200 through 198.51.100.210, and the group address is translated to 233.252.0.2/32.
- For the interface to R4, the source address is translated to an address in the range 10.10.10.100 through 10.10.10.110, and the group address is translated to 233.252.0.2/32.

Figure 20: NAT Translations for Multicast Flows



This example describes the following configurations:

- Destination NAT pool `dst-nat-pool` that contains the IP address 233.252.0.2/32.
- Destination NAT rule set `rs1` with rule `r1` to match packets arriving on interface `xe-2/0/1.0` with the destination IP address 233.252.0.1/32. For matching packets, the destination address is translated to the IP address in the `dst-nat-pool` pool.
- Source NAT pool `src-nat-shift-1` that contains the IP address range 198.51.100.200/32 through 198.51.100.210/32. For this pool, the beginning of the original source IP address range is 203.0.113.100/32 and is specified with the `host-address-base` option.

- Source NAT rule set `rs-shift1` with rule `r1` to match packets from the trust zone to interface `xe-1/0/1.0` with a source IP address in the `203.0.113.96/28` subnet. For matching packets that fall within the source IP address range specified by the `src-nat-shift-1` configuration, the source address is translated to the IP address in the `src-nat-shift-1` pool.
- Source NAT pool `src-nat-shift-2` that contains the IP address range `10.10.10.100/32` through `10.10.10.110/32`. For this pool, the beginning of the original source IP address range is `203.0.113.100/32` and is specified with the `host-address-base` option.
- Source NAT rule set `rs-shift2` with rule `r1` to match packets from the trust zone to interface `xe-2/0/0.0` with a source IP address in the `203.0.113.96/28` subnet. For matching packets that fall within the source IP address range specified by the `src-nat-shift-2` configuration, the source address is translated to the IP address in the `src-nat-shift-2` pool.
- Proxy ARP for the addresses `203.0.113.100` through `203.0.113.110` on interface `xe-1/0/0.0`, addresses `198.51.100.200` through `198.51.100.210` on interface `xe-1/0/1.0`, and addresses `10.10.10.100` through `10.10.10.110` on interface `xe-2/0/0.0`. This allows the Juniper Networks security device to respond to ARP requests received on the interface for those addresses.
- Security policy to permit traffic from the trust zone to the untrust zone.
- Security policy to permit traffic from the untrust zone to the translated destination IP address in the trust zone.

Topology

Configuration

IN THIS SECTION

- [Procedure | 233](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security nat source pool src-nat-shift-1 address 198.51.100.200/32 to 198.51.100.210/32
set security nat source pool src-nat-shift-1 host-address-base 203.0.113.100/32
set security nat source pool src-nat-shift-2 address 10.10.10.100/32 to 10.10.10.110/32
set security nat source pool src-nat-shift-2 host-address-base 203.0.113.100/32
set security nat source rule-set rs-shift1 from zone trust
set security nat source rule-set rs-shift1 to interface xe-1/0/1.0
set security nat source rule-set rs-shift1 rule r1 match source-address 203.0.113.96/28
set security nat source rule-set rs-shift1 rule r1 then source-nat pool src-nat-shift1
set security nat source rule-set rs-shift2 from zone trust
set security nat source rule-set rs-shift2 to interface xe-2/0/0.0
set security nat source rule-set rs-shift2 rule r2 match source-address 203.0.113.96/28
set security nat source rule-set rs-shift2 rule r2 then source-nat pool src-nat-shift2
set security nat destination pool dst-nat-pool address 233.252.0.2/32
set security nat destination rule-set rs1 from interface xe-2/0/1.0
set security nat destination rule-set rs1 rule r1 match destination-address 233.252.0.1/32
set security nat destination rule-set rs1 rule r1 then destination-nat pool dst-nat-pool
set security nat proxy-arp interface xe-1/0/0.0 address 203.0.113.100/32 to 203.0.113.110/32
set security nat proxy-arp interface xe-1/0/1.0 address 198.51.100.200/32 to 198.51.100.210/32
set security nat proxy-arp interface xe-2/0/0.0 address 10.10.10.100/32 to 10.10.10.110/32
set security policies from-zone trust to-zone untrust policy internet-access match source-address any
set security policies from-zone trust to-zone untrust policy internet-access match destination-address any
set security policies from-zone trust to-zone untrust policy internet-access match application any
set security policies from-zone trust to-zone untrust policy internet-access then permit
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access match source-address any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access match destination-address 233.252.0.1/21
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access match application any
set security policies from-zone untrust to-zone trust policy dst-nat-pool-access then permit
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure the destination and source NAT translations for multicast flows:

1. Create a destination NAT pool.

```
[edit security nat destination]
user@host# set pool dst-nat-pool address 233.252.0.2/32
```

2. Create a destination NAT rule set.

```
[edit security nat destination]
user@host# set rule-set rs1 from interface xe-2/0/1.0
```

3. Configure a rule that matches packets and translates the destination address to the address in the destination NAT pool.

```
[edit security nat destination]
user@host# set rule-set rs1 rule r1 match destination-address 233.252.0.1/32
user@host# set rule-set rs1 rule r1 then destination-nat pool dst-nat-pool
```

4. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool src-nat-shift-1 address 198.51.100.200 to 198.51.100.210
```

5. Specify the beginning of the original source IP address range.

```
[edit security nat source]
user@host# set pool src-nat-shift-1 host-address-base 203.0.113.100
```

6. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs-shift1 from zone trust
user@host# set rule-set rs-shift1 to interface xe-1/0/1.0
```

7. Configure a rule that matches packets and translates the destination address to the address in the source NAT pool.

```
[edit security nat source]
user@host# set rule-set rs-shift1 rule r1 match source-address 203.0.113.96/28
user@host# set rule-set rs-shift1 rule r1 then source-nat pool src-nat-shift1
```

8. Create a source NAT pool.

```
[edit security nat source]
user@host# set pool src-nat-shift-2 address 10.10.10.100 to 10.10.10.110
```

9. Specify the beginning of the original source IP address range.

```
[edit security nat source]
user@host# set pool src-nat-shift-2 host-address-base 203.0.113.100/32
```

10. Create a source NAT rule set.

```
[edit security nat source]
user@host# set rule-set rs-shift2 from zone trust
user@host# set rule-set rs-shift2 to interface xe-2/0/0.0
```

11. Configure a rule that matches packets and translates the destination address to the address in the source NAT pool.

```
[edit security nat source]
user@host# set rule-set rs-shift2 rule r2 match source-address 203.0.113.96/28
user@host# set rule-set rs-shift2 rule r2 then source-nat pool src-nat-shift2
```

12. Configure proxy ARP.

```
[edit security nat]
user@host# set proxy-arp interface xe-1/0/0.0 address 203.0.113.100 to 203.0.113.110
user@host# set proxy-arp interface xe-1/0/1.0 address 198.51.100.200 to 198.51.100.210
user@host# set proxy-arp interface xe-2/0/0.0 address 10.10.10.100 to 10.10.10.110
```

13. Configure a security policy that allows traffic from the trust zone to the untrust zone.

```
[edit security policies from-zone trust to-zone untrust]
user@host# set policy internet-access match source-address any destination-address any
application any
user@host# set policy internet-access then permit
```

14. Configure a security policy that allows traffic from the untrust zone to the trust zone.

```
[edit security policies from-zone untrust to-zone trust]
user@host# set policy dst-nat-pool-access match source-address any destination-address
233.252.0.1/32 application any
user@host# set policy dst-nat-pool-access then permit
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show security policies` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
source {
  pool src-nat-shift-1 {
    address {
      198.51.100.200/32 to 198.51.100.210/32;
    }
    host-address-base 203.0.113.100/32;
  }
  pool src-nat-shift-2 {
    address {
```

```

        10.10.10.100/32 to 10.10.10.110/32;
    }
    host-address-base 203.0.113.100/32;
}
rule-set trust-to-untrust {
    from zone trust;
    to zone untrust;
    rule source-nat-rule {
        match {
            source-address 0.0.0.0/0;
        }
        then {
            source-nat {
                interface;
            }
        }
    }
}
rule-set rs-shift1 {
    from zone trust;
    to interface xe-1/0/1.0;
    rule r1 {
        match {
            source-address 203.0.113.96/28;
        }
        then {
            source-nat {
                pool {
                    src-nat-shift1;
                }
            }
        }
    }
}
rule-set rs-shift2 {
    from zone trust;
    to interface xe-2/0/0.0;
    rule r2 {
        match {
            source-address 203.0.113.96/28;
        }
        then {
            source-nat {

