

Junos® OS

Attack Detection and Prevention User Guide for Security Devices

Published
2024-12-18

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 Attack Detection and Prevention User Guide for Security Devices
Copyright © 2024 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

Attack Detection and Prevention Overview | 2

Screens Options for Attack Detection and Prevention | 3

Understanding Screens Options on SRX Series Devices | 3

Example: Configuring Multiple Screening Options | 13

Requirements | 14

Overview | 14

Configuration | 14

Verification | 18

Understanding Screen Options on the SRX5000 Module Port Concentrator | 20

Understanding IPv6 Support for Screens | 27

Understanding Screen IPv6 Tunneling Control | 32

Example: Improving Tunnel Traffic Security with IP Tunneling Screen Options | 36

Requirements | 36

Overview | 36

Configuration | 37

Verification | 42

2

Denial of Service Attacks

DoS Attack Overview | 47

Firewall DoS Attacks | 49

Understanding Session Table Flood Attacks | 49

Understanding Source-Based Session Limits | 49

Example: Setting Source-Based Session Limits | 51

Requirements | 51

Overview | 51

Configuration | 52

Verification | 54

Understanding Destination-Based Session Limits | 55

Example: Setting Destination-Based Session Limits | 56

Requirements | 56

Overview | 56

Configuration | 56

Verification | 58

Understanding SYN-ACK-ACK Proxy Flood Attacks | 59

Protecting Your Network Against a SYN-ACK-ACK Proxy Flood Attack | 59

Requirements | 60

Overview | 60

Configuration | 60

Verification | 62

Network DoS Attacks | 63

Network DoS Attacks Overview | 64

Understanding SYN Flood Attacks | 64

Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection | 68

Requirements | 68

Overview | 68

Configuration | 68

Verification | 70

Example: Enabling SYN Flood Protection for Webservers in the DMZ | 71

Requirements | 72

Overview | 72

Configuration | 75

Verification | 80

Understanding Allowlists for SYN Flood Screens | 80

Example: Configuring Allowlists for SYN Flood Screens | 81

Requirements | 81

Overview | 81

Configuration | **81**

Verification | **83**

Understanding Allowlist for UDP Flood Screens | **84**

Example: Configuring Allowlist for UDP Flood Screens | **84**

Requirements | **85**

Overview | **85**

Configuration | **85**

Verification | **87**

Understanding Allowlist for All Screen Options | **88**

Understanding SYN Cookie Protection | **89**

Detecting and Protecting Your Network Against SYN Flood Attacks by Enabling SYN Cookie Protection | **92**

Requirements | **92**

Overview | **92**

Configuration | **92**

Verification | **94**

Understanding ICMP Flood Attacks | **96**

Protecting Your Network Against ICMP Flood Attacks by Enabling ICMP Flood Protection | **98**

Requirements | **98**

Overview | **98**

Configuration | **98**

Verification | **100**

Understanding UDP Flood Attacks | **101**

Protecting Your Network Against UDP Flood Attacks by Enabling UDP Flood Protection | **102**

Requirements | **103**

Overview | **103**

Configuration | **103**

Verification | **105**

Understanding Land Attacks | **106**

Protecting Your Network Against Land Attacks by Enabling Land Attack Protection | **107**

Requirements | **107**

- Overview | 107
- Configuration | 107
- Verification | 109

OS-Specific DoS Attack | 110

OS-Specific DoS Attacks Overview | 111

Understanding Ping of Death Attacks | 111

Example: Protecting Against a Ping of Death Attack | 112

- Requirements | 112
- Overview | 112
- Configuration | 113
- Verification | 113

Understanding Teardrop Attacks | 114

Understanding WinNuke Attacks | 115

Example: Protecting Against a WinNuke Attack | 116

- Requirements | 117
- Overview | 117
- Configuration | 117
- Verification | 118

3

Suspicious Packets

Suspicious Packet Attributes Overview | 120

ICMP and SYN Fragment Attacks | 120

Understanding ICMP Fragment Protection | 121

Example: Blocking Fragmented ICMP Packets | 122

- Requirements | 122
- Overview | 122
- Configuration | 123
- Verification | 123

Understanding Large ICMP Packet Protection | 124

Example: Blocking Large ICMP Packets | 125

- Requirements | 125

- Overview | 125
- Configuration | 126
- Verification | 126

Understanding SYN Fragment Protection | 127

Example: Dropping IP Packets Containing SYN Fragments | 128

- Requirements | 128
- Overview | 128
- Configuration | 129
- Verification | 129

IP Packet Protection | 130

Understanding IP Packet Fragment Protection | 130

Example: Dropping Fragmented IP Packets | 132

- Requirements | 132
- Overview | 132
- Configuration | 133
- Verification | 133

Understanding Bad IP Option Protection | 134

Example: Blocking IP Packets with Incorrectly Formatted Options | 135

- Requirements | 135
- Overview | 135
- Configuration | 136
- Verification | 136

Understanding Unknown Protocol Protection | 137

Example: Dropping Packets Using an Unknown Protocol | 138

- Requirements | 138
- Overview | 138
- Configuration | 139
- Verification | 139

Understanding Allowlists for IP Block Fragment Screen | 140

Network Reconnaissance

Reconnaissance Deterrence Overview | 142

IP Address Sweep and Port Scan | 142

Understanding Network Reconnaissance Using IP Options | 143

Example: Detecting Packets That Use IP Screen Options for Reconnaissance | 147

Requirements | 147

Overview | 147

Configuration | 148

Verification | 150

Understanding IP Address Sweeps | 151

Example: Blocking IP Address Sweeps | 152

Requirements | 153

Overview | 153

Configuration | 153

Verification | 154

Understanding TCP Port Scanning | 155

Understanding UDP Port Scanning | 156

Enhancing Traffic Management by Blocking Port Scans | 158

Requirements | 158

Overview | 158

Configuration | 159

Verification | 160

Operating System Identification Probes | 162

Understanding Operating System Identification Probes | 162

Understanding Domain Name System Resolve | 163

Understanding TCP Headers with SYN and FIN Flags Set | 163

Example: Blocking Packets with SYN and FIN Flags Set | 164

Requirements | 165

Overview | 165

Configuration | 165

Verification | 166

Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set | 168

Example: Blocking Packets With FIN Flag Set and Without ACK Flag Set | 169

Requirements | 169

Overview | 169

Configuration | 170

Verification | 170

Understanding TCP Header with No Flags Set | 172

Example: Blocking Packets with No Flags Set | 173

Requirements | 173

Overview | 173

Configuration | 173

Verification | 174

Attacker Evasion Techniques | 176

Understanding Attacker Evasion Techniques | 176

Understanding FIN Scans | 177

Thwarting a FIN Scan | 177

Understanding TCP SYN Checking | 177

Setting TCP SYN Checking | 180

Setting TCP Strict SYN Checking | 180

Understanding IP Spoofing | 180

Example: Blocking IP Spoofing | 181

Requirements | 181

Overview | 181

Configuration | 181

Verification | 182

Understanding IP Spoofing in Layer 2 Transparent Mode on Security Devices | 184

Configuring IP Spoofing in Layer 2 Transparent Mode on Security Devices | 185

Understanding IP Source Route Options | 186

Example: Blocking Packets with Either a Loose or a Strict Source Route Option Set | 189

Requirements | 189

Overview | 189

Configuration | 190

Verification | 190

Example: Detecting Packets with Either a Loose or a Strict Source Route Option Set | 192

Requirements | 192

Overview | 192

Configuration | 192

Verification | 193

5

Configuration Statements and Operational Commands

Junos CLI Reference Overview | 197

About This Guide

Use this guide to configure the screen options in Junos OS on the SRX Series Firewalls to detect and prevent internal and external attacks, including SYN flood attacks, UDP flood attacks, and port scan attacks.

1

CHAPTER

Overview

[Attack Detection and Prevention Overview | 2](#)

[Screens Options for Attack Detection and Prevention | 3](#)

Attack Detection and Prevention Overview

Juniper Networks provides various detection and defense mechanisms at the zone and policy levels to combat exploits at all stages of their execution:

Attack detection and prevention, also known as stateful firewall, detects and prevents attacks in network traffic. An exploit can be either an information-gathering probe or an attack to compromise, disable, or harm a network or network resource. In some cases, the distinction between the two objectives of an exploit can be unclear. For example, a barrage of TCP SYN segments might be an IP address sweep with the intent of triggering responses from active hosts, or it might be a SYN flood attack with the intent of overwhelming a network so that it can no longer function properly. Furthermore, because an attacker usually precedes an attack by performing reconnaissance on the target, we can consider information-gathering efforts as a precursor to an impending attack—that is, they constitute the first stage of an attack. Thus, the term *exploit* encompasses both reconnaissance and attack activities, and the distinction between the two is not always clear.

- Screen options at the zone level.
- Firewall policies at the inter-, intra-, and super-zone policy levels (*super-zone* here means in global policies, where no security zones are referenced).

To secure all connection attempts, Junos OS uses a dynamic packet-filtering method known as stateful inspection. Using this method, Junos OS identifies various components in the IP packet and TCP segment headers—source and destination IP addresses, source and destination port numbers, and packet sequence numbers—and maintains the state of each TCP session and pseudo UDP session traversing the firewall. (Junos OS also modifies session states based on changing elements such as dynamic port changes or session termination.) When a responding TCP packet arrives, Junos OS compares the information reported in its header with the state of its associated session stored in the inspection table. If they match, the responding packet is allowed to pass the firewall. If the two do not match, the packet is dropped.

Junos OS screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone.

Screens Options for Attack Detection and Prevention

IN THIS SECTION

- [Understanding Screens Options on SRX Series Devices | 3](#)
- [Example: Configuring Multiple Screening Options | 13](#)
- [Understanding Screen Options on the SRX5000 Module Port Concentrator | 20](#)
- [Understanding IPv6 Support for Screens | 27](#)
- [Understanding Screen IPv6 Tunneling Control | 32](#)
- [Example: Improving Tunnel Traffic Security with IP Tunneling Screen Options | 36](#)

Attack detection and prevention detects and defend the network against attacks. Using Screen options, Junos security platforms can protect against different internal and external attacks, For more information, see the following topics:

Understanding Screens Options on SRX Series Devices

IN THIS SECTION

- [Statistics-based screens | 4](#)
- [Signature-based screens | 6](#)
- [Understanding Central Point Architecture Enhancements for Screens | 9](#)
- [Implementation of Screen Options on SRX Series Devices | 10](#)

On all SRX Series Firewalls, the screens are divided into two categories:

Statistics-based screens

Table 1 on page 4 lists all the statistics-based screen options.

Table 1: Statistics-Based Screen Options

Screen Option Name	Description
ICMP flood	<p>Use the ICMP flood IDS option to protect against ICMP flood attacks. An ICMP flood attack typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.</p> <p>The threshold value defines the number of ICMP packets per second (pps) allowed to be send to the same destination address before the device rejects further ICMP packets.</p>
UDP flood	<p>Use the UDP flood IDS option to protect against UDP flood attacks. A UDP flood attack occurs when an attacker sends IP packets containing a UDP datagram with the purpose of slowing down the resources, such that valid connections can no longer be handled.</p> <p>The threshold value defines the number of UDP packets per second allowed to be send to the same destination IP address. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.</p>
TCP SYN flood source	<p>Use the TCP SYN flood source IDS option to set the source threshold value. The threshold value defines the number of SYN segments to be received per second before the device begins dropping connection requests.</p> <p>The applicable range is 4 through 500,000 SYN pps.</p>
TCP SYN flood destination	<p>Use the SYN flood destination IDS option to set the destination threshold value. The threshold value defines the number of SYN segments received per second before the device begins dropping connection requests.</p> <p>The applicable range is 4 through 500,000 SYN pps.</p>
TCP SYN flood	<p>Use the TCP SYN flood IDS option to detect and prevent SYN flood attacks. Such attacks occur when the connecting host continuously sends TCP SYN requests without replying to the corresponding ACK responses.</p>

Table 1: Statistics-Based Screen Options (Continued)

Screen Option Name	Description
TCP port scan	Use the TCP port scan IDS option to prevent the port scan attacks. The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.
TCP SYN-ACK-ACK proxy	Use the TCP SYN-ACK-ACK proxy screen option to prevent SYN-ACK-ACK attack. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, SRX Series Firewalls running Junos OS reject further connection requests from that IP address.
ICMP IP sweep	<p>Use the ICMP IP sweep IDS option to detect and prevent an IP sweep attack. An IP sweep attack occurs when an attacker sends ICMP echo requests (pings) to multiple destination addresses. If a target host replies, the reply reveals the target's IP address to the attacker. If the device receives 10 ICMP echo requests within the number of microseconds specified in this statement, it flags this as an IP sweep attack, and rejects the eleventh and all further ICMP packets from that host for the remainder of the second.</p> <p>The threshold value defines the maximum number of microseconds during which up to 10 ICMP echo requests from the same host are allowed into the device.</p>
TCP SYN flood alarm	Use the TCP SYN flood alarm IDS option to set the alarm threshold value. The threshold value defines the number of half-complete proxy connections per second at which the device makes entries in the event alarm log. The range is 1 through 500,000 requests per second.
TCP SYN flood attack	Use the TCP SYN flood attack IDS option to set the attack threshold value. The threshold value defines the number of SYN packets per second required to trigger the SYN proxy response. The range is 1 through 500,000 proxied pps.

Table 1: Statistics-Based Screen Options (Continued)

Screen Option Name	Description
UDP udp sweep	<p>Use the UDP udp sweep IDS option to detect and prevent UDP sweep attacks. In a UDP sweep attack, an attacker sends UDP packets to the target device. If the device responds to those packets, the attacker gets an indication that a port in the target device is open, which makes the port vulnerable to attack. If a remote host sends UDP packets to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as a UDP sweep attack.</p> <p>If the alarm-without-drop option is not set, the device rejects the eleventh and all further UDP packets from that host for the remainder of the specified threshold period.</p> <p>The threshold value defines the number of microseconds for which the device accepts 10 UDP packets from the same remote source to different destination addresses.</p>

Starting with Junos OS Release 15.1X49-D20 and Junos OS Release 17.3R1, the firewall generates only one log message every second irrespective of the number of packets that trigger the source or destination session limit. This behavior applies to flood protection screens with TCP-Synflood-src-based, TCP-Synflood-dst-based, and UDP flood protection.

Signature-based screens

[Table 2 on page 6](#) lists all the signature-based screen options.

Table 2: Signature-Based Screen Options

Screen Option Name	Description
TCP Winnuke	Enable or disable the TCP WinNuke attacks IDS option. WinNuke is a denial-of-service (DoS) attack targeting any computer on the Internet running Windows.
TCP SYN fragment	Use the TCP SYN fragment attack IDS option to drop any packet fragments used for the attack. A SYN fragment attack floods the target host with SYN packet fragments. The host caches these fragments, waiting for the remaining fragments to arrive so it can reassemble them. The flood of connections that cannot be completed eventually fills the host's memory buffer. No further connections are possible, and damage to the host's operating system can occur.

Table 2: Signature-Based Screen Options (Continued)

Screen Option Name	Description
TCP no flag	Use the TCP tcp no flag IDS option to drop illegal TCP packets with a missing or malformed flag field. The threshold value defines the number of TCP headers without flags set. A normal TCP segment header has at least one control flag set.
TCP SYN FIN	Use the TCP SYN FIN IDS option to detect an illegal combination of flags that attackers can use to consume sessions on the target device, thus resulting in a denial-of-service (DoS) condition.
TCP land	Enable or disable the TCP land attack IDS option. Land attacks occur when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and the source IP address.
TCP FIN no ACK	Use the FIN bit with no ACK bit IDS option to detect an illegal combination of flags, and reject packets that have this combination.
ICMP ping of death	<p>Use the ping of death IDS option to detect and reject oversized and irregular ICMP packets. Although the TCP/IP specification requires a specific packet size, many ping implementations allow larger packet sizes. Larger packets can trigger a range of adverse system reactions, including crashing, freezing, and restarting.</p> <p>Ping of death occurs when IP packets are sent that exceed the maximum legal length (65,535 bytes).</p>
ICMP fragment	Use the ICMP fragment IDS option to detect and drop any ICMP frame with the More Fragments flag set or with an offset indicated in the offset field.
ICMP large	Use the ICMP large IDS option to detect and drop any ICMP frame with an IP length greater than 1024 bytes.
IP unknown protocol	Use the IP unknown protocol IDS option to discard all received IP frames with protocol numbers greater than 137 for IPv4 and 139 for IPv6. Such protocol numbers are undefined or reserved.

Table 2: Signature-Based Screen Options *(Continued)*

Screen Option Name	Description
IP bad option	Use the IP bad IDS option to detect and drop any packet with an incorrectly formatted IP option in the IP packet header. The device records the event in the screen counters list for the ingress interface. This screen option is applicable to IPv4 and IPv6.
IP strict source route option	Use the IP strict source route IDS option to detect packets where the IP option is 9 (strict source routing), and record the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field. Currently, this screen option is applicable only to IPv4.
IP loose source route option	Use the IP loose source route IDS option to detect packets where the IP option is 3 (loose source routing), and record the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified. The type 0 routing header of the loose source route option is the only related header defined in IPv6.
IP source route option	Use the IP source route IDS option to detect packets and record the event in the screen counters list for the ingress interface.
IP stream option	Use the IP stream IDS option to detect packets where the IP option is 8 (stream ID), and record the event in the screen counters list for the ingress interface. This option provides a way for the 16-bit SATNET stream identifier to be carried through networks that do not support streams. Currently, this screen option is applicable only to IPv4.
IP block fragment	Enable or disable the IP packet fragmentation blocking. When this feature is enabled, Junos OS denies IP fragments on a security zone and blocks all IP packet fragments that are received at interfaces bound to that zone.

Table 2: Signature-Based Screen Options (Continued)

Screen Option Name	Description
IP record route option	Use the IP record route IDS option to detect packets where the IP option is 7 (record route), and record the event in the screen counters list for the ingress interface. This option records the IP addresses of the network devices along the path that the IP packet travels. Currently, this screen option is applicable only to IPv4.
IP timestamp option	Use the IP timestamp IDS option to detect packets where the IP option list includes option 4 (Internet timestamp), and record the event in the screen counters list for the ingress interface. This option records the time (in Universal Time) when each network device receives the packet during its trip from the point of origin to its destination. Currently, this screen option is applicable only to IPv4.
IP security option	Use the IP security IDS option to detect packets where the IP option is 2 (security), and record the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
IP spoofing	Use the IP address spoofing IDS option to prevent spoofing attacks. IP spoofing occurs when an invalid source address is inserted in the packet header to make the packet appear to come from a trusted source.
IP tear drop	Use the IP tear drop IDS option to block teardrop attacks. Teardrop attacks occur when fragmented IP packets overlap and cause the host attempting to reassemble the packets to crash. The tear drop option directs the device to drop any packets that have such a discrepancy. Teardrop attacks exploit the reassembly of fragmented IP packets.

Understanding Central Point Architecture Enhancements for Screens

Starting with Junos OS Release 15.1X49-D30 and Junos OS Release 17.3R1, on SRX5400, SRX5600, and SRX5800 devices, the central point architecture is enhanced to achieve a higher number of connections per second (CPS). Due to the enhancements, the central point session and central point packet processing have been moved from the central point to the Services Processing Unit (SPU).

Previously, the central point had a session limit and if no resources (session limit entries) were available, then the packet was always permitted by the session limit. Now, both the central point and the SPU have session limits. If there are no resources available in the central point, but resources are available in the SPU, then the central point cannot limit the sessions but the SPU can limit the sessions.

The following scenarios describe when the central point and the SPU determine whether to permit or drop a packet.

- When the central point has no session limit entry and the SPU has a session limit entry:
 1. If the session limit counter of the SPU is larger than the threshold value, the packet is dropped.
 2. If the session limit counter of the SPU is not larger than the threshold value, the packet is permitted.
- When the SPU does not have a session limit entry:
 1. If the session limit counter of the SPU is larger than the threshold value, the packet is permitted.
 2. If the session limit counter of the SPU is not larger than the threshold, the packet is permitted.