```



```

    }
}

```

```

[edit]
user@host# show security policies
from-zone trust to-zone untrust {
  policy trust-to-untrust {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
  policy internet-access {
    match {
      source-address any;
      destination-address any;
      application any;
    }
    then {
      permit;
    }
  }
}
from-zone untrust to-zone trust {
  policy dst-nat-pool-access {
    match {
      source-address any;
      destination-address 233.252.0.1/21;
      application any;
    }
    then {
      permit;
    }
  }
}
}
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Destination NAT Pool Usage | 240](#)
- [Verifying Destination NAT Rule Usage | 240](#)
- [Verifying Source NAT Pool Usage | 241](#)
- [Verifying Source NAT Rule Usage | 241](#)
- [Verifying NAT Application to Traffic | 241](#)

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

Verifying Destination NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the destination NAT pool.

Action

From operational mode, enter the `show security nat destination pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Destination NAT Rule Usage

Purpose

Verify that there is traffic matching the destination NAT rule.

Action

From operational mode, enter the `show security nat destination rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying Source NAT Pool Usage

Purpose

Verify that there is traffic using IP addresses from the source NAT pool.

Action

From operational mode, enter the `show security nat source pool all` command. View the Translation hits field to check for traffic using IP addresses from the pool.

Verifying Source NAT Rule Usage

Purpose

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. View the Translation hits field to check for traffic that matches the rule.

Verifying NAT Application to Traffic

Purpose

Verify that NAT is being applied to the specified traffic.

Action

From operational mode, enter the `show security flow session` command.

IPv6 NAT

IN THIS SECTION

- [IPv6 NAT Overview | 242](#)
- [IPv6 NAT PT Overview | 245](#)
- [IPv6 NAT-PT Communication Overview | 246](#)
- [Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping | 247](#)
- [Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping | 252](#)
- [Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping | 257](#)
- [Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping | 263](#)

IPv6 NAT helps to translate IPv4 addresses to IPv6 addresses of network devices. IPv6 NAT also helps to translate the address between IPv6 hosts. IPv6 NAT supports source NAT, destination NAT, and static NAT.

IPv6 NAT Overview

IN THIS SECTION

- [Source NAT Translations Supported by IPv6 NAT | 243](#)
- [Destination NAT Mappings Supported by IPv6 NAT | 243](#)
- [Static NAT Mappings Supported by IPv6 NAT | 244](#)

IPv6 has a vastly larger address space than the impending exhausted IPv4 address space. IPv4 has been extended using techniques such as Network Address Translation (NAT), which allows for ranges of private addresses to be represented by a single public address, and temporary address assignment.

There are a lot of technologies to provide the transition mechanism for the legacy IPv4 host to keep the connection to the Internet. IPv6 NAT provides address translation between IPv4 and IPv6 addressed network devices. It also provides address translation between IPv6 hosts. NAT between IPv6 hosts is done in a similar manner and for similar purposes as IPv4 NAT.

IPv6 NAT in Junos OS provides the following NAT types:

- Source NAT
- Destination NAT
- Static NAT

Source NAT Translations Supported by IPv6 NAT

Source NAT is the translation of the source IP address of a packet leaving the Juniper Networks device. Source NAT is used to allow hosts with private IP addresses to access a public network.

IPv6 NAT in Junos OS supports the following source NAT translations:

- Translation of one IPv6 subnet to another IPv6 subnet without port address translation
- Translation of IPv4 addresses to IPv6 prefix + IPv4 addresses
- Translation of IPv6 hosts to IPv6 hosts with or without port address translation
- Translation of IPv6 hosts to IPv4 hosts with or without port address translation
- Translation of IPv4 hosts to IPv6 hosts with or without port address translation

Destination NAT Mappings Supported by IPv6 NAT

Destination NAT is the translation of the destination IP address of a packet entering the Juniper Networks device. Destination NAT is used to redirect traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

IPv6 NAT in Junos OS supports the following destination NAT translations:

- Prefix translation between IPv4 and IPv6 prefix
- Mapping of one IPv6 subnet to another IPv6 subnet
- Mapping of one IPv6 subnet to an IPv6 host
- Mapping of one IPv6 subnet to one IPv4 subnet
- Mapping of one IPv4 subnet to one IPv6 subnet

- Mapping of one IPv6 host (and optional port number) to one special IPv6 host (and optional port number)
- Mapping of one IPv6 host (and optional port number) to one special IPv4 host (and optional port number)
- Mapping of one IPv4 host (and optional port number) to one special IPv6 host (and optional port number)

Static NAT Mappings Supported by IPv6 NAT

Static NAT defines a one-to-one mapping from one IP subnet to another IP subnet. The mapping includes destination IP address translation in one direction and source IP address translation in the reverse direction. From the NAT device, the original destination address is the virtual host IP address while the mapped-to address is the real host IP address.

IPv6 NAT in Junos OS supports the following static NAT translations:

- Translation of one IPv6 subnet to another IPv6 subnet
- Translation of one IPv6 host to another IPv6 host
- Translation of one IPv4 address a.b.c.d to IPv6 address Prefix::a.b.c.d
- Translation of IPv4 hosts to IPv6 hosts

See ["Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping"](#) on page 247.

- Translation of IPv6 hosts to IPv4 hosts

See ["Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping"](#) on page 257.

- Mapping of one IPv6 prefix to one IPv4 prefix

See ["Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping"](#) on page 263.

- Mapping of one IPv4 prefix to one IPv6 prefix

See ["Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping"](#) on page 252.

IPv6 NAT PT Overview

Starting in Junos OS Release 20.2R1 you can run IPv6 NAT-PT Next Gen Services on MX240, MX480, and MX960 routers.

IPv6 Network Address Translation-Protocol Translation (NAT-PT) provides address allocation and protocol translation between IPv4 and IPv6 addressed network devices. The translation process is based on the Stateless IP/ICMP Translation (SIIT) method; however, the state and the context of each communication are retained during the session lifetime. IPv6 NAT-PT supports Internet Control Message Protocol (ICMP), TCP, and UDP packets.

IPv6 NAT-PT supports the following types of NAT-PT:

- **Traditional NAT-PT**—In traditional NAT-PT, the sessions are unidirectional and outbound from the IPv6 network. Traditional NAT-PT allows hosts within an IPv6 network to access hosts in an IPv4 network. There are two variations to traditional NAT-PT: basic NAT-PT and NAPT-PT.

In basic NAT-PT, a block of IPv4 addresses at an IPv4 interface is set aside for translating addresses as IPv6 hosts as they initiate sessions to the IPv4 hosts. The basic NAT-PT translates the source IP address and related fields such as IP, TCP, UDP, and ICMP header checksums for packets outbound from the IPv6 domain. For inbound packets, it translates the destination IP address and the checksums.

Network Address Port Translation-Protocol Translation (NAPT-PT) can be combined with basic NAT-PT so that a pool of external addresses is used in conjunction with port translation. NAPT-PT allows a set of IPv6 hosts to share a single IPv4 address. NAPT-PT translates the source IP address, source transport identifier, and related fields such as IP, TCP, UDP, and ICMP header checksums, for packets outbound from the IPv6 network. The transport identifier can be a TCP/UDP port or an ICMP query ID. For inbound packets, it translates the destination IP address, destination transport identifier, and the IP and the transport header checksums.

- **Bidirectional NAT-PT**—In bidirectional NAT-PT, sessions can be initiated from hosts in the IPv4 network as well as the IPv6 network. IPv6 network addresses are bound to IPv4 addresses, either statically or dynamically as connections are established in either direction. The static configuration is similar to static NAT translation. Hosts in IPv4 realm access hosts in the IPv6 realm using DNS for address resolution. A DNS ALG must be employed in conjunction with bidirectional NAT-PT to facilitate name-to-address mapping. Specifically, the DNS ALG must be capable of translating IPv6 addresses in DNS queries and responses into their IPv4 address bindings, and vice versa, as DNS packets traverse between IPv6 and IPv4 realms.

NOTE: The devices partially support the bidirectional NAT-PT specification. It supports flow of bidirectional traffic assuming that there are other ways to convey the mapping between

the IPv6 address and the dynamically allocated IPv4 address. For example, a local DNS can be configured with the mapped entries for IPv4 nodes to identify the addresses.

NAT-PT Operation—The devices support the traditional NAT-PT and allow static mapping for the user to communicate from IPv4 to IPv6. The user needs to statically configure the DNS server with an IPv4 address for the hostname and then create a static NAT on the device for the IPv6-only node to communicate from an IPv4-only node to an IPv6-only node based on the DNS.

SEE ALSO

[NAT46 Next Gen Services Configuration Examples](#)

IPv6 NAT-PT Communication Overview

NAT-PT communication with static mapping— Network Address Translation-Protocol Translation (NAT-PT) can be done in two directions, from IPv6 to IPv4 and vice versa. For each direction, static NAT is used to map the destination host to a local address and a source address NAT is used to translate the source address. There are two types of static NAT and source NAT mapping: one-to-one mapping and prefix-based mapping.

NAT-PT communication with DNS ALG—A DNS-based mechanism dynamically maps IPv6 addresses to IPv4-only servers. NAT-PT uses the DNS ALG to transparently do the translations. For example, a company using an internal IPv6 network needs to be able to communicate with external IPv4 servers that do not yet have IPv6 addresses.

To support the dynamic address binding, a DNS should be used for name resolution. The IPv4 host looks up the name of the IPv6 node in its local configured IPv4 DNS server, which then passes the query to the IPv6 DNS server through a device using NAT-PT.

The DNS ALG in NAT device :

- Translates the IPv6 address resolution back to IPv4 address resolution.
- Allocates an IPv6 address for the mapping.
- Stores a mapping of the allocated IPv4 address to the IPv6 address returned in the IPv6 address resolution so that the session can be established from any-IPv4 hosts to the IPv6 host.

SEE ALSO[IPv6 NAT PT Overview | 245](#)

Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Default Destination Address Prefix Static Mapping

IN THIS SECTION

- [Requirements | 247](#)
- [Overview | 247](#)
- [Configuration | 248](#)
- [Verification | 251](#)

This example shows how to configure an IPv4-initiated connection to an IPv6 node using default destination address prefix static mapping.

Requirements

Before you begin, configure interfaces and assign them to security zones.

Overview

The following example describes how to configure an IPv4-initiated connection to an IPv6 node that has a static mapping 126-based IPv6 address defined on its interface and static mapping /126 set up on the device. This example assumes that the IPv6 addresses to be mapped to IPv4 addresses make the IPv4 addresses part of the IPv6 address space.

Configuring an IPv4-initiated connection to an IPv6 node is useful when the devices on the IPv4 network must be interconnected to the devices on the IPv6 network and during migration of an IPv4 network to an IPv6 network. The mapping can be used for DNS ALG for reverse lookup of IPv4 addresses from IPv6 addresses, for the traffic initiated from the IPv6 network. This process also provides connectivity for sessions initiated from IPv4 nodes with IPv6 nodes on the other side of the NAT/PT device.

Configuration

IN THIS SECTION

- [Procedure](#) | 248

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/0.0
set security nat static rule-set test_rs rule test_rule match destination-address 10.1.1.0/30
set security nat static rule-set test_rs rule test_rule then static-nat prefix 2001:db8::/126
set security nat source pool myipv6_prefix address 2001:db8::/126
set security nat source rule-set myipv6_rs from interface ge-0/0/0.0
set security nat source rule-set myipv6_rs to interface ge-0/0/1.0
set security nat source rule-set myipv6_rs rule ipv6_rule match source-address 10.1.1.45/30
set security nat source rule-set myipv6_rs rule ipv6_rule match destination-address
2001:db8::/96
set security nat source rule-set myipv6_rs rule ipv6_rule then source-nat pool myipv6_prefix
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure an IPv4-initiated connection to an IPv6 node using static destination address one-to-one mapping:

1. Configure the static NAT rule set for an interface.

```
[edit security nat static]
user@host# set rule-set test_rs from interface ge-0/0/0.0
```

2. Define the rule to match the destination address prefix.

NOTE: The destination address number in the match rule must be a number equal to the static-nat prefix range.

There is no limitation on the source address number in the match rule.

```
[edit security nat static rule-set test_rs]
user@host# set rule test_rule match destination-address 10.1.1.0/30
```

3. Define the static NAT prefix for the device.