NOTE: An extra message is sent to the central point to maintain accurate session counts might impact the number of connections per second (CPS) for screens. This impacts the source or destination session limit.

Global traffic statistics lacking a central point might impact some global view screens. Only the SYN cookie has no global view, and the global traffic statistics are handled by the SPU, so the counter might be not accurate as before. For other statistics-based screens, handled by both the central point and the SPU, the counters are accurate.

Previously, statistics-based screens were handled only by the central point and the log and the SNMP trap could be rate-limited strictly. Now both the central point and the SPU can generate the log and the SNMP trap independently. Therefore, the log and the SNMP trap might be larger than before.

Implementation of Screen Options on SRX Series Devices

The below table lists all the screen options implemented on SRX Series Firewalls and are supported on all SRX Series Firewalls.

Table 3: Screen Options Implemented on SRX Series Devices

Screens	Implemented on NP/CP/SPU	Support in Hash mode	Support in SOF mode
icmp-flood	NP	Yes	Yes
udp-flood	NP	Yes	Yes

Table 3: Screen Options Implemented on SRX Series Devices (Continued)

Screens	Implemented on NP/CP/SPU	Support in Hash mode	Support in SOF mode
winnuke	NP	Yes	Yes
tcp-port-scan	CP+SPU	Yes	Yes
udp-port-scan	CP+SPU	Yes	Yes
address-sweep	CP+SPU	Yes	Yes
tcp-sweep	CP+SPU	Yes	Yes
udp-sweep	CP+SPU	Yes	Yes
tear-drop	SPU	Yes	NO
syn-flood	SPU	Yes	Yes
syn-flood-src	NP	Yes	Yes
syn-flood-dst	NP	Yes	Yes
ip-spoofing	SPU	Yes	Yes
ping-of-death	NP	Yes	Yes
ip-option-src-route	NP	Yes	Yes
land	NP	Yes	Yes
syn-fragment	NP	Yes	Yes

Table 3: Screen Options Implemented on SRX Series Devices (Continued)

Screens	Implemented on NP/CP/SPU	Support in Hash mode	Support in SOF mode
tcp-no-flag	NP	Yes	Yes
unknown-protocol	NP	Yes	Yes
ip-option-bad	NP	Yes	Yes
ip-option-record-route	NP	Yes	Yes
ip-option-timestamp	NP	Yes	Yes
ip-option-security	NP	Yes	Yes
ip-option-loose-src-route	NP	Yes	Yes
ip-option-strict-src-route	NP	Yes	Yes
ip-option-stream	NP	Yes	Yes
icmp-fragment	NP	Yes	Yes
icmp-large-pkt	NP	Yes	Yes
syn-fin	NP	Yes	Yes
fin-no-ack	NP	Yes	Yes
src-session-limit	CP+SPU	Yes	Yes
syn-ack-ack-proxy	SPU	Yes	Yes

Table 3: Screen Options Implemented on SRX Series Devices (Continued)

Screens	Implemented on NP/CP/SPU	Support in Hash mode	Support in SOF mode
block-fragment	NP	Yes	Yes
dst-session-limit	CP+SPU	Yes	Yes
ipv6-ext-header	SPU	Yes	No
ipv6-ext-hbyh-option	SPU	Yes	No
ipv6-ext-dst-option	SPU	Yes	No
ipv6-ext-header-limit	SPU	Yes	No
ipv6-malformed-header	SPU	Yes	No
icmpv6-malformed-packet	SPU	Yes	No
ip-tunnel-summary	SPU	Yes	No



NOTE: All the screen functionalities supported on the IOC1 card are supported on the IOC2 and IOC3 cards. On the SRX5000 line of devices and on the SRX4600 device, the Network Processor Unit (NPU) in an IOC2 card is replaced by the Lookup Unit (LU).

Example: Configuring Multiple Screening Options

IN THIS SECTION



Requirements | 14

- Overview | 14
- Configuration | 14
- Verification | 18

This example shows how to create one intrusion detection service (IDS) profile for multiple screening options.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In a security zone, you can apply one IDS profile to multiple screening options. In this example we are configuring the following screening options:

- ICMP screening
- IP screening
- TCP screening
- UDP screening

These screening options are assigned to an untrust zone.

Configuration

IN THIS SECTION

- CLI Quick Configuration | 15
- Procedure | 16

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 screen ids-option screen-config icmp ip-sweep threshold 1000
set security screen ids-option screen-config icmp fragment
set security screen ids-option screen-config icmp large
set security screen ids-option screen-config icmp flood threshold 200
set security screen ids-option screen-config icmp ping-death
set security screen ids-option screen-config ip bad-option
set security screen ids-option screen-config ip stream-option
set security screen ids-option screen-config ip spoofing
set security screen ids-option screen-config ip strict-source-route-option
set security screen ids-option screen-config ip unknown-protocol
set security screen ids-option screen-config ip tear-drop
set security screen ids-option screen-config tcp syn-fin
set security screen ids-option screen-config tcp tcp-no-flag
set security screen ids-option screen-config tcp syn-frag
set security screen ids-option screen-config tcp port-scan threshold 1000
set security screen ids-option screen-config tcp syn-ack-ack-proxy threshold 500
set security screen ids-option screen-config tcp syn-flood alarm-threshold 500
set security screen ids-option screen-config tcp syn-flood attack-threshold 500
set security screen ids-option screen-config tcp syn-flood source-threshold 50
set security screen ids-option screen-config tcp syn-flood destination-threshold 1000
set security screen ids-option screen-config tcp syn-flood timeout 10
set security screen ids-option screen-config tcp land
set security screen ids-option screen-config tcp winnuke
set security screen ids-option screen-config tcp tcp-sweep threshold 1000
set security screen ids-option screen-config udp flood threshold 500
set security screen ids-option screen-config udp udp-sweep threshold 1000
set security zones security-zone untrust screen screen-config
```

Enter `commit` from configuration mode.

Procedure

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*.

To configure an IDS profile for multiple screening options:

1. Configure the ICMP screening options.

```
[edit security screen ids-option screen-config]
user@host# set icmp ip-sweep threshold 1000
user@host# set icmp fragment
user@host# set icmp large
user@host# set icmp flood threshold 200
user@host# set icmp ping-death
```

2. Configure the IP screening options.

```
[edit security screen ids-option screen-config]
user@host# set ip bad-option
user@host# set ip stream-option
user@host# set ip spoofing
user@host# set ip strict-source-route-option
user@host# set ip unknown-protocol
user@host# set ip tear-drop
```

3. Configure the TCP screening options.

```
[edit security screen ids-option screen-config]
user@host# set tcp syn-fin
user@host# set tcp tcp-no-flag
user@host# set tcp syn-frag
user@host# set tcp port-scan threshold 1000
user@host# set tcp syn-ack-ack-proxy threshold 500
user@host# set tcp syn-flood alarm-threshold 500
user@host# set tcp syn-flood attack-threshold 500
user@host# set tcp syn-flood source-threshold 50
user@host# set tcp syn-flood destination-threshold 1000
```

```

user@host# set tcp syn-flood timeout 10
user@host# set tcp land
user@host# set tcp winnuke
user@host# set tcp tcp-sweep threshold 1000

```

4. Configure the UDP screening options.

```

[edit security screen ids-option screen-config]
user@host# set udp flood threshold 500
user@host# set udp udp-sweep threshold 1000

```

5. Attach the IDS profile to the zone.

```

[edit]
user@host# set security zones security-zone untrust screen screen-config

```

Results

From configuration mode, confirm your configuration by entering the `show security screen ids-option screen-config` and `show security zones` 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 screen ids-option screen-config
icmp {
    ip-sweep threshold 1000;
    fragment;
    large;
    flood threshold 200;
    ping-death;
}
ip {
    bad-option;
    stream-option;
    spoofing;
    strict-source-route-option;
    unknown-protocol;
    tear-drop;
}

```

```

tcp {
    syn-fin;
    tcp-no-flag;
    syn-frag;
    port-scan threshold 1000;
    syn-ack-ack-proxy threshold 500;
    syn-flood {
        alarm-threshold 500;
        attack-threshold 500;
        source-threshold 50;
        destination-threshold 1000;
        timeout 10;
    }
    land;
    winnuke;
    tcp-sweep threshold 1000;
}
udp {
    flood threshold 500;
    udp-sweep threshold 1000;
}

```

```

[edit]
user@host# show security zones
security-zone untrust {
    screen screen-config;
}

```

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

Verification

IN THIS SECTION

- [Verifying the IDS Profile for Multiple Screening Options | 19](#)

Verifying the IDS Profile for Multiple Screening Options

Purpose

Verify that the IDS profile for multiple screening options is configured properly.

Action

Enter the `show security screen ids-option screen-config Screen object status` and `show security zones command` from operational mode.

```
user@host> show security screen ids-option screen-config
Screen object status:
```

Name	Value
ICMP flood threshold	200
UDP flood threshold	500
TCP winnuke	enabled
TCP port scan threshold	1000
ICMP address sweep threshold	1000
TCP sweep threshold	1000
UDP sweep threshold	1000
IP tear drop	enabled
TCP SYN flood attack threshold	500
TCP SYN flood alarm threshold	500
TCP SYN flood source threshold	50
TCP SYN flood destination threshold	1000
TCP SYN flood timeout	10
IP spoofing	enabled
ICMP ping of death	enabled
TCP land attack	enabled
TCP SYN fragment	enabled
TCP no flag	enabled
IP unknown protocol	enabled
IP bad options	enabled
IP strict source route option	enabled
IP stream option	enabled
ICMP fragmentation	enabled
ICMP large packet	enabled
TCP SYN FIN	enabled
TCP SYN-ACK-ACK proxy threshold	500

```
user@host> show security zones
```

```
Security zone: untrust
```

```
Send reset for non-SYN session TCP packets: Off
```

```
Policy configurable: Yes
```

```
Screen: screen-config
```

```
Interfaces bound: 0
```

```
Interfaces:
```



NOTE: On all SRX Series Firewalls, the TCP synchronization flood alarm threshold value does not indicate the number of packets dropped, however the value does show the packet information after the alarm threshold has been reached.