```
[edit security nat static rule-set test_rs]
user@host# set rule test_rule then static-nat prefix 2001:db8::/126
```

4. Configure the source NAT pool with an IPv6 address prefix.

```
[edit security nat source]
user@host# set pool myipv6_prefix address 2001:db8::/126
```

5. Configure the source NAT rule set for the interface.

```
[edit security nat source]
user@host# set rule-set myipv6_rs from interface ge-0/0/0.0
user@host# set rule-set myipv6_rs to interface ge-0/0/1.0
```

6. Configure the IPv6 source NAT source address.

NOTE: The source address number in the match rule must be an address number equal to the source pool range. For example, $2^{(32 - 30)} = 2^{(128 - 126)} = >$.

There is no limitation on the destination address number in the match rule.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match source-address 10.1.1.45/30
```

7. Configure the IPv6 source NAT destination address.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match destination-address 2001:db8::/96
```

8. Define the configured source NAT IPv6 pool in the rule.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule then source-nat pool myipv6_prefix
```

Results

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

```
source {
  pool myipv6_prefix {
    address {
      2001:db8::/126;
    }
  }
  rule-set myipv6_rs {
    from interface ge-0/0/0.0;
    to interface ge-0/0/1.0;
    rule ipv6_rule {
      match {
        source-address 10.1.1.45/30;
        destination-address 2001:db8::/96;
      }
      then {
        source-nat {
          pool {
            myipv6_prefix;
          }
        }
      }
    }
  }
}
```

```
static {  
  rule-set test_rs {  
    from interface ge-0/0/0.0;  
    rule test_rule {  
      match {  
        destination-address 10.1.1.0/30;  
      }  
      then {  
        static-nat {  
          prefix {  
            2001:db8::/126;  
          }  
        }  
      }  
    }  
  }  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That Static NAT Is Configured | 251](#)
- [Verifying That Source NAT Is Configured | 252](#)

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

Verifying That Static NAT Is Configured

Purpose

Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action

From operational mode, enter the `show security nat static` command.

Verifying That Source NAT Is Configured

Purpose

Verify whether source NAT is configured.

Action

From operational mode, enter the `show security nat source` command.

Example: Configuring an IPv4-Initiated Connection to an IPv6 Node Using Static Destination Address One-to-One Mapping

IN THIS SECTION

- [Requirements | 252](#)
- [Overview | 252](#)
- [Configuration | 253](#)
- [Verification | 256](#)

This example shows how to configure an IPv4-initiated connection to an IPv6 node using static destination address one-to-one mapping.

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones.

Overview

The following example describes how to configure an IPv4 node to communicate with an IPv6 node using one-to-one static NAT on the device.

The communication of an IPv4 node with an IPv6 node is useful for IPv4 hosts accessing an IPv6 server, for new servers that support IPv6 only and that need to be connected to the IPv6 network, and for migrating of old hosts to the new server when most of the machines have already moved to IPv6. For example, you can use this feature to connect an IPv4-only node to an IPv6-only printer. This mapping

can also be used for DNS ALG for reverse lookup of IPv4 addresses from IPv6 addresses for traffic that is initiated from the IPv6 network.

In this example, the source IPv4 address matching the prefix 10.10.10.1/30 is added with the IPv6 prefix 2001:db8::/96 to form the translated source IPv6 address and the destination IPv4 address 10.1.1.25/32 is translated to IPv6 address 2001:db8::25/128.

Configuration

IN THIS SECTION

- [Procedure | 253](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address 10.1.1.25/32
set security nat static rule-set test_rs rule test_rule then static-nat prefix 2001:db8::25/128
set security nat source pool myipv6_prefix address 2001:db8::/96
set security nat source rule-set myipv6_rs from interface ge-0/0/1
set security nat source rule-set myipv6_rs to interface ge-0/0/2
set security nat source rule-set myipv6_rs rule ipv6_rule match source-address 10.10.10.1/30
set security nat source rule-set myipv6_rs rule ipv6_rule match destination-address 2001:db8::25
set security nat source rule-set myipv6_rs rule ipv6_rule then source-nat pool myipv6_prefix
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an IPv4-initiated connection to an IPv6 node using static destination address one-to-one mapping:

1. Configure the static NAT rule set for an interface.

```
[edit security nat static]
user@host# set rule-set test_rs from interface ge-0/0/1
```

2. Define the rule and the destination address.

```
[edit security nat static rule-set test_rs]
user@host# set rule test_rule match destination-address 10.1.1.25/32
```

3. Define the static NAT prefix.

```
[edit security nat static rule-set test_rs]
user@host# set rule test_rule then static-nat prefix 2001:db8::25/128
```

4. Configure a source NAT pool with an IPv6 prefix address.

```
[edit security]
user@host# set nat source pool myipv6_prefix address 2001:db8::/96
```

5. Configure the source NAT rule set.

```
[edit security nat source]
user@host# set rule-set myipv6_rs from interface ge-0/0/1
user@host# set rule-set myipv6_rs to interface ge-0/0/2
```

6. Configure the source NAT source address.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match source-address 10.10.10.1/30
```

7. Configure the source NAT destination address.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule match destination-address 2001:db8::25
```

8. Define a configured source NAT IPv6 pool in the rule.

```
[edit security nat source rule-set myipv6_rs]
user@host# set rule ipv6_rule then source-nat pool myipv6_prefix
```

Results

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

```
[edit]
user@host# show security nat
source {
    pool myipv6_prefix {
        address {
            2001:db8::/96;
        }
    }
    rule-set myipv6_rs {
        from interface ge-0/0/1.0;
        to interface ge-0/0/2.0;
        rule ipv6_rule {
            match {
                source-address 10.10.10.1/30;
                destination-address 2001:db8::25;
            }
            then {
                source-nat {
                    pool {
                        myipv6_prefix;
                    }
                }
            }
        }
    }
}
static {
    rule-set test_rs {
        from interface ge-0/0/1.0;
        rule test_rule {
```

```
        match {
            destination-address 10.1.1.25/32;
        }
        then {
            static-nat prefix 2001:db8::25/128;
        }
    }
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That Static NAT Is Configured | 256](#)
- [Verifying That Source NAT Is Configured | 256](#)

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

Verifying That Static NAT Is Configured

Purpose

Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action

From operational mode, enter the `show security nat static` command.

Verifying That Source NAT Is Configured

Purpose

Verify whether source NAT is configured.

Action

From operational mode, enter the `show security nat source` command.

Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Default Destination Address Prefix Static Mapping

IN THIS SECTION

- [Requirements | 257](#)
- [Overview | 257](#)
- [Configuration | 258](#)
- [Verification | 261](#)

This example shows how to configure an IPv6-initiated connection to an IPv4 node using default destination address prefix static mapping. This example does not show how to configure the NAT translation for the reverse direction.

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones.

Overview

The following example describes the communication of an IPv6 node with an IPv4 node that has prefix-based static NAT defined on the device. The static NAT assumes that the IPv4 network is a special IPv6 network (that is, an IPv4-mapped IPv6 network), and hides the entire IPv4 network behind an IPv6 prefix.

The communication of an IPv6 node with an IPv4 node is useful when IPv6 is used in the network and must be connected to the IPv4 network, or when both IPv4 and IPv6 are used in the network and a mechanism is required to interconnect the two networks during migration. This also provides connectivity for sessions initiated from IPv6 nodes with IPv4 nodes on the other side of the NAT/PT device.

Configuration

IN THIS SECTION

- [Procedure](#) | 258

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule-set test_rs from interface ge-0/0/1
set security nat static rule-set test_rs rule test_rule match destination-address 2001:db8::1/96
set security nat static rule-set test_rs rule test_rule then static-nat inet
set security nat source pool myipv4 address 203.0.113.2 to 203.0.113.5
set security nat source rule-set myipv4_rs from interface ge-0/0/1
set security nat source rule-set myipv4_rs to interface ge-0/0/2
set security nat source rule-set myipv4_rs rule ipv4_rule match destination-address 10.1.1.15/30
set security nat source rule-set myipv4_rs rule ipv4_rule match source-address 2001:db8::2/96
set security nat source rule-set myipv4_rs rule ipv4_rule then source-nat pool myipv4
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an IPv6-initiated connection to an IPv4 node using default destination address prefix static mapping:

1. Configure the static NAT for an interface.

```
[edit security nat static]
user@host# set rule test_rs from interface ge-0/0/1
```

2. Define the rule and destination address with the prefix for the static NAT translation defined on the device.

```
[edit security nat static rule-set test_rs]  
user@host# set rule test_rule match destination-address 2001:db8::1/96
```

3. Define the static NAT as inet to translate to an IPv4 address.

```
[edit security nat static rule-set test_rs]  
user@host# set rule test_rule then static-nat inet
```

4. Configure the IPv4 source NAT pool address.

```
[edit security nat source]  
user@host# set pool myipv4 address 203.0.113.2 to 203.0.113.5
```

5. Configure the source NAT rule set.

```
[edit security nat source ]  
user@host# set rule-set myipv4_rs from interface ge-0/0/1  
user@host# set rule-set myipv4_rs to interface ge-0/0/2
```

6. Configure the IPv4 source NAT destination address.

```
[edit security nat source rule-set myipv4_rs]  
user@host# set rule ipv4_rule match destination-address 10.1.1.15/30
```

7. Define the source address with the prefix for the source NAT defined on the device.

```
[edit security nat source rule-set myipv4_rs]  
user@host# set rule ipv4_rule match source-address 2001:db8::2/96
```

8. Define a configured source NAT IPv4 pool in the rule.

```
[edit security nat source rule-set myipv4_rs]
user@host# set rule ipv4_rule then source-nat pool myipv4
```

Results

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

```
[edit]
user@host# show security nat
source {
    pool myipv4 {
        address {
            203.0.113.2/32 to 203.0.113.5/32;
        }
    }
}
rule-set myipv4_rs {
    from interface ge-0/0/1.0;
    to interface ge-0/0/2.0;
    rule ipv4_rule {
        match {
            source-address 2001:db8::/96;
            destination-address 10.1.1.15/30;
        }
        then {
            source-nat {
                pool {
                    myipv4;
                }
            }
        }
    }
}
static {
    rule-set test_rs {
        from interface ge-0/0/1.0;
        rule test_rule {
```

```

        match {
            destination-address 2001:db8::1/96;
        }
        then {
            static-nat inet;
        }
    }
}
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That Static NAT Is Configured | 261](#)
- [Verifying That Source NAT Is Configured | 262](#)

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

Verifying That Static NAT Is Configured

Purpose

Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action

From operational mode, enter the `show security nat static rule` command.