The synchronization cookie or proxy never drops packets; therefore the alarm-without-drop (not drop) action is shown in the system log.

Understanding Screen Options on the SRX5000 Module Port Concentrator

IN THIS SECTION

- [Statistics-Based Screens | 21](#)
- [Differences Between IOC1 and IOC2 | 22](#)
- [Signature-Based Screens | 26](#)

The SRX5000 line Module Port Concentrator (SRX5K-MPC) supports Junos OS screen options. Screen options secure a zone by inspecting, then allowing or denying, all connection attempts that require crossing an interface bound to that zone.

Using screen options, your security device can protect against different internal and external attacks, including SYN flood attacks, UDP flood attacks, and port scan attacks. Junos OS applies screen checks to traffic prior to the security policy processing, resulting in less resource utilization.

The screen options are divided into the following two categories:

- Statistics-based screens

- Signature-based screens

Statistics-Based Screens

All screen features implemented on an SRX5K-MPC are independent of Layer 2 or Layer 3 mode. The flood protections are used to defend against SYN flood attacks, session table flood attacks, firewall denial-of-service (DoS) attacks, and network DoS attacks.

The following four types of threshold-based flood protection are performed on each processor for both IPv4 and IPv6:

- UDP-based flood protection
- ICMP-based flood protection
- TCP source-based SYN flood protection
- TCP destination-based SYN flood protection



NOTE: If one of the two types of TCP SYN flood protections is configured on a zone, the second type of TCP SYN flood protection is automatically enabled on the same zone. These two types of protections always work together.

Each type of flood protection is threshold-based, and the threshold is calculated per zone on each microprocessor. If the flood is detected on a microprocessor chip, that particular microprocessor takes action against the offending packets based on the configuration:

- Default action (report and drop)—Screen logging and reporting are done on an SPU, so offending packets need to be forwarded to the central point or SPU for this purpose. To protect SPUs from flooding, only the first offending packet for each screen in a zone is sent to the SPU for logging and reporting in each second. The rest of the offending packets are counted and dropped in a microprocessor.

For example, assume UDP flooding is configured at a logical interface with a threshold set to 5000 packets per second. If UDP packets come in at the rate of 20,000 per second, then about 5000 UDP packets are forwarded to the central point or SPU each second, and the remaining packets are detected as flooding. However, only one UDP flooding packet is sent to the SPU for logging and reporting in each second. The remaining packets are dropped in the microprocessor.

- Alarm only (alarm-without-drop)—An offending packet detected by screen protection is not dropped. It skips the rest of the screen checks and is forwarded to the central point or SPU with the screen result copied to its meta-header. It is not counted as a dropped packet.

Differences Between IOC1 and IOC2

The behavior of screens is the same whether the device has either IOC1 or an IOC2 card. However, there are differences in the threshold values for the statistics-based screens. [Table 4 on page 22](#) lists the statistics-based screen options and the behavior of the screens depending on whether the device has either IOC1 or an IOC2 card.

Table 4: Statistics-Based Screen Options

Screen Option Name	Description	IOC1	IOC2
ICMP flood	<p>Sets the ICMP flood threshold value. The ICMP flood screen option is used to protect against ICMP flood attacks. An ICMP flood attack typically occurs when ICMP echo requests use all resources in responding, such that valid network traffic can no longer be processed.</p> <p>The threshold value defines the number of ICMP packets per second allowed to ping the same destination address before the device rejects further ICMP packets.</p>	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised.	<p>On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps.</p> <p>On SRX5000 line devices, when the IOC2 card is in services-offload mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.</p>

Table 4: Statistics-Based Screen Options *(Continued)*

Screen Option Name	Description	IOC1	IOC2
UDP flood	<p>Sets the UDP flood threshold value. The UDP flood screen option is used to protect against UDP flood attacks. UDP flood attack occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the resources, such that valid connections can no longer be handled.</p> <p>The threshold value defines the number of UDP pps allowed to ping the same destination IP address/port pair. When the number of packets exceeds this value within any 1-second period, the device generates an alarm and drops subsequent packets for the remainder of that second.</p>	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised.	<p>On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps.</p> <p>On SRX5000 line devices, when the IOC2 card is in services-offload mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.</p>

Table 4: Statistics-Based Screen Options *(Continued)*

Screen Option Name	Description	IOC1	IOC2
TCP SYN flood source	<p>Sets the TCP SYN flood source threshold value. The threshold value defines the number of SYN segments to be received per second before the device begins dropping connection requests.</p> <p>The applicable range is 4 through 500,000 SYN pps.</p>	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised..	<p>On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps.</p> <p>On SRX5000 line devices, when the IOC2 card is in services-offload mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.</p>

Table 4: Statistics-Based Screen Options *(Continued)*

Screen Option Name	Description	IOC1	IOC2
TCP SYN flood destination	<p>Sets the TCP SYN flood destination threshold value. The threshold value defines the number of SYN segments received per second before the device begins dropping connection requests.</p> <p>The applicable range is 4 through 500,000 SYN pps.</p>	If the Incoming traffic exceeds the threshold pps, either the packets are dropped or an alarm is raised.	<p>On SRX5000 line devices with IOC2 card, there is a change in the screen configuration for lookup (LU) chips. There are four LU chips in each IOC2 card. If the incoming traffic exceeds the threshold value pps, the packets are dropped. For example, if the user specify the threshold value of 1000 pps, we configure 250 pps on each LU chip internally, so that the threshold value of 1000 pps gets distributed equally among the 4 LU chips. As an expected result, the user gets the overall threshold value of 1000 pps.</p> <p>On SRX5000 line devices, when the IOC2 card is in services-offload mode, only one LU chip will function. If the incoming traffic rate exceeds the threshold value, the packets are dropped as a result of the expected behavior.</p>



NOTE: On SRX5400, SRX5600, and SRX5800 line devices, the screen threshold value is set for each IOC in the DUT for the LAG/LACP and RLAG/RETH child links. When you have cross-IOC child interfaces as a part of LAG/LACP or RETH/RLAG interfaces and the ingress traffic is also traversing multiple child links across IOCs, set the threshold

value to match the total number of packets passed by the screen from multiple IOCs with the expected total number of packets per second (pps) at the egress interface.

Signature-Based Screens

The SRX5K-MPC provides signature-based screen options along with sanity checks on the received packet.

Sometimes packets received by the device are malformed or invalid, and they might cause damage to the device and network. These packets must be dropped during initial stages of processing.

For both signature-based screen options and sanity checks, the packet contents, including packet header, status and control bits, and extension headers (for IPv6), are examined. You can configure the screens according to your requirements, whereas packet sanity checks are performed by default.

The packet sanity checks and screen options are performed on packets received on ingress interfaces.

The processor does sanity checks and runs some screen features to detect the malformed and malicious ingress packets received from physical interfaces. Packets that fail a sanity check are counted and dropped.

The following packet sanity checks are supported:

- IPv4 sanity check
- IPv6 sanity check

The following screen features are supported:

- IP-based screen
- UDP-based screen
- TCP-based screen
- ICMP-based screen

The screen features are applicable to both IPv4 and IPv6 packets, with the exception of IP options screens, which only apply to IPv4 packets. If a packet is detected by one screen option, it skips the rest of the screen checks and is forwarded to the central point or Services Processing Unit (SPU) for logging and statistics collection.



NOTE: On SRX5400, SRX5600, and SRX5800 devices, the first path signature screen is performed first, followed by the fast path bad-inner-header screen.

Understanding IPv6 Support for Screens

IN THIS SECTION

- IPv6 Extension Header Checking and Filtering | 27
- Maximum Number of Extension Headers | 29
- Bad Option Extension Headers | 29
- ICMPv6 Checking and Filtering | 30
- IPv6 Packet Header Checking and Filtering | 31

Juniper Networks provides various detection and defense mechanisms at the zone and policy levels to combat exploits at all stages of their execution. Screen options are at the zone level. Junos OS screen options secure a zone by inspecting it, and then allowing or denying all connection attempts that require crossing an interface bound to that zone.

You can configure screen options to check and filter packets based on IPv6 extension headers, packet headers, and ICMPv6 traffic. Based on your configuration, the screen can drop packets, create logs, and provide increased statistics for IPv6 traffic.

IPv6 Extension Header Checking and Filtering

You can use the `ipv6-extension-header` statement to selectively screen one or more extension headers. [Table 5 on page 27](#) lists common IPv6 extension headers and their type values.

Table 5: IPv6 Extension Headers and Type Values

Header Name	Header Type Value	Internet Standards
Authentication	51	RFC 2460
Encapsulating Security Payload	50	RFC 2460
Host Identify Protocol	139	RFC 5201

Table 5: IPv6 Extension Headers and Type Values *(Continued)*

Header Name	Header Type Value	Internet Standards
Destination Options <ul style="list-style-type: none"> • ILNP nonce option • Home address option • Line identification option • Tunnel encapsulation limit option 	60	RFC 2460
Fragment	44	RFC 2460
Hop-by-Hop Options <ul style="list-style-type: none"> • CALIPSO option • RPL option • SFM DPD option • Jumbo payload option • Quick start option • Router alert option 	0	RFC 2460
Mobility	135	RFC 6275
No next	59	RFC 2460
Routing	43	RFC 2460
Shim6	140	RFC 5533

Maximum Number of Extension Headers

You can specify the maximum number of permitted extension headers in a packet by using the `ipv6-extension-header-limit` statement. Although the maximum number of extension headers in a packet is not explicitly specified, the order of extension headers is recommended in RFC 2460:

1. Hop-by-Hop Options header
2. Destination Options header
3. Routing header
4. Fragment extension header
5. Authentication header
6. Encapsulating Security Payload header
7. Destination Options header

Each extension header should occur at most once, except for the destination options header, which should occur at most twice (once before a routing header and once before the upper-layer protocol header).

The maximum extension header number based on RFC 2460 is 7. Other extension headers have been defined by subsequent RFCs. We recommend the maximum extension header number to be in the range of 0 through 32.

Bad Option Extension Headers

You can configure screens to detect and drop any packet with an incorrectly formatted IP option in the IP packet header (IPv4 or IPv6). The device records the event in the screen counters list for the ingress interface. [Table 6 on page 29](#) lists key criteria that the device uses to screen packets for bad options.