```

user@host> show security nat static rule test_rule
Static NAT rule: test_rule          Rule-set: test_rs
Rule-Id                           : 2
Rule position                      : 2
From interface                    : ge-0/0/1.0
Destination addresses             : 2001:db8::1
Host addresses                    : 0.0.0.0

```



```

Netmask                : 96
Host routing-instance   : N/A
Translation hits        : 0
  Successful sessions    : 0
  Failed sessions       : 0
Number of sessions      : 0

```

Verifying That Source NAT Is Configured

Purpose

Verify whether source NAT is configured.

Action

From operational mode, enter the `show security nat source rule` command.

```

user@host> show security nat source rule ipv4_rule
source NAT rule: ipv4_rule          Rule-set: myipv4_rs
Rule-Id                            : 2
Rule position                       : 2
From interface                     : ge-0/0/1.0
To interface                       : ge-0/0/2.0
Match
  Source addresses                  : 2001:db8:: - 2001:db8::ffff:ffff
  Destination addresses             : 10.1.1.15 - 10.1.1.15
Action                             : myipv4
  Persistent NAT type               : N/A
  Persistent NAT mapping type       : address-port-mapping
  Inactivity timeout                : 0
  Max session number                : 0
Translation hits                   : 0
  Successful sessions               : 0
  Failed sessions                   : 0
Number of sessions                  : 0

```

From operational mode, enter the `show security nat source pool` command.

```

user@host> show security nat source pool myipv4
Pool name                         : myipv4

```

```

Pool id          : 5
Routing instance : default
Host address base : 0.0.0.0
Port             : [1024, 63487]
Twin port        : [63488, 65535]
Port overloading  : 1
Address assignment : no-paired
Total addresses   : 4
Translation hits   : 0
Address range     Single Ports  Twin Ports
203.0.113.2 - 203.0.113.5      0          0

```

Example: Configuring an IPv6-Initiated Connection to an IPv4 Node Using Static Destination Address One-to-One Mapping

IN THIS SECTION

- [Requirements | 263](#)
- [Overview | 263](#)
- [Configuration | 264](#)
- [Verification | 267](#)

This example shows how to configure an IPv6-initiated connection to an IPv4 node using static destination address one-to-one mapping.

Requirements

Before you begin, configure the interfaces and assign the interfaces to security zones.

Overview

The following example describes the communication of an IPv6 node with an IPv4 node that has a one-to-one static NAT address defined on the device. The communication of an IPv6 node with an IPv4 node allows IPv6 hosts to access an IPv4 server when neither of the devices has a dual stack and must depend on the NAT/PT device to communicate. This enables some IPv4 legacy server applications to work even after the network has migrated to IPv6.

Configuration

IN THIS SECTION

- [Procedure](#) | 264

Procedure

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, copy and paste the commands into the CLI at the `[edit]` hierarchy level, and then enter `commit` from configuration mode.

```
set security nat static rule test_rs from interface ge-0/0/1
set security nat static rule test_rs rule test_rule match destination-address 2001:db8::15/128
set security nat static rule test_rs rule test_rule then static-nat prefix 10.2.2.15/32
set security nat source pool myipv4 address 203.0.113.2 to 203.0.113.3
set security nat source rule myipv4_rs from interface ge-0/0/1
set security nat source rule myipv4_rs to interface ge-0/0/2
set security nat source rule myipv4_rs rule ipv4_rule match source-address 2001:db8::/96
set security nat source rule myipv4_rs rule ipv4_rule match destination-address 10.2.2.15
set security nat source rule myipv4_rs rule ipv4_rule then source-nat pool myipv4
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode*.

To configure an IPv6-initiated connection to an IPv4 node using static destination address one-to-one mapping:

1. Configure the static NAT rule set for an interface.

```
[edit security nat static]
user@host# set rule-set test_rs from interface ge-0/0/1
```

2. Define a rule to match the destination address.

```
[edit security nat static rule-set test_rs]
user@host# set rule test_rule match destination-address 2001:db8::15/128
```

3. Define the static NAT prefix to the rule.

```
[edit security nat static rule-set test_rs]
user@host# set rule test_rule then static-nat prefix 10.2.2.15/32
```

4. Configure a source NAT pool with an IPv4 addresses.

```
[edit security nat]
user@host# set source pool myipv4 address 203.0.113.2 203.0.113.3
```

5. Configure the IPv4 address for the interface.

```
[edit security nat source ]
user@host# set rule-set myipv4_rs from interface ge-0/0/1
```

6. Configure the source address to the IPv4 source NAT address.

```
[edit security nat source rule-set myipv4_rs]
user@host# set rule ipv4_rule match source-address 2001:db8::/96
```

7. Configure the destination address to IPv4 source NAT address.

```
[edit security nat source rule-set myipv4_rs]
user@host# set rule ipv4_rule match destination-address 10.2.2.15
```

8. Define the configured source NAT IPv4 pool in the rule.

```
[edit security nat source rule-set myipv4_rs]
user@host# set rule ipv4_rule then source-nat pool myipv4
```

Results

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

```
[edit]
user@host# show security nat
source {
    pool myipv4 {
        address {
            203.0.113.2/32 to 203.0.113.3/32;
        }
    }
}
rule-set myipv4_rs {
    from interface ge-0/0/1.0;
    to interface ge-0/0/2.0;
    rule ipv4_rule {
        match {
            source-address 2001:db8::/96;
            destination-address 10.2.2.15/32;
        }
        then {
            source-nat {
                pool {
                    myipv4;
                }
            }
        }
    }
}
static {
    rule-set test_rs {
        from interface ge-0/0/1.0;
        rule test_rule {
            match {
                destination-address 2001:db8::15/128;
            }
            then {
                static-nat prefix 10.2.2.15/32;
            }
        }
    }
}
```

```
}  
}  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying That Static NAT Is Configured | 267](#)
- [Verifying That Source NAT Is Configured | 267](#)

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

Verifying That Static NAT Is Configured

Purpose

Verify whether static NAT is configured with an interface, a destination address, and a prefix.

Action

From operational mode, enter the `show security nat static` command.

Verifying That Source NAT Is Configured

Purpose

Verify whether source NAT is configured.

Action

From operational mode, enter the `show security nat source` command.

Release History Table

Release	Description
20.2R1	Starting in Junos OS Release 20.2R1 you can run IPv6 NAT-PT Next Gen Services on MX240, MX480, and MX960 routers.

RELATED DOCUMENTATION

[Source NAT | 32](#)

[Destination NAT | 125](#)

[Static NAT | 157](#)

IPv6 Dual-Stack Lite

IN THIS SECTION

- [Understanding IPv6 Dual-Stack Lite | 268](#)
- [Example: Configuring IPv6 Dual-Stack Lite | 271](#)

IPv6 Dual-Stack Lite (DS-Lite) is a technology to help Internet service providers to migrate to an IPv6 access network without changing end-user software. IPv4 users continue to access IPv4 internet content with minimum disruption to their home networks while enabling IPv6 users to access IPv6 content.

Understanding IPv6 Dual-Stack Lite

IPv6 dual-stack lite (DS-Lite) is a technology that enables Internet service providers to move to an IPv6 network while simultaneously handling IPv4 address depletion.


IPv4 addresses are becoming depleted; therefore, broadband service providers (DSL, cable, and mobile) need new addresses to support new users. Providing IPv6 addresses alone is often not workable

because most of the systems that make up the public Internet are still enabled and support only IPv4, and many users' systems do not yet fully support IPv6.

DS-Lite allows service providers to migrate to an IPv6 access network without changing end-user software. The device that accesses the Internet remains the same, thus allowing IPv4 users to continue accessing IPv4 internet content with minimum disruption to their home networks, while enabling IPv6 users to access IPv6 content.

Figure 21 on page 269 illustrates the DS-Lite architecture which uses IPv6-only links between the provider and the user while maintaining the IPv4 (or dual-stack) hosts in the user network.

Figure 21: DS-Lite NAT (IPv4-in-IPv6)



The DS-Lite deployment model consists of the following components:

- Software initiator for the DS-Lite home router--Encapsulates the IPv4 packet and transmits it across an IPv6 tunnel.
- Software concentrator for DS-Lite carrier-grade Network Address Translation (NAT)--Decapsulates the IPv4-in-IPv6 packet and also performs IPv4-IPv4 NAT translations.

When a user's device sends an IPv4 packet to an external destination, DS-Lite encapsulates the IPv4 packet in an IPv6 packet for transport into the provider network. These IPv4-in-IPv6 tunnels are called *softwires*. Tunneling IPv4 over IPv6 is simpler than translation and eliminates performance and redundancy concerns.

The softwires terminate in a software concentrator at some point in the service provider network, which decapsulates the IPv4 packets and sends them through a carrier-grade Network Address Translation (NAT) device. There, the packets undergo source NAT processing to hide the original source address.

IPv6 packets originated by hosts in the subscriber's home network are transported natively over the access network.

The DS-Lite carrier-grade NAT translates IPv4-to-IPv4 addresses to multiple subscribers through a single global IPv4 address. Overlapping address spaces used by subscribers are disambiguated through the identification of tunnel endpoints. One concentrator can be the endpoint of multiple softwires.

The IPv4 packets originated by the end hosts have private (and possibly overlapping) IP addresses. Therefore, NAT must be applied to these packets. If end hosts have overlapping addresses, *Network Address Port Translation* (NAPT) is needed.

Using NAPT, the system adds the source address of the encapsulating IPv6 packet in the subscriber network to the inside IPv4 source address and port. Because each user's IPv6 address is unique, the combination of the IPv6 source address with the IPv4 source address and port creates an unambiguous mapping.

The system takes the following actions when it receives a responding IPv4 packet from outside the subscriber network:

- Encapsulates the IPv4 packet in an IPv6 packet using the mapped IPv6 address as the IPv6 destination address.
- Forwards the packet to the user.

[Table 15 on page 270](#) lists the maximum number of software initiators and software concentrators per device. Platform support depends on the Junos OS release in your installation.

Table 15: Software Initiator and Software Concentrator Capacity

Description	SRX650	SRX1500	SRX3400 SRX3600	SRX4100 SRX4200	SRX4600	SRX5400 SRX5600 SRX5800
Maximum software initiators connected per device	50,000	300	100,000	200,000	200,000	100,000
Maximum software concentrator numbers per device	32	32	32	32	32	32

NOTE: The most recent IETF draft documentation for DS-Lite uses new terminology:

- The term software initiator has been replaced by B4.
- The term software concentrator has been replaced by AFTR.

Junos OS documentation generally uses the original terms when discussing configuration in order to be consistent with the CLI statements used to configure DS-Lite.

For more information, see the following documents:

- draft-ietf-softwire-dual-stack-lite-06, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*, August 2010.
- RFC 2473, *Generic Packet Tunneling in IPv6 Specification*, December 1998.
- RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, August 1999.
- RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP, BCP 127*, January 2007.
- RFC 4925, *Softwire Problem Statement*, July 2007.
- RFC 5382, *NAT Behavioral Requirements for TCP, BCP 142*, October 2008.
- RFC 5508, *NAT Behavioral Requirements for ICMP, BCP 148*, April 2009.
- <http://www.potaroo.net/tools/ipv4/index.html>
- <http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml>

Example: Configuring IPv6 Dual-Stack Lite

IN THIS SECTION

- Requirements | 272
- Overview | 272
- Configuration | 272
- Verification | 273

When an ISP begins to allocate IPv6 addresses and IPv6-capable equipment to new subscriber homes, dual-stack lite (DS-Lite) provides a method for the private IPv4 addresses behind the IPv6 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 at the customer edge to encapsulate IPv4 packets into IPv6 packets with minimum disruption to their home network, while enabling IPv6 customers to access IPv6 content. The softwire concentrator decapsulates the IPv4-in-IPv6 packets and also performs IPv4-IPv4 NAT translations.

This example shows you how to configure a softwire concentrator for IPv4-in-IPv6 addresses.

Requirements

Before you begin:

- Review the overview section on DS-Lite. See ["Understanding IPv6 Dual-Stack Lite" on page 268](#).
- Review how ICMPv6 packets are handled by the SRX Series Firewalls. See [Understanding How SRX Series Devices Handle ICMPv6 Packets](#).

Overview

This configuration example shows how to configure a software concentrator, the software name, the concentrator address, and the software type.

NOTE: The software concentrator IPv6 address can match an IPv6 address configured on a physical interface or an IPv6 address configured on a loopback interface.

Configuration

IN THIS SECTION

- [Procedure | 272](#)

Procedure

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, and then enter `commit` from configuration mode.

```
set security softwires software-name my_sc1 software-concentrator 2001:db8::1 software-type IPv4-in-IPv6
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For instructions on how to do that, see *Using the CLI Editor in Configuration Mode* in the [CLI User Guide](#).

To configure a DS-Lite software concentrator to convert IPv4 packets into IPv6 packets:

1. Assign a name for the software concentrator.

```
[edit security]
user@host# edit softwires software-name my_sc1
```

2. Specify the address of the software concentrator.

```
[edit security softwires software-name my_sc1]
user@host# set software-concentrator 2001:db8::1
```

3. Specify the software type for IPv4 to IPv6.

```
[edit security softwires software-name my_sc1 software-concentrator 2001:db8::1]
user@host# set software-type IPv4-in-IPv6
```

Results

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

```
[edit security softwires software-name my_sc1]
user@host# show
software-concentrator 2001:db8::1;
software-type ipv4-in-ipv6;
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

From operational mode, enter the `show security softwires` command. If a software is not connected, the operational output looks like the following sample:

```
user@host# show security softwires
Software Name          SC Address          Status    Number of SI connected
```

my-sc1	2001:db8::1	Active	0
--------	-------------	--------	---

If a softwire is connected, the operational output looks like the following sample:

```
user@host# show security softwires
```

Softwire Name	SC Address	Status	Number of SI connected
my-sc1	2001:db8::1	Connected	1

RELATED DOCUMENTATION

[Understanding IPv6 Address Space, Addressing, Address Format, and Address Types](#)

[Understanding How SRX Series Devices Handle ICMPv6 Packets](#)

[About the IPv6 Basic Packet Header](#)

NAT for VRF Routing Instance

IN THIS SECTION

- [NAT Overview | 275](#)
- [Example: Configuring Source NAT to convert the private IP address of a VRF instance to the private IP address of another VRF instance | 275](#)
- [Example: Configuring Destination NAT to Convert Public IP Address to VRF's Single Private IP Address of a VRF instance | 283](#)
- [Example: Configuring Static NAT to Convert the Private IP Address of a VRF Instance to Public IP Address | 290](#)

NAT Overview

Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. NAT was described in RFC 1631 to solve IPv4 address depletion problems. NAT is a useful tool for firewalls, traffic redirect, load sharing, and network migrations.

In an SD-WAN deployment, SRX Series Firewalls are deployed in the hub and spoke locations. Different sites are connected to the spoke SRX Series Firewall. Packets are sent from these sites to public Internet servers or remote sites. At the hub, after the security processing is complete, the packet is examined to determine whether the destination is a public Internet server or an MPLS next-hop device. If the destination is a public Internet server, NAT converts the virtual routing and forwarding (VRF) private IP address to a public IP address and establishes a session. Similarly, NAT is required for traffic from public Internet servers to reach a VRF private network.

The following types of NAT are supported on Juniper Networks devices:

- Static NAT
- Destination NAT
- Source NAT

Example: Configuring Source NAT to convert the private IP address of a VRF instance to the private IP address of another VRF instance

IN THIS SECTION

- [Requirements | 275](#)
- [Overview | 276](#)
- [Configuration | 276](#)

This example describes how to configure a source NAT between two MPLS networks.

Requirements

- Understand how SRX Series Firewalls work in an SD-WAN deployment for NAT. See ["NAT Overview" on page 275](#).

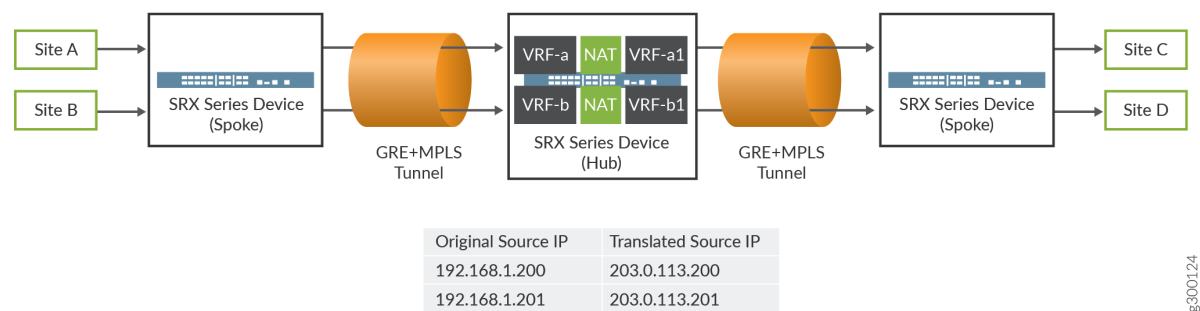
- Understand Virtual Routing and Forwarding Instances. See [Virtual Routing and Forwarding Instances in SD-WAN Deployments](#).

Overview

Source NAT is the translation of the source IP address of a packet leaving the Juniper Networks device. Source NAT is used to allow hosts with private IP addresses to access a public network.

In this example, the SRX Series Firewall connects two MPLS private networks to convert the private IP address from one VRF's private IP address to another VRF's private IP address. In [Figure 22 on page 276](#), the spoke SRX Series Firewall is configured with VRF-a and VRF-b routing instances, which are connected to the hub SRX Series Firewall. Site C and site D are connected to another spoke SRX Series Firewall. In the hub SRX Series Firewall, the source IP addresses 192.168.1.200 and 192.168.1.201 from VRF-a and VRF-b routing instances are translated to 203.0.113.200 and 203.0.113.201.

Figure 22: Source NAT conversion



Configuration

IN THIS SECTION

●

Verification | 281

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set routing-instances VRF-a instance-type vrf
set routing-instances VRF-a route-distinguisher 30:200
set routing-instances VRF-a vrf-target target:100:100
set routing-instances VRF-a vrf-table-label
set routing-instances VRF-b instance-type vrf
set routing-instances VRF-b route-distinguisher 40:200
set routing-instances VRF-b vrf-target target:200:100
set routing-instances VRF-b vrf-table-label
set routing-instances VRF-a1 instance-type vrf
set routing-instances VRF-a1 route-distinguisher 60:200
set routing-instances VRF-a1 vrf-target target:300:100
set routing-instances VRF-a1 vrf-table-label
set routing-instances VRF-b1 instance-type vrf
set routing-instances VRF-b1 route-distinguisher 50:200
set routing-instances VRF-b1 vrf-target target:400:100
set routing-instances VRF-b1 vrf-table-label
set security nat source pool vrf-a_p address 203.0.113.200
set security nat source rule-set vrf-a_rs from routing-instance VRF-a
set security nat source rule-set vrf-a_rs to routing-instance VRF-a1
set security nat source rule-set vrf-a_rs rule rule1 match source-address 192.168.1.200
set security nat source rule-set vrf-a_rs rule rule1 then source-nat pool vrf-a_p
set security nat source pool vrf-b_p address 203.0.113.201
set security nat source rule-set vrf-b_rs from routing-instance VRF-b
set security nat source rule-set vrf-b_rs to routing-instance VRF-b1
set security nat source rule-set vrf-b_rs rule rule2 match source-address 192.168.1.201
set security nat source rule-set vrf-b_rs rule rule2 then source-nat pool vrf-b_p
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure source NAT mapping:

1. Layer 3 VPNs require a VRF table for distributing routes within the networks. Create a VRF instance and specify the value **vrf**.

```
[edit routing-instances]
user@host#set VRF-a instance-type vrf
user@host#set VRF-b instance-type vrf
user@host#set VRF-a1 instance-type vrf
user@host#set VRF-b1 instance-type vrf
```

2. Assign a route distinguisher to the routing instance.

```
[edit routing-instances]
user@host#set VRF-a route-distinguisher 30:200
user@host#set VRF-b route-distinguisher 40:200
user@host#set VRF-a1 route-distinguisher 60:200
user@host#set VRF-b1 route-distinguisher 50:200
```

3. Create a community policy to import or export all routes.