Table 6: Bad Option Extension Header Screening Criteria

Screening Criteria	Internet Standards	Description
Routing extension header is after fragment header	RFC 2460	The order of extension headers in a packet is defined; accordingly, the fragment extension header must be after the routing header.

Table 6: Bad Option Extension Header Screening Criteria (Continued)

Screening Criteria	Internet Standards	Description
Wrong router alert parameter	RFC 2711	<p>This option is located in the hop-by-hop header and in the Junos OS implementation:</p> <ul style="list-style-type: none"> • There can be only one option of this type per hop-by-hop header • The header length must be 2. • There can be only one router alert option in one extension header.
More than one back-to-back pad option	draft-krishnan-ipv6-hopbyhop-00	This type of traffic is screened as error packets.
Non-zero payload in PadN option	RFC 4942	The system checks that the PadN only has zero octets in its payload.
Padding beyond the next eight-octet boundary	RFC 4942	The system checks for padding beyond the next eight octet boundary. There is no legitimate reason for padding beyond the next eight octet boundary.
Jumbo payload with non-zero IPv6 header payload	RFC 2675	The payload length field in the IPv6 header must be set to zero in every packet that carries the jumbo payload option.

ICMPv6 Checking and Filtering

You can enable ICMPv6 checking and filtering. The system then checks whether the ICMPv6 packet received matches the defined criteria and performs the specified action on matching packets. Some of the key defined criteria are as follows:

- Information message of unknown type—Many types of ICMPv6 information messages are defined, such as echo request (value 128), echo reply (value 129), and router solicitation (value 133). The maximum type definition is 149. Any value higher than 149 is treated as an unknown type and screened accordingly.

- Does not meet the ICMPv6 ND packet format rules (RFC 4861)—There are standard rules, such as the IP Hop limit field has a value of 255, ICMP checksum must be valid, the ICMP code must be 0, and so on.
- Malformed ICMPv6 packet filtering—For instance, the ICMPv6 packet is too big (message type 2), the next header is set to routing (43), and routing header is set to hop-by hop.

IPv6 Packet Header Checking and Filtering

You can enable the checking and filtering of IPv6 packet headers using the `ipv6-malformed-header` statement. Once enabled, the system verifies any incoming IPv6 packet to check if it matches any of the defined criteria. The system then performs the specified action (drop or alarm-without-drop) on matching packets. [Table 7 on page 31](#) lists key criteria that the device uses to screen packets.

Table 7: IPv6 Packet Header Screening Criteria

Screening Criteria	Internet Standards	Description
Deprecated site-local source and destination addresses	RFC 3879	The IPv6 site-local unicast prefix (1111111011 binary or FEC0::/10) is not supported.
Illegal multicast address scope values	RFC 4291	The unassigned multicast address scope values are treated as illegal.
Documentation-only prefix (2001:DB8::/32)	RFC 3849	IANA is to record the allocation of the IPv6 global unicast address prefix (2001:DB8::/32) as a documentation-only prefix in the IPv6 address registry. No end party is to be assigned this address.
Deprecated IPv4-compatible IPv6 source and destination addresses (::/96)	RFC 4291	The IPv4-compatible IPv6 address has been deprecated and is not supported.
ORCHID source and destination addresses (2001:10::/28)	RFC 5156	Addresses of the Overlay Routable Cryptographic Hash Identifiers (2001:10::/28) are used as identifiers and cannot be used for routing at the IP layer. Addresses within this block must not appear on the public Internet.

Table 7: IPv6 Packet Header Screening Criteria (*Continued*)

Screening Criteria	Internet Standards	Description
An IPv4 address embedded inside the IPv6 address (64:ff9b::/96) is an illegal, unacceptable IPv4 address	RFC 6052	The IPv6 address, 64:ff9b::/96, is reserved as “Well-known Prefix” for use in algorithmic mapping.

Understanding Screen IPv6 Tunneling Control

Several IPv6 transition methodologies are provided to utilize the tunneling of IPv6 packets over IPv4 networks that do not support IPv6. For this reason, these methods use public gateways and bypass the policies of the operator.

The security of tunneled packets is a major concern for service providers, because tunneled packets are easily accessed by attackers. Numerous IPv6 transition methodologies have evolved for sending tunneled packets through a network; however, because some of them operate on public gateways, they bypass the policies of the operator. This means that packet transmission is exposed to attackers. To overcome and secure transfer of packets, the IPv6 end nodes are required to de-capsulate the encapsulated data packets. Screen is one of the latest available technologies for blocking or allowing tunneling traffic based on user preferences.

You can configure the following screen options to check and filter packets based on IPv6 extension headers, packet headers, and Bad-Inner-Header IPv6 or IPv4 address validation. Based on your configuration, the screen can drop packets, create logs, and provide increased statistics for IP tunneling.

- **GRE 4in4 Tunnel:** The GRE 4in4 Tunnel screen matches the following signature: | IPv4 outer header | GRE header | IPv4 inner header

An outer IPv4 header must be **Protocol 47 GRE Encapsulation**. A GRE header must have **protocol E-type 0x0800 IPv4**. If these conditions are met, this packet is classified as GRE 4in4 tunnel signature.

- **GRE 4in6 Tunnel:** The GRE 4in6 Tunnel screen matches the following signature: IPv6 outer main header | IPv6 extension header(s) | GRE header | IPv4 inner header

An outer IPv6 main header or an IPv6 extension header must have a **Next Header of value 47 for GRE**. A GRE header must have **protocol E-type 0x0800 IPv4**. If these conditions are met, this packet is classified as GRE 4in6 tunnel signature.

- **GRE 6in4 Tunnel:** The GRE 6in4 Tunnel screen matches the following signature: IPv4 outer header | GRE header | IPv6 inner header

An outer IPv4 header must be **Protocol 47 GRE Encapsulation**. A GRE header must have **protocol E-type 0x086DD IPv6**. If these conditions are met, this packet is classified as GRE 6in4 tunnel signature.

- **GRE 6in6 Tunnel:** The GRE 6in6 Tunnel screen matches the following signature: IPv6 outer main header | IPv6 extension header(s) | GRE header | IPv6 inner header

An outer IPv6 main header or an IPv6 extension header must have a Next Header of **value 47 for GRE**. A GRE header must have **protocol E-type 0x086DD` IPv6**. If these conditions are met, this packet is classified as GRE 6in6 tunnel signature.

- **IPinIP 6to4relay Tunnel :** The IPinIP 6to4relay Tunnel screen matches the following signature: | IPv4 outer header | IPv6 inner header

An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation**. An outer header source address or destination address must be in network **192.88.99.0/24**. An inner IPv6 header source address or destination address must be in network **2002::/16**. If these conditions are met, this packet is classified as IPinIP 6to4relay tunnel signature.

- **IPinIP 6in4 Tunnel :** The IPinIP 6in4 Tunnel screen matches the following signature: | IPv4 outer header | IPv6 inner header

An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation**. If this condition is met, this packet is classified as IPinIP 6in4 tunnel signature.



NOTE: Typically, when IPv6 packets need to be transported in a complete IPv4 network, the IPv6 packets utilizes a point-to-point 6in4 tunnel.

- **IPinIP 6over4 Tunnel :** The IPinIP 6over4 Tunnel screen matches the following signature: | IPv4 outer header | IPv6 inner header

An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation:W**. An inner header source address or destination address must be in **fe80::/64** network. If these conditions are met, this packet is classified as IPinIP 6over4 tunnel signature.

- **IPinIP 4in6 Tunnel :** The IPinIP 4in6 Tunnel screen matches the following signature: | IPv6 outer main header | IPv6 extension header(s) | IPv4 inner header

An outer IPv6 header or an IPv6 extension header must have a Next Header of **value 04 for IPv4**. If these conditions are met, this packet is classified as IPinIP 4in6 tunnel signature.

- **IPinIP ISATAP Tunnel:** The IPinIP ISATAP Tunnel screen matches the following signature: | IPv6 outer main header | IPv6 inner header

An outer IPv4 header must be **Protocol 41 IPv6 Encapsulation**. An inner IPv6 header source address or destination address must be in **fe80::200:5efe/96** or **fe80::5efe/96** network. If these conditions are met, this packet is classified as IPinIP ISATAP tunnel signature.

- **IPinIP DS-Lite Tunnel:** The IPinIP DS-Lite Tunnel screen matches the following signature: | IPv6 outer main header | IPv6 extension header(s) | IPv4 inner header

An outer IPv6 header or an IPv6 extension header must have a Next Header of **value 04 for IPv4**. An inner IPv4 source address or destination address must be in **192.0.0.0/29** network. If these conditions are met, this packet is classified as IPinIP DS-Lite tunnel signature.

- **IPinIP 6in6 Tunnel:** The IPinIP 6in6 Tunnel screen matches the following signature: | IPv6 outer main header | IPv6 extension header(s) | IPv6 inner main header

An outer IPv6 main header or an IPv6 extension header must have a Next Header of **value 41 for IPv6**. An inner IPv6 main header must be **Version 6**. If these two conditions are met, this packet is classified as IPinIP 6in6 tunnel signature.

- **IPinIP 4in4 Tunnel:** The IPinIP 4in4 Tunnel screen matches the following signature: | IPv6 outer header | IPv4 inner header . An outer IPv4 header must have a Protocol of **value 04 for IPv4**. An inner IPv4 header must be **Version 4**.

- **IPinUDP Teredo Tunnel:** The IPinUDP Teredo Tunnel matches the following signature: IPv4 outer header | UDP header | IPv6 inner header

An outer IPv4 header must have a **Protocol of 17 for UDP payload**. A UDP header source or destination port must be **3544**. An inner IPv6 header source address or destination address must be in network **2001:0000:/32**.

- **IP Tunnel Bad Inner-Header Check:** The Bad Inner Header Tunnel screen checks the tunnel traffic inner header information for consistency. The packet drops when any of the following is detected:

- Inner header does not match outer header.
- Inner header TTL or Hop Limit must not be 0 or 255.
- Inner header IPv6 address checking.
- Inner header IPv4 address checking.
- Outer and Inner header length checks:
- Inner header IPv4 and IPv6 TCP/UDP/ICMP header length check:

TCP/UDP/ICMP header length must fit inside of inner IPv4/IPv6/EH6 header length when inner IP(v4/v6) is not a first, next, or last fragment.