```
[edit routing-instances]
user@host#set VRF-a vrf-target target:100:100
user@host#set VRF-b vrf-target target:200:100
user@host#set VRF-a1 vrf-target target:300:100
user@host#set VRF-b1 vrf-target target:400:100
```

4. Assign a single VPN label for all the routes in the VRF.

```
[edit routing-instances]
user@host#set VRF-a vrf-table-label
user@host#set VRF-a1 vrf-table-label
user@host#set VRF-b vrf-table-label
user@host#set VRF-b1 vrf-table-label
```

5. Create a source NAT pool.

```
[edit security nat source]
user@host#set vrf-a_p address 203.0.113.200
user@host#set vrf-b_p address 203.0.113.201
```

6. Create a source NAT rule set.

```
[edit security nat source]
user@host#set rule-set vrf-a_rs from routing-instance VRF-a
user@host#set rule-set vrf-a_rs to routing-instance VRF-a1
user@host#set rule-set vrf-b_rs from routing-instance VRF-b
user@host#set rule-set vrf-b_rs to routing-instance VRF-b1
```

7. Configure a rule that matches packets and translates the source IP address to an IP address in the source NAT pool.

```
[edit security nat source]
user@host# set rule-set vrf-a_rs rule rule1 match source-address 192.168.1.200
user@host# set rule-set vrf-a_rs rule rule1 then source-nat pool vrf-a_p
user@host# set rule-set vrf-b_rs rule rule2 match source-address 192.168.1.201
user@host# set rule-set vrf-b_rs rule rule2 then source-nat pool vrf-b_p
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show routing-instances` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
  source {
    pool vrf-a_p {
      address {
        203.0.113.200/32;
      }
    }
    pool vrf-b_p {
```

```

        address {
            203.0.113.201/32;
        }
    }
    rule-set vrf-a_rs {
        from routing-instance VRF-a;
        to routing-instance VRF-a1;
        rule rule1 {
            match {
                source-address 192.168.1.200/32;
            }
            then {
                source-nat {
                    pool {
                        vrf-a_p;
                    }
                }
            }
        }
    }
    rule-set vrf-b_rs {
        from routing-instance VRF-b;
        to routing-instance VRF-b1;
        rule rule2 {
            match {
                source-address 192.168.1.201/32;
            }
            then {
                source-nat {
                    pool {
                        vrf-b_p;
                    }
                }
            }
        }
    }
}

```

[edit]

user@host# show routing-instances

VRF-a {

```
        instance-type vrf;  
        route-distinguisher 30:200;  
        vrf-target target:100:100;  
        vrf-table-label;  
    }  
    VRF-a1 {  
        instance-type vrf;  
        route-distinguisher 60:200;  
        vrf-target target:300:100;  
        vrf-table-label;  
    }  
    VRF-b {  
        instance-type vrf;  
        route-distinguisher 40:200;  
        vrf-target target:200:100;  
        vrf-table-label;  
    }  
    VRF-b1 {  
        instance-type vrf;  
        route-distinguisher 50:200;  
        vrf-target target:400:100;  
        vrf-table-label;  
    }  
}
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Source NAT Rule Usage | 281](#)

Verifying Source NAT Rule Usage

Purpose

Verify that there is traffic matching the source NAT rule.

Action

From operational mode, enter the `show security nat source rule all` command. In the Translation hits field, verify whether there is traffic that matches the source NAT rule.

```

user@host>show security nat source rule all
Total rules: 2
Total referenced IPv4/IPv6 ip-prefixes: 2/0
source NAT rule: rule1                Rule-set: vrf-a_rs
  Rule-Id                            : 1
  Rule position                       : 1
  From routing instance               : VRF-a
  To routing instance                 : VRF-a1
  Match
    Source addresses                  : 192.168.1.200 - 192.168.1.200
  Action                             : vrf-a_p
    Persistent NAT type               : N/A
    Persistent NAT mapping type       : address-port-mapping
    Inactivity timeout                : 0
    Max session number               : 0
  Translation hits                    : 0
    Successful sessions               : 0
    Failed sessions                   : 0
  Number of sessions                  : 0
source NAT rule: rule2                Rule-set: vrf-b_rs
  Rule-Id                            : 2
  Rule position                       : 2
  From routing instance               : VRF-b
  To routing instance                 : VRF-b1
  Match
    Source addresses                  : 192.168.1.201 - 192.168.1.201
  Action                             : vrf-b_p
    Persistent NAT type               : N/A
    Persistent NAT mapping type       : address-port-mapping
    Inactivity timeout                : 0
    Max session number               : 0
  Translation hits                    : 0
    Successful sessions               : 0
    Failed sessions                   : 0
  Number of sessions                  : 0

```

Example: Configuring Destination NAT to Convert Public IP Address to VRF's Single Private IP Address of a VRF instance

IN THIS SECTION

- [Requirements | 283](#)
- [Overview | 283](#)
- [Configuration | 284](#)
- [Verification | 289](#)

This example describes how to configure the destination NAT mapping of a public IP address to the single VRF's private address for directing the packets to the correct VRF instance.

Requirements


- Understand how SRX Series Firewalls work in an SD-WAN deployment for NAT. See ["NAT Overview" on page 275](#).
- Understand Virtual Routing and Forwarding Instances. See [Virtual Routing and Forwarding Instances in SD-WAN Deployments](#).

Overview

Destination NAT is the translation of the destination IP address of a packet entering the Juniper Networks device. Destination NAT is used to redirect traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

In this example, an SRX Series Firewall is configured with destination NAT to convert a public IP address to the VRF private IP address of a VRF instance. The public IP address can be configured per VRF instance. In [Figure 23 on page 284](#), the SRX Series Firewall is configured with two VRF instances, VRF-a and VRF-b. The SRX Series Firewall converts the public IP address to private IP address of a VRF instance.

Figure 23: Destination NAT



Configuration

IN THIS SECTION

Procedure | 285

Results | 287

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set routing-instances VRF-a instance-type vrf
set routing-instances VRF-a route-distinguisher 30:200
set routing-instances VRF-a vrf-target target:100:100
set routing-instances VRF-a vrf-table-label
set routing-instances VRF-b instance-type vrf
set routing-instances VRF-b route-distinguisher 40:200
set routing-instances VRF-b vrf-target target:200:100
set routing-instances VRF-b vrf-table-label
set security nat destination pool vrf-a_p routing-instance VRF-a
set security nat destination pool vrf-a_p address 192.168.1.200
set security nat destination rule-set rs from interface ge-0/0/0
set security nat destination rule-set rs rule vrf-a_r match destination-address 203.0.113.200
set security nat destination rule-set rs rule vrf-a_r then destination-nat pool vrf-a_p
set security nat destination pool vrf-b_p routing-instance VRF-b
set security nat destination pool vrf-b_p address 192.168.1.201
set security nat destination rule-set rs from interface ge-0/0/1
set security nat destination rule-set rs rule vrf-b_r match destination-address 203.0.113.201
set security nat destination rule-set rs rule vrf-b_r then destination-nat pool vrf-b_p
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure destination NAT mapping for a single VRF:

1. Layer 3 VPNs require a VRF table for distributing routes within the networks. Create a VRF instance and specify the value **vrf**.

```
[edit routing-instances]
user@host#set VRF-a instance-type vrf
user@host#set VRF-b instance-type vrf
```


2. Assign a route distinguisher to the routing instance.

```
[edit routing-instances]
user@host#set VRF-a route-distinguisher 30:200
user@host#set VRF-b route-distinguisher 40:200
```

3. Create a community policy to import or export all routes.

```
[edit routing-instances]
user@host#set VRF-a vrf-target target:100:100
user@host#set VRF-b vrf-target target:200:100
```

4. Assign a single VPN label for all the routes in the VRF.

```
[edit routing-instances]
user@host#set VRF-a vrf-table-label
user@host#set VRF-b vrf-table-label
```

5. Specify a destination NAT IP address pool.

```
[edit security nat destination]
user@host# set pool vrf-a_p address 192.168.1.200
user@host# set pool vrf-b_p address 192.168.1.201
```

6. Assign the routing instance to the destination pool.

```
[edit security nat destination]
user@host# set pool vrf-a_p routing-instance VRF-a
user@host# set pool vrf-b_p routing-instance VRF-b
```

7. Create a destination NAT rule set.

```
[edit security nat destination]
user@host# set rule-set rs from interface ge-0/0/0
user@host# set rule-set rs from interface ge-0/0/1
```

8. Configure a rule that matches packets and translates the destination IP address to an IP address in the destination NAT IP address pool.

```
[edit security nat destination]
user@host# set rule-set rs rule vrf-a_r match destination-address 203.0.113.200
user@host# set rule-set rs rule vrf-a_r then destination-nat pool vrf-a_p
user@host# set rule-set rs rule vrf-b_r match destination-address 203.0.113.201
user@host# set rule-set rs rule vrf-b_r then destination-nat pool vrf-b_p
```

Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show routing-instances` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
destination {
    pool vrf-a_p {
        routing-instance {
            VRF-a;
        }
        address 192.168.1.200/32;
    }
    pool vrf-b_p {
        routing-instance {
            VRF-b;
        }
        address 192.168.1.201/32;
    }
    rule-set rs {
        from interface [ ge-0/0/0.0 ge-0/0/1.0 ];
        rule vrf-a_r {
            match {
                destination-address 203.0.113.200/32;
            }
            then {
                destination-nat {
                    pool {
                        vrf-a_p;
                    }
                }
            }
        }
    }
}
```

```

    }
  }
}
rule vrf-b_r {
  match {
    destination-address 203.0.113.201/32;
  }
  then {
    destination-nat {
      pool {
        vrf-b_p;
      }
    }
  }
}
}
}
}
}

```

```

[edit]
user@host# show routing-instances
  VRF-a {
    instance-type vrf;
    route-distinguisher 30:200;
    vrf-target target:100:100;
    vrf-table-label;
  }
  VRF-b {
    instance-type vrf;
    route-distinguisher 40:200;
    vrf-target target:200:100;
    vrf-table-label;
  }
}

```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Destination NAT Rule Usage | 289](#)

Verifying Destination NAT Rule Usage

Purpose

Verify that there is traffic matching the destination NAT rule.

Action

From operational mode, enter the `show security nat destination rule all` command. In the Translation hits field, verify whether there is traffic that matches the destination NAT rule.

```
user@host> show security nat destination rule all
Total destination-nat rules: 2
Total referenced IPv4/IPv6 ip-prefixes: 2/0
Destination NAT rule: vrf-a_r           Rule-set: rs
  Rule-Id           : 1
  Rule position     : 1
  From interface    : ge-0/0/0.0
                   : ge-0/0/1.0
  Destination addresses : 203.0.113.200 - 203.0.113.200
  Action            : vrf-a_p
  Translation hits   : 0
    Successful sessions : 0
    Failed sessions    : 0
  Number of sessions : 0
Destination NAT rule: vrf-b_r           Rule-set: rs
  Rule-Id           : 2
  Rule position     : 2
  From interface    : ge-0/0/0.0
                   : ge-0/0/1.0
  Destination addresses : 203.0.113.201 - 203.0.113.201
  Action            : vrf-b_p
  Translation hits   : 0
```

```
Successful sessions      : 0
Failed sessions          : 0
Number of sessions       : 0
```

Example: Configuring Static NAT to Convert the Private IP Address of a VRF Instance to Public IP Address

IN THIS SECTION

- [Requirements | 290](#)
- [Overview | 290](#)
- [Configuration | 291](#)
- [Verification | 295](#)

This example describes how to configure a static NAT mapping of VRF single private IP address to a public IP address.


Requirements

Understand how SRX Series Firewalls work in an SD-WAN deployment for NAT. See ["NAT Overview" on page 275](#).

Overview

In this example, an SRX Series Firewall is configured with static NAT to convert the VRF private IP address of a VRF instance to a public IP address of a VRF instance. Static NAT can be applied on the source NAT and destination NAT. In [Figure 24 on page 291](#), the SRX Series Firewall is configured with two VRF instances, VRF-a and VFR-b. The SRX Series Firewall converts the private IP address of a VRF instance to a public IP address.

Figure 24: Static NAT



Configuration

IN THIS SECTION

- Procedure | 292
- Results | 294

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set routing-instances VRF-a instance-type vrf
set routing-instances VRF-a route-distinguisher 30:200
set routing-instances VRF-a vrf-target target:100:100
set routing-instances VRF-a vrf-table-label
set routing-instances VRF-b instance-type vrf
set routing-instances VRF-b route-distinguisher 40:200
set routing-instances VRF-b vrf-target target:200:100
set routing-instances VRF-b vrf-table-label
set security nat static rule-set rs from interface ge-0/0/0
set security nat static rule-set rs rule vrf-a_r match static-address 203.0.113.200
set security nat static rule-set rs rule vrf-a_r then static-nat prefix 192.168.1.200
set security nat static rule-set rs rule vrf-a_r then static-nat prefix routing-instance VRF-a
set security nat static rule-set rs from interface ge-0/0/1
set security nat static rule-set rs rule vrf-b_r match static-address 203.0.113.201
set security nat static rule-set rs rule vrf-b_r then static-nat prefix 192.168.1.201
set security nat static rule-set rs rule vrf-b_r then static-nat prefix routing-instance VRF-b
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure static NAT mapping for the IP address of a single VRF:

1. Layer 3 VPNs require a VRF table for distributing routes within the networks. Create a VRF instance and specify the value **vrf**.

```
[edit routing-instances]
user@host#set VRF-a instance-type vrf
user@host#set VRF-b instance-type vrf
```

2. Assign a route distinguisher to the routing instance.

```
[edit routing-instances]
user@host#set VRF-a route-distinguisher 30:200
user@host#set VRF-b route-distinguisher 40:200
```

3. Create a community policy to import or export all routes.

```
[edit routing-instances]
user@host#set VRF-a vrf-target target:100:100
user@host#set VRF-b vrf-target target:200:100
```

4. Assign a single VPN label for all the routes in the VRF.

```
[edit routing-instances]
user@host#set VRF-a vrf-table-label
user@host#set VRF-b vrf-table-label
```

5. Create a static NAT rule set.

```
[edit security nat static]
user@host# set rule-set rs from interface ge-0/0/0
user@host# set rule-set rs from interface ge-0/0/1
```

6. Configure a rule that matches packets and translates the destination address in the packets to a private IP address.

```
[edit security nat static]
user@host# set rule-set rs rule vrf-a_r match static-address 203.0.113.200
user@host# set rule-set rs rule vrf-a_r then static-nat prefix 192.168.1.200
user@host# set rule-set rs rule vrf-a_r then static-nat prefix routing-instance VRF-a
user@host# set rule-set rs rule vrf-b_r match static-address 203.0.113.201
user@host# set rule-set rs rule vrf-b_r then static-nat prefix 192.168.1.201
user@host# set rule-set rs rule vrf-b_r then static-nat prefix routing-instance VRF-b
```


Results

From configuration mode, confirm your configuration by entering the `show security nat` and `show routing-instances` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
[edit]
user@host# show security nat
static {
    rule-set rs {
        from interface [ ge-0/0/0.0 ge-0/0/1.0 ];
        rule vrf-a_r {
            match {
                destination-address 203.0.113.200/32;
            }
            then {
                static-nat {
                    prefix {
                        192.168.1.200/32;
                    }
                    routing-instance VRF-a;
                }
            }
        }
        rule vrf-b_r {
            match {
                destination-address 203.0.113.201/32;
            }
            then {
                static-nat {
                    prefix {
                        192.168.1.201/32;
                    }
                    routing-instance VRF-b;
                }
            }
        }
    }
}
```

```
}  
}
```

```
[edit]  
user@host# show routing-instances  
  VRF-a {  
    instance-type vrf;  
    route-distinguisher 30:200;  
    vrf-target target:100:100;  
    vrf-table-label;  
  }  
  VRF-b {  
    instance-type vrf;  
    route-distinguisher 40:200;  
    vrf-target target:200:100;  
    vrf-table-label;  
  }
```

If you are done configuring the device, enter `commit` from configuration mode.

Verification

IN THIS SECTION

- [Verifying Static NAT Rule Usage](#) | 295

Verifying Static NAT Rule Usage

Purpose

Verify that there is traffic matching the static NAT rule.

Action

From operational mode, enter the `show security nat static rule` command. In the Translation hits field, verify whether there is traffic that matches the static NAT rule.

```

user@host> show security nat static rule all
Total static-nat rules: 2
Total referenced IPv4/IPv6 ip-prefixes: 4/0
Static NAT rule: vrf-a_r           Rule-set: rs
  Rule-Id                : 1
  Rule position           : 1
  From interface          : ge-0/0/0.0
                        : ge-0/0/1.0
  Destination addresses   : 203.0.113.200
  Host addresses          : 192.168.1.200
  Netmask                  : 32
  Host routing-instance   : VRF-a
  Translation hits        : 0
    Successful sessions   : 0
    Failed sessions      : 0
  Number of sessions      : 0
Static NAT rule: vrf-b_r           Rule-set: rs
  Rule-Id                : 2
  Rule position           : 2
  From interface          : ge-0/0/0.0
                        : ge-0/0/1.0
  Destination addresses   : 203.0.113.201
  Host addresses          : 192.168.1.201
  Netmask                  : 32
  Host routing-instance   : VRF-b
  Translation hits        : 0
    Successful sessions   : 0
    Failed sessions      : 0
  Number of sessions      : 0

```

RELATED DOCUMENTATION

[Flow Management in SRX Series Devices Using VRF Routing Instance](#)

Understanding ALG Support for VRF Routing Instance

Configuring Security Policies for a VRF Routing Instance

NAT for VRF group

IN THIS SECTION

- [Overview | 297](#)
- [Example: Configuring Source NAT to convert the private IP address of a VRF Group to the private IP address of different VRF instance | 297](#)
- [Example: Configuring Destination NAT to Convert Public IP Address of a VRF Group to the private IP address of different VRF instance | 303](#)

Overview

In SD-WAN network, NAT is used when you convert the private IP to global IP pool in a VRF group. An SRX Series Firewall can be configured using the following VRF group NAT to translate the given IPs belonging to a given VRF group to different IPs belonging to different VRF instances:

- VRF group destination NAT
- VRF group source NAT
- VRF group static NAT

Example: Configuring Source NAT to convert the private IP address of a VRF Group to the private IP address of different VRF instance

IN THIS SECTION

- [Requirements | 298](#)
- [Overview | 298](#)
- [Configuration | 298](#)

This example describes how to configure a source NAT between two MPLS networks.

Requirements


- Understand how SRX Series Firewalls work in an SD-WAN deployment for NAT.
- Understand Virtual-Group in NAT, Virtual Routing and Forwarding Instances. See [Virtual Routing and Forwarding Instances in SD-WAN Deployments](#).

Overview

Source NAT is the translation of the source IP address of a packet leaving the Juniper Networks device. Source NAT is used to allow hosts with private IP addresses to access a public network.

In [Figure 25 on page 298](#), SRX Series Firewall is configured with VRF group vpn-A and vpn-B, which are connected to the interfaces ge-0/0/1.0 and ge-0/0/1.1 on SRX Series Firewall. In the hub SRX Series Firewall, the source IP addresses 192.168.1.200 and 192.168.1.201 from VRF group vpn-A and vpn-B are translated to 203.0.113.200 and 203.0.113.201.