- TCP: The minimum TCP header size must fit in the previous encapsulation length.

- ICMP: The minimum ICMP header size must fit in the previous encapsulation length.
- Fragmented packets: For fragmented packets, if the tunnel information needs to be checked for a screen and is not in the first fragment, then checking is not performed except the parts of the tunnel encapsulation that are included in the first fragment. Length checks are performed on first fragment packets using the actual packet buffer length, but the length checks are ignored because the inner header is larger than the outer header.
 - When the outer header is first fragment, do not examine the past physical packet length of the fragment.
 - When the inner header is a first fragment, do not examine the past length of the fragment.

For non-first fragment packets, checking is not performed in Bad Inner Header Tunnel screen.

- When outer header is a non-first fragment, examine the packet for screens that only use IP header signatures, because the payload cannot be examined.
- When inner header is a non-first fragment, do not examine the next packet.
- The IPv4 inner header checks that IPv4 header is from 20 to 50 bytes.



NOTE: On all SRX Series Firewalls, when a packet allow or drop session is established, the bad-inner-header screen is performed on every packet, because this screen is a fast path screen.

On SRX300, SRX320, SRX340, SRX345, SRX380, SRX1500, SRX4100, SRX4200 devices and vSRX Virtual Firewall instances., the fast-path bad-inner-header screen is always performed first, followed by the first path signature screen.

Starting with Junos OS Release 12.3X48-D10 and Junos OS Release 17.3R1, the syslog messages RT_SCREEN_IP and RT_SCREEN_IP_LS for the IP tunneling screen have been updated. The updated messages include the tunnel screen attacks and log-without-drop criteria. The following list illustrates some examples of these new system log messages for each of the tunnel types:

- RT_SCREEN_IP: Tunnel GRE 6in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: alarm-without-drop
- RT_SCREEN_IP: Tunnel GRE 6in6! source: 1212::12, destination: 1111::11, zone name: untrust, interface name: ge-0/0/1.0, action: drop
- RT_SCREEN_IP: Tunnel GRE 4in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: drop
- RT_SCREEN_IP_LS: [lsys: LSYS1] Tunnel GRE 6in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: alarm-without-drop

- RT_SCREEN_IP_LS: [lsys: LSYS1] Tunnel GRE 6in6! source: 1212::12, destination: 1111::11, zone name: untrust, interface name: ge-0/0/1.0, action: drop
- RT_SCREEN_IP_LS: [lsys: LSYS1] Tunnel GRE 4in4! source: 12.12.12.1, destination: 11.11.11.1, zone name: untrust, interface name: ge-0/0/1.0, action: drop

Example: Improving Tunnel Traffic Security with IP Tunneling Screen Options

IN THIS SECTION

- [Requirements | 36](#)
- [Overview | 36](#)
- [Configuration | 37](#)
- [Verification | 42](#)

This example shows how to configure the tunnel screens to enable the screens to control, allow, or block the transit of tunneled traffic.

Requirements

This example uses the following hardware and software components:

- An SRX Series Firewall
- Junos OS Release 12.3X48-D10 and later

Before you begin:

- Understand the IPv6 Tunneling control. See ["Understanding Screen IPv6 Tunneling Control" on page 32](#).

Overview

You can configure the following IP tunneling screen options to check and filter packets, based on IPv6 extension headers, packet headers, and bad-inner-header IPv6 or IPv4 address validation. Based on your configuration, the screen can drop packets, create logs, and provide increased statistics for IP tunneling. The following tunneling screen options are assigned to an untrust zone.

- GRE 4in4 Tunnel
- GRE 4in6 Tunnel
- GRE 6in4 Tunnel
- GRE 6in6 Tunnel
- IPinUDP Teredo Tunnel
- IPinIP 4in4 Tunnel
- IPinIP 4in6 Tunnel
- IPinIP 6in4 Tunnel
- IPinIP 6in6 Tunnel
- IPinIP 6over4 Tunnel
- IPinIP 6to4relay Tunnel
- IPinIP ISATAP Tunnel
- IPinIP DS-Lite Tunnel
- Bad Inner Header Tunnel

Configuration

IN THIS SECTION

- [Configuring GRE Tunnel Screens | 38](#)
- [Configuring an IPinUDP Teredo Tunnel Screen | 39](#)
- [Configuring an IPinIP Tunnel Screen | 39](#)
- [Configuring a Bad-Inner-Header Tunnel Screen | 41](#)
- [Results | 41](#)

To configure the IP tunneling screen options, perform these tasks:

Configuring GRE Tunnel Screens

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 screen ids-option screen1 ip tunnel gre gre-4in4
set security screen ids-option screen1 ip tunnel gre gre-4in6
set security screen ids-option screen1 ip tunnel gre gre-6in4
set security screen ids-option screen1 ip tunnel gre gre-6in6
set security zones security-zone untrust screen screen1
```

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 GRE tunnel screen:

1. Configure a GRE tunnel screen to check the tunnel traffic inner header information for consistency and validate the signature type screen.

```
[edit security screen ids-option screen1 ip tunnel gre]
user@host# set gre-4in4
user@host# set gre-4in6
user@host# set gre-6in4
user@host# set gre gre-6in6
```

2. Configure the screens in the security zones.

```
user@host#set security zones security-zone untrust screen screen1
```

Configuring an IPinUDP Teredo Tunnel Screen

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 screen ids-option screen1 ip tunnel ip-in-udp teredo
set security zones security-zone untrust screen screen1
```

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 IPinUDP Teredo tunnel screen:

1. Configure an IPinUDP Teredo tunnel screen to check the tunnel traffic inner header information for consistency and validate the signature type screen.

```
[edit security screen ids-option screen1 ip tunnel]
user@host# set ip-in-udp teredo
```

2. Configure the screens in the security zones.

```
user@host# set security zones security-zone untrust screen screen1
```

Configuring an IPinIP Tunnel Screen

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 screen ids-option screen1 ip tunnel ipip dslite
set security screen ids-option screen1 ip tunnel ipip ipip-4in4
```

```

set security screen ids-option screen1 ip tunnel ipip ipip-4in6
set security screen ids-option screen1 ip tunnel ipip ipip-6in4
set security screen ids-option screen1 ip tunnel ipip ipip-6in6
set security screen ids-option screen1 ip tunnel ipip ipip-6over4
set security screen ids-option screen1 ip tunnel ipip ipip-6to4relay
set security screen ids-option screen1 ip tunnel ipip ipip-isatap
set security zones security-zone untrust screen screen1

```

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 IPinIP tunnel screen:

1. Configure an IPinIP tunnel screen to check the tunnel traffic inner header information for consistency and validate the signature type screen.

```

[edit security screen ids-option screen1 ip tunnel ipip]
user@host# set dslite
user@host# set ipip-4in4
user@host# set ipip-4in6
user@host# set ipip-6in4
user@host# set ipip-6in6
user@host# set ipip-6over4
user@host# set ipip-6to4relay
user@host# set ipip-isatap

```

2. Configure the screens in the security zones.

```

user@host# set security zones security-zone untrust screen screen1

```

Configuring a Bad-Inner-Header Tunnel Screen

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 screen ids-option screen1 ip tunnel bad-inner-header
set security zones security-zone untrust screen screen1
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy.

To configure a bad-inner-header tunnel screen:

1. Configure a bad-inner-header tunnel screen to check the tunnel traffic inner header information for consistency.

```
[edit security screen ids-option screen1 ip tunnel]
user@host# set bad-inner-header
```

2. Configure the screens in the security zones.

```
user@host# set security zones security-zone untrust screen screen1
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security screen statistics zone untrust ip tunnel` commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

For brevity, this `show` output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
user@host# show security screen
...
```

```
ids-option screen1 {  
  ip{  
    tunnel {  
      gre {  
        gre-4in4;  
        gre-4in6;  
        gre-6in4;  
        gre-6in6;  
      }  
      ip-in-udp {  
        teredo;  
      }  
      ipip {  
        ipip-4in4;  
        ipip-4in6;  
        ipip-6in4;  
        ipip-6in6;  
        ipip-6over4;  
        ipip-6to4relay;  
        isatap;  
        dslite;  
      }  
      bad-inner-header;  
    }  
  }  
}
```

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

Verification

IN THIS SECTION

- [Verifying the Security Screen Configuration | 43](#)
- [Verifying IP Tunnel Screens in the Security Zones | 43](#)

Confirm that the configuration is working properly.

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen1` command.

```
user@host> show security screen ids-option screen1
show security screen ids-option screen1:
      Name                               Value
IP Tunnel Bad Inner Header              enabled
IP Tunnel GRE 6in4                      enabled
IP Tunnel GRE 4in6                      enabled
IP Tunnel GRE 6in6                      enabled
IP Tunnel GRE 4in4                      enabled
IP Tunnel IPinUDP Teredo                enabled
IP Tunnel IPIP 6to4 Relay                enabled
IP Tunnel IPIP 6in4                     enabled
IP Tunnel IPIP 6over4                   enabled
IP Tunnel IPIP 4in6                     enabled
IP Tunnel IPIP 4in4                     enabled
IP Tunnel IPIP 6in6                     enabled
IP Tunnel IPIP ISATAP                   enabled
IP Tunnel IPIP DS-Lite                   enabled
```

Meaning

The `show security screen ids-option screen1` command displays screen object status as enabled.

Verifying IP Tunnel Screens in the Security Zones

Purpose

Verify that the IP tunneling screen options are configured properly in the security zones.

Action

From operational mode, enter the `show security screen statistics zone untrust ip tunnel` command.

```
user@host> show security screen statistics zone untrust ip tunnel

IP Tunnel Screen statistics:

      IDS attack type                Statistics
IP tunnel GRE 6in4                   0
IP tunnel GRE 4in6                   0
IP tunnel GRE 6in6                   0
IP tunnel GRE 4in4                   0
IP tunnel IPIP 6to4 relay             0
IP tunnel IPIP 6in4                  0
IP tunnel IPIP 6over4                0
IP tunnel IPIP 4in6                  0
IP tunnel IPIP 4in4                  0
IP tunnel IPIP 6in6                  0
IP tunnel IPIP ISATAP                0
IP tunnel IPIP DS-Lite               0
IP tunnel IPinUDP Teredo             0
IP tunnel bad inner header           0
```

Meaning

The `show security screen statistics zone untrust ip tunnel` command displays the IP tunnel screen statistics summary.

Change History Table

Feature support is determined by the platform and release you are using. Use [Feature Explorer](#) to determine if a feature is supported on your platform.

Release	Description
15.1X49-D30	Starting with Junos OS Release 15.1X49-D30 and Junos OS Release 17.3R1, on SRX5400, SRX5600, and SRX5800 devices, the central point architecture is enhanced to achieve a higher number of connections per second (CPS).

15.1X49-D20	Starting with Junos OS Release 15.1X49-D20 and Junos OS Release 17.3R1, the firewall generates only one log message every second irrespective of the number of packets that trigger the source or destination session limit. This behavior applies to flood protection screens with TCP-Synflood-src-based, TCP-Synflood-dst-based, and UDP flood protection.
12.3X48-D10	Starting with Junos OS Release 12.3X48-D10 and Junos OS Release 17.3R1, the syslog messages RT_SCREEN_IP and RT_SCREEN_IP_LS for the IP tunneling screen have been updated.

2

CHAPTER

Denial of Service Attacks

[DoS Attack Overview](#) | 47

[Firewall DoS Attacks](#) | 49

[Network DoS Attacks](#) | 63

[OS-Specific DoS Attack](#) | 110

DoS Attack Overview

IN THIS SECTION

- [Firewall DoS Attacks Overview | 47](#)
- [Understanding Firewall Filters on the SRX5000 Module Port Concentrator | 48](#)

The intent of a denial-of-service (DoS) attack is to overwhelm the targeted victim with a tremendous amount of bogus traffic so that the victim becomes so preoccupied processing the bogus traffic that legitimate traffic cannot be processed. The target can be the firewall, the network resources to which the firewall controls access, or the specific hardware platform or operating system of an individual host.

If a DoS attack originates from multiple source addresses, it is known as a distributed denial-of-service (DDoS) attack. Typically, the source address of a DoS attack is spoofed. The source addresses in a DDoS attack might be spoofed, or the actual addresses of compromised hosts might be used as “zombie agents” to launch the attack.

The device can defend itself and the resources it protects from DoS and DDoS attacks.

Firewall DoS Attacks Overview

The intent of a denial-of-service (DoS) attack is to overwhelm the targeted victim with a tremendous amount of bogus traffic so that the victim becomes so preoccupied processing the bogus traffic that legitimate traffic cannot be processed.

If attackers discover the presence of the Juniper Networks firewall, they might launch a DoS attack against it instead of the network behind it. A successful DoS attack against a firewall amounts to a successful DoS attack against the protected network in that it thwarts attempts of legitimate traffic to traverse the firewall.

An attacker might use session table floods and SYN-ACK-ACK proxy floods to fill up the session table of Junos OS and thereby produce a DoS.

Understanding Firewall Filters on the SRX5000 Module Port Concentrator

The SRX5000 line Module Port Concentrator (SRX5K-MPC) for the SRX5400, SRX5600, and SRX5800 supports a firewall filter to provide filter based forwarding and packet filtering at logical interfaces including the chassis loopback interface. A firewall filter is used to secure networks, to protect Routing Engines and Packet Forwarding Engines, and to ensure class of service (CoS).

The firewall filter provides:

- Filter-based forwarding at logical interfaces
- Protection of a Routing Engine from DoS attacks
- Blocking of certain types of packets to reach a Routing Engine and packet counter

The firewall filter examines packets and performs actions according to the configured filter policy. The policy is composed of match conditions and actions. The match conditions cover various fields of Layer 3 packet and Layer 4 header information. In association with the match conditions, various actions are defined in the firewall filter policy, and these actions include accept, discard, log counter, and so on.

After configuring the firewall filter, you can apply a logical interface to the firewall filter in the ingress or egress, or in both directions. All packets passing through the logical interface are checked by the firewall filter. As part of the firewall filter configuration, a policer is defined and applied to the logical interface. A policer restricts the traffic bandwidth at the logical interface.



NOTE: Firewall filtering on an SRX5K-MPC does not support aggregated Ethernet interfaces.



NOTE: On SRX5400, SRX5600 and SRX5800 devices with an SRX5K-MPC, applying a policer at the loopback (lo0) interface ensures that the Packet Forwarding Engine discards certain types of packets and prevents them from reaching the Routing Engine.

RELATED DOCUMENTATION

[Network DoS Attacks Overview | 64](#)

[OS-Specific DoS Attacks Overview | 111](#)

Firewall DoS Attacks

IN THIS SECTION

- [Understanding Session Table Flood Attacks | 49](#)
- [Understanding Source-Based Session Limits | 49](#)
- [Example: Setting Source-Based Session Limits | 51](#)
- [Understanding Destination-Based Session Limits | 55](#)
- [Example: Setting Destination-Based Session Limits | 56](#)
- [Understanding SYN-ACK-ACK Proxy Flood Attacks | 59](#)
- [Protecting Your Network Against a SYN-ACK-ACK Proxy Flood Attack | 59](#)

DoS attack protection leverages stateful inspection to look for and then allow or deny all connection attempts that require crossing an interface on their way to and from the intended destination. For more information, see the following topics:

Understanding Session Table Flood Attacks

A successful DoS attack overwhelms its victim with such a massive barrage of false simulated traffic that it becomes unable to process legitimate connection requests. DoS attacks can take many forms—SYN flood, SYN-ACK-ACK flood, UDP flood, ICMP flood, and so on—but they all seek the same objective, which is to fill up their victim's session table.


When the session table is full, that host cannot create any new sessions and begins rejecting new connection requests. The source-based session limits screen option and the destination-based session limit screen option help mitigate such attacks.

Understanding Source-Based Session Limits

In addition to limiting the number of concurrent sessions from the same source IP address, you can also limit the number of concurrent sessions to the same destination IP address. One benefit of setting a source-based session limit is that it can stem an attack such as the Nimda virus (which is actually both a

virus and a worm) that infects a server and then begins generating massive amounts of traffic from that server. Because all the virus-generated traffic originates from the same IP address, a source-based session limit ensures that the firewall can curb such excessive amounts of traffic. See [Figure 1 on page 50](#).

Figure 1: Limiting Sessions Based on Source IP Address



Another benefit of source-based session limiting is that it can mitigate attempts to fill up the firewall's session table if all the connection attempts originate from the same source IP address.

Determining what constitutes an acceptable number of connection requests requires a period of observation and analysis to establish a baseline for typical traffic flows. You also need to consider the maximum number of concurrent sessions required to fill up the session table of the particular Juniper Networks platform you are using. To see the maximum number of sessions that your session table supports, use the CLI command `show security flow session summary`, and then look at the last line in the output, which lists the number of current (allocated) sessions, the maximum number of sessions, and the number of failed session allocations:

```
userhost# show security flow session summary
Unicast-sessions: 0
Multicast-sessions: 0
Failed-sessions: 0
Sessions-in-use: 0
  Valid sessions: 0
  Pending sessions: 0
  Invalidated sessions: 0
  Sessions in other states: 0
Maximum-sessions: 2097152
```

The default maximum for source-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.



NOTE: Junos OS supports source-based session limits for both IPv4 and IPv6 traffic.

Example: Setting Source-Based Session Limits

IN THIS SECTION

- [Requirements | 51](#)
- [Overview | 51](#)
- [Configuration | 52](#)
- [Verification | 54](#)

This example shows how to limit the amount of sessions based on source IP.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

IN THIS SECTION

- [Topology | 52](#)

The following example shows how to limit the number of sessions that any one server in the DMZ and in zone a can initiate. Because the DMZ contains only web servers, none of which should initiate traffic, you set the source-session limit at the lowest possible value, which is one session. On the other hand, zone a contains personal computers, servers, printers, and so on, many of which do initiate traffic. For zone a, you set the source-session limit to a maximum of 80 concurrent sessions.

Topology

Configuration

IN THIS SECTION

- [Procedure | 52](#)

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 screen ids-option 1-limit-session limit-session source-ip-based 1
set security zones security-zone dmz screen 1-limit-session
set security screen ids-option 80-limit-session limit-session source-ip-based 80
set security zones security-zone zone_a screen 80-limit-session
```

Step-by-Step Procedure

1. Specify the number of concurrent sessions based on source IP for the DMZ zone.

```
[edit security]
user@host# set screen ids-option 1-limit-session limit-session source-ip-based 1
```

2. Set the security zone for the DMZ to the configuration limit.

```
[edit security]
user@host# set zones security-zone dmz screen 1-limit-session
```

3. Specify the number of concurrent sessions based on source IP for the zone a zone.

```
[edit security]
user@host# set screen ids-option 80-limit-session limit-session source-ip-based 80
```

4. Set the security zone for zone a to the configuration limit.

```
[edit security]
user@host# set zones security-zone zone_a screen 80-limit-session
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zones` 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 screen
  ids-option 1-limit-session {
    limit-session {
      source-ip-based 1;
    }
  }
  ids-option 80-limit-session {
    limit-session {
      source-ip-based 80;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone dmz {
    screen 1-limit-session;
  }
  security-zone zone_a {
    screen 80-limit-session;
  }
```


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

Verification

IN THIS SECTION

- [Verifying Source-Based Session Limits | 54](#)

Verifying Source-Based Session Limits

Purpose

Verify source-based session limits.

Action

Enter the `show security screen ids-option 1-limit-session`, `show security screen ids-option 80-limit-session`, and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option 1-limit-session
```

```
Screen object status:
```

Name	Value
Session source limit threshold	1

```
user@host> show security screen ids-option 80-limit-session
```

```
Screen object status:
```

Name	Value
Session source limit threshold	80

```
user@host> show security zones
```

```
Security zone: dmz
```

```
Send reset for non-SYN session TCP packets: Off
```

```
Policy configurable: Yes
```

```
Screen: 1-limit-session
```

```
Interfaces bound: 0
```

```
Interfaces:
```

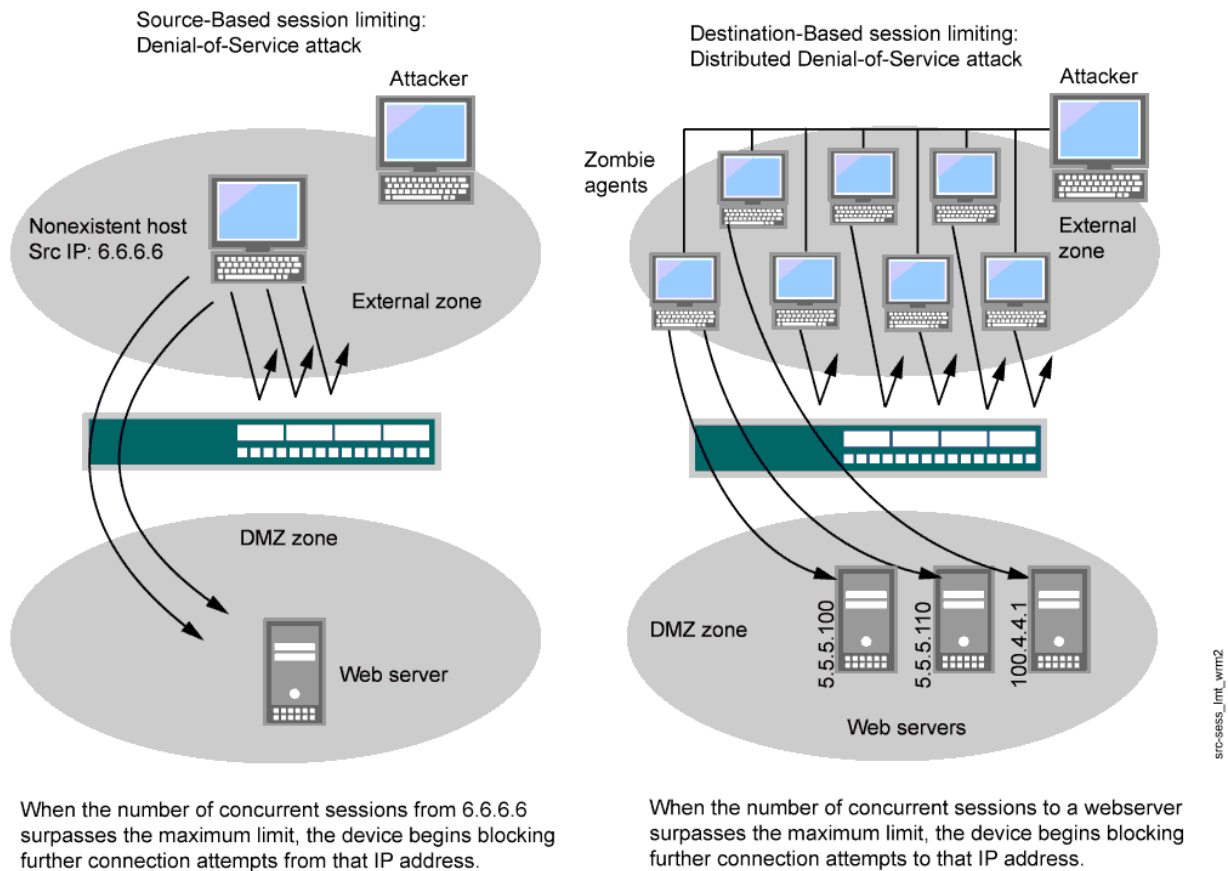
Meaning

The sample output shows the source session limit values for DMZ zone and zone a.

Understanding Destination-Based Session Limits

In addition to limiting the number of concurrent sessions from the same source IP address, you can also limit the number of concurrent sessions to the same destination IP address. A wily attacker can launch a distributed denial-of-service (DDoS) attack. In a DDoS attack, the malicious traffic can come from hundreds of hosts, known as “zombie agents,” that are surreptitiously under the control of an attacker. In addition to the SYN, UDP, and ICMP flood detection and prevention screen options, setting a destination-based session limit can ensure that Junos OS allows only an acceptable number of concurrent connection requests—no matter what the source—to reach any one host. See [Figure 2 on page 55](#).

Figure 2: Distributed DOS Attack



The default maximum for destination-based session limits is 128 concurrent sessions, a value that might need adjustment to suit the needs of your network environment and the platform.

Example: Setting Destination-Based Session Limits

IN THIS SECTION

- [Requirements | 56](#)
- [Overview | 56](#)
- [Configuration | 56](#)
- [Verification | 58](#)

This example shows how to set the destination-based session limits.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you limit the amount of traffic to a webserver at 1.2.2.5. The server is in the DMZ. The example assumes that after observing the traffic flow from the external zone to this server for a month, you have determined that the average number of concurrent sessions it receives is 2000. Also, you set the new session limit at 2000 concurrent sessions. Although traffic spikes might sometimes exceed that limit, the example assumes that you are opting for firewall security over occasional server inaccessibility.

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 screen ids-option 2000-limit-session limit-session destination-ip-based 2000
set security zones security-zone external_zone screen 2000-limit-session
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. To set the destination-based session limits:

1. Specify the number of concurrent sessions.

```
[edit]
user@host# set security screen ids-option 2000-limit-session limit-session destination-ip-
based 2000
```

2. Set the security zone for the external zone.

```
[edit]
user@host# set security zones security-zone external_zone screen 2000-limit-session
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zones` 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 screen
  ids-option 2000-limit-session {
    limit-session {
      destination-ip-based 2000;
```

```
    }
}
```

```
[edit]
user@host# show security zones
    security-zone external_zone {
        screen 2000-limit-session;
    }
```

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

Verification

IN THIS SECTION

- [Verifying Destination-Based Session Limits | 58](#)

Verifying Destination-Based Session Limits

Purpose

Verify destination-based session limits.

Action

Enter the `show security screen ids-option 2000-limit-session` and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option 2000-limit-session
node0:
```

```
-----
Screen object status:
```

Name	Value
Session destination limit threshold	2000
Value	

```

user@host> show security zones
Security zone: external_zone
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 2000-limit-session
  Interfaces bound: 0
  Interfaces:

```

Meaning

The sample output shows the destination session limit values for external zone.

Understanding SYN-ACK-ACK Proxy Flood Attacks

When an authentication user initiates a Telnet or an FTP connection, the user sends a SYN segment to the Telnet or FTP server. Junos OS intercepts the SYN segment, creates an entry in its session table, and proxies a SYN-ACK segment to the user. The user then replies with an ACK segment. At this point, the initial three-way handshake is complete. Junos OS sends a login prompt to the user. If the user, with malicious intent, does not log in but instead continues initiating SYN-ACK-ACK sessions, the firewall session table can fill up to the point where the device begins rejecting legitimate connection requests.

To prevent such an attack, you can enable the SYN-ACK-ACK proxy protection screen option. After the number of connections from the same IP address reaches the SYN-ACK-ACK proxy threshold, Junos OS rejects further connection requests from that IP address. By default, the threshold is 512 connections from any single IP address. You can change this threshold (to any number between 1 and 250,000) to better suit the requirements of your network environment.



NOTE: Junos OS supports SYN-ACK-ACK proxy protection for both IPv4 and IPv6 addresses.

Protecting Your Network Against a SYN-ACK-ACK Proxy Flood Attack

IN THIS SECTION



Requirements | 60

- [Overview | 60](#)
- [Configuration | 60](#)
- [Verification | 62](#)

This example shows how to protect your network against a SYN-ACK-ACK proxy flood attack.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable protection against a SYN-ACK-ACK proxy flood. The value unit is connections per source address. The default value is 512 connections from any single address.

Configuration

IN THIS SECTION

- [Procedure | 60](#)

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 screen ids-option 1000-syn-ack-ack-proxy tcp syn-ack-ack-proxy threshold 1000
set security zones security-zone zone screen 1000-syn-ack-ack-proxy
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. To protect against a SYN-ACK-ACK proxy flood attack:

1. Specify the source session limits.

```
[edit]
user@host# set security screen ids-option 1000-syn-ack-ack-proxy tcp syn-ack-ack-proxy
threshold 1000
```

2. Set the security zone for zone screen.

```
[edit]
user@host# set security zones security-zone zone screen 1000-syn-ack-ack-proxy
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zones` 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 screen
  ids-option 1000-syn-ack-ack-proxy {
    tcp {
      syn-ack-ack-proxy threshold 1000;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone zone {
    screen 1000-syn-ack-ack-proxy;
  }
```

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

Verification

IN THIS SECTION

- [Verifying SYN-ACK-ACK Proxy Flood Attack | 62](#)

Verifying SYN-ACK-ACK Proxy Flood Attack

Purpose

Verify SYN-ACK-ACK proxy flood attack.

Action

Enter the `show security screen ids-option 1000-syn-ack-ack-proxy` and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option 1000-syn-ack-ack-proxy
```

```
node0:
```

```
-----
```

```
Screen object status:
```

Name	Value
TCP SYN-ACK-ACK proxy threshold	1000

```
user@host> show security zones
```

```
Security zone: zone
```

```
Send reset for non-SYN session TCP packets: Off
```

```
Policy configurable: Yes
```

```
Screen: 1000-syn-ack-ack-proxy
```

```
Interfaces bound: 0
```

```
Interfaces:
```

Meaning

The sample output shows that there is no attack from SYN-ACK-ACK-proxy flood.

RELATED DOCUMENTATION

[OS-Specific DoS Attack | 110](#)

Network DoS Attacks

IN THIS SECTION

- [Network DoS Attacks Overview | 64](#)
- [Understanding SYN Flood Attacks | 64](#)
- [Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection | 68](#)
- [Example: Enabling SYN Flood Protection for Webservers in the DMZ | 71](#)
- [Understanding Allowlists for SYN Flood Screens | 80](#)
- [Example: Configuring Allowlists for SYN Flood Screens | 81](#)
- [Understanding Allowlist for UDP Flood Screens | 