Figure 25: Source NAT using VRF group



Configuration

IN THIS SECTION

- [Verification | 302](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security l3vpn vrf-group vpn-A vrf VRF-A1
set security l3vpn vrf-group vpn-A vrf VRF-A2
set security l3vpn vrf-group vpn-B vrf VRF-B1
set security l3vpn vrf-group vpn-B vrf VRF-B2
set security nat source pool vrf-a_p address 203.0.113.200
set security nat source rule-set vrf-a_rs from routing-group vpn-A
set security nat source rule-set vrf-a_rs to interface ge-0/0/1.0
set security nat source rule-set vrf-a_rs rule rule1 match source-address 192.168.1.200
set security nat source rule-set vrf-a_rs rule rule1 then source-nat pool vrf-a_p
set security nat source pool vrf-b_p address 203.0.113.201
set security nat source rule-set vrf-b_rs from routing-group vpn-B
set security nat source rule-set vrf-b_rs to interface ge-0/0/1.1
set security nat source rule-set vrf-b_rs rule rule2 match source-address 192.168.1.201
set security nat source rule-set vrf-b_rs rule rule2 then source-nat pool vrf-b_p
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure source NAT mapping:

1. In Layer 3 VPNs create a VRF group `vpn-A` with VRF instances `A1` and `A2`.

```
[edit security]
user@host#set l3vpn vrf-group vpn-A vrf VRF-A1
user@host#set l3vpn vrf-group vpn-A vrf VRF-A2
```

2. Create another VRF group vpn-B with VRF instances B1 and B2.

```
[edit security]
user@host#set l3vpn vrf-group vpn-B vrf VRF-B1
user@host#set l3vpn vrf-group vpn-B vrf VRF-B2
```

3. Create a source NAT pool.

```
[edit security nat source pool]
user@host#set vrf-a_p address 203.0.113.200
user@host#set vrf-b_p address 203.0.113.201
```

4. Create a source NAT rule set.

```
[edit security nat source]
user@host#set rule-set vrf-a_rs from routing-group vpn-A
user@host#set rule-set vrf-a_rs to interface ge-0/0/1.0
user@host#set rule-set vrf-b_rs from routing-group vpn-B
user@host#set rule-set vrf-b_rs to interface ge-0/0/1.1
```

5. Configure a rule that matches packets and translates the source IP address to an IP address in the source NAT pool.

```
[edit security nat source]
user@host# set rule-set vrf-a_rs rule rule1 match source-address 192.168.1.200
user@host# set rule-set vrf-a_rs rule rule1 then source-nat pool vrf-a_p
user@host# set rule-set vrf-b_rs rule rule2 match source-address 192.168.1.201
user@host# set rule-set vrf-b_rs rule rule2 then source-nat pool vrf-b_p
```

Results

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

```
[edit]
user@host# show security nat
  source {
    pool vrf-a_p {
      address {
        203.0.113.200/32;
      }
    }
    pool vrf-b_p {
      address {
        203.0.113.201/32;
      }
    }
  }
  rule-set vrf-a_rs {
    from routing-group vpn-A;
    to interface ge-0/0/1.0;
    rule rule1 {
      match {
        source-address 192.168.1.200/32;
      }
      then {
        source-nat {
          pool {
            vrf-a_p;
          }
        }
      }
    }
  }
  rule-set vrf-b_rs {
    from routing-group vpn-B;
    to interface ge-0/0/1.1;
    rule rule2 {
      match {
```



```

From routing-Group      : vpn-A
To interface            : ge-0/0/1.0
Match
  Source addresses      : 192.168.1.200 - 192.168.1.200
Action                  : vrf-a_p
  Persistent NAT type   : N/A
  Persistent NAT mapping type : address-port-mapping
  Inactivity timeout    : 0
  Max session number    : 0
Translation hits        : 0
  Successful sessions   : 0
  Failed sessions      : 0
Number of sessions      : 0
rule: rule2             Rule-set: vrf-b_rs
  Rule-Id               : 2
  Rule position         : 2
  From routing-Group    : vpn-B
  To interface          : ge-0/0/1.1
  Match
    Source addresses    : 192.168.1.201 - 192.168.1.201
  Action                : vrf-b_p
    Persistent NAT type : N/A
    Persistent NAT mapping type : address-port-mapping
    Inactivity timeout  : 0
    Max session number  : 0
Translation hits        : 0
  Successful sessions   : 0
  Failed sessions      : 0
Number of sessions      : 0

```

Example: Configuring Destination NAT to Convert Public IP Address of a VRF Group to the private IP address of different VRF instance

IN THIS SECTION

● [Requirements | 304](#)

● [Overview | 304](#)

- Configuration | 305
- Verification | 308

This example describes how to configure the destination NAT mapping of a public IP address of a VRF group to the single VRF's private address for directing the packets to the correct VRF instance.

Requirements

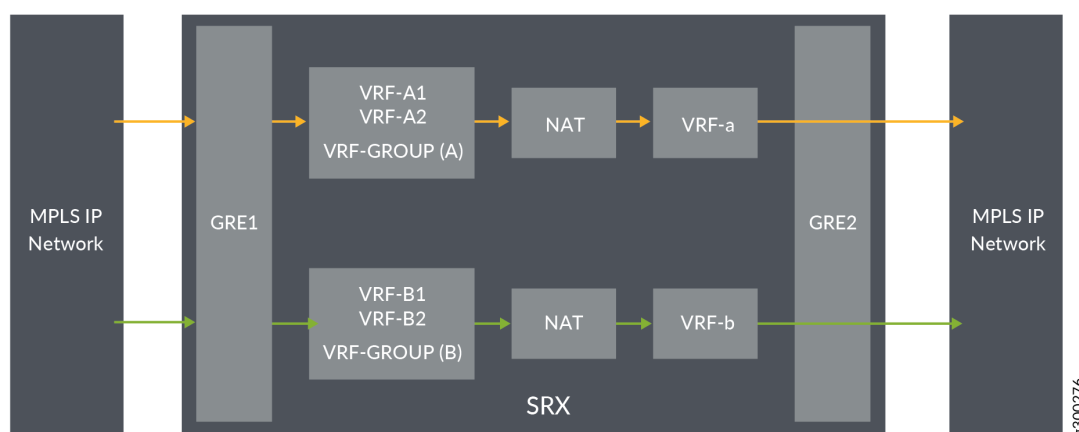
- Understand how SRX Series Firewalls work in an SD-WAN deployment for NAT.
- Understand Virtual Routing and Forwarding Instances. See [Virtual Routing and Forwarding Instances in SD-WAN Deployments](#).

Overview

Destination NAT is the translation of the destination IP address of a packet entering the Juniper Networks device. Destination NAT is used to redirect traffic destined to a virtual host (identified by the original destination IP address) to the real host (identified by the translated destination IP address).

In [Figure 26 on page 304](#), the SRX Series Firewall is configured destination NAT to convert from IP's that belong to different VRF groups, to different set of IP's with routing instance pointing to different VRF. After the destination NAT rule search, NAT updates the destination routing table to point to right VRF instance for flow to do destination route look-up in right table.

Figure 26: Destination NAT using VRF Group



Configuration

IN THIS SECTION

- [Procedure | 305](#)
- [Results | 307](#)

Procedure

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, copy and paste the commands into the CLI at the [edit] hierarchy level, and then enter `commit` from configuration mode.

```
set security l3vpn vrf-group vpn-A vrf VRF-A1
set security l3vpn vrf-group vpn-A vrf VRF-A2
set security l3vpn vrf-group vpn-B vrf VRF-B1
set security l3vpn vrf-group vpn-B vrf VRF-B2
set security nat destination pool vrf-a_p routing-instance VRF-a
set security nat destination pool vrf-a_p address 192.168.1.200
set security nat destination rule-set rs from routing-group vpn-A
set security nat destination rule-set rs rule vrf-a_r match destination-address 203.0.113.200
set security nat destination rule-set rs rule vrf-a_r then destination-nat pool vrf-a_p
set security nat destination pool vrf-b_p routing-instance VRF-b
set security nat destination pool vrf-b_p address 192.168.1.201
set security nat destination rule-set rs from routing-group vpn-B
set security nat destination rule-set rs rule vrf-b_r match destination-address 203.0.113.201
set security nat destination rule-set rs rule vrf-b_r then destination-nat pool vrf-b_p
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure destination NAT mapping for a single VRF:

1. In Layer 3 VPNs create a VRF group vpn-A with VRF instances A1 and A2.

```
[edit security]
user@host#set l3vpn vrf-group vpn-A vrf VRF-A1
user@host#set l3vpn vrf-group vpn-A vrf VRF-A2
```

2. Create another VRF group vpn-B with VRF instances B1 and B2.

```
[edit security]
user@host#set l3vpn vrf-group vpn-B vrf VRF-B1
user@host#set l3vpn vrf-group vpn-B vrf VRF-B2
```

3. Specify a destination NAT IP address pool.

```
[edit security nat destination]
user@host# set pool vrf-a_p address 192.168.1.200
user@host# set pool vrf-b_p address 192.168.1.201
```

4. Assign the routing instance to the destination pool.

```
[edit security nat destination]
user@host# set pool vrf-a_p routing-instance VRF-a
user@host# set pool vrf-b_p routing-instance VRF-b
```

5. Create a destination NAT rule set.

```
[edit security nat destination]
user@host# set rule-set rs from routing-group vpn-A
user@host# set rule-set rs from routing-group vpn-B
```

6. Configure a rule that matches packets and translates the destination IP address to an IP address in the destination NAT IP address pool.

```
[edit security nat destination]
user@host# set rule-set rs rule vrf-a_r match destination-address 203.0.113.200
user@host# set rule-set rs rule vrf-a_r then destination-nat pool vrf-a_p
```

```

user@host# set rule-set rs rule vrf-b_r match destination-address 203.0.113.201
user@host# set rule-set rs rule vrf-b_r then destination-nat pool vrf-b_p

```

Results

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

```

[edit]
user@host# show security nat
  destination {
    pool vrf-a_p {
      routing-instance {
        VRF-a;
      }
      address 192.168.1.200/32;
    }
    pool vrf-b_p {
      routing-instance {
        VRF-b;
      }
      address 192.168.1.201/32;
    }
  }
  rule-set rs {
    from routing-group [ vpn-A vpn-B ];
    rule vrf-a_r {
      match {
        destination-address 203.0.113.200/32;
      }
      then {
        destination-nat {
          pool {
            vrf-a_p;
          }
        }
      }
    }
    rule vrf-b_r {
      match {
        destination-address 203.0.113.201/32;
      }
    }
  }

```



```

Action                : vrf-a_p
Translation hits      : 0
  Successful sessions  : 0
  Failed sessions     : 0
  Number of sessions   : 0
Destination NAT rule: vrf-b_r          Rule-set: rs
Rule-Id               : 2
Rule position         : 2
From routing-group    : vpn-A
Destination addresses  : 203.0.113.201 - 203.0.113.201
Action                : vrf-b_p
Translation hits      : 0
  Successful sessions  : 0
  Failed sessions     : 0
  Number of sessions   : 0

```

RELATED DOCUMENTATION

[Flow Processing using Virtual Routing and Forwarding Group](#)

[*Configuring Security Policies Using VRF Group*](#)

[Understanding ALG Support for VRF group](#)

4

CHAPTER

Configuration Statements and Operational Commands

[Junos CLI Reference Overview](#) | 311

Junos CLI Reference Overview

We've consolidated all Junos CLI commands and configuration statements in one place. Learn about the syntax and options that make up the statements and commands and understand the contexts in which you'll use these CLI elements in your network configurations and operations.

- *Junos CLI Reference*

Click the links to access Junos OS and Junos OS Evolved configuration statement and command summary topics.

- *Configuration Statements*
- *CLI Commands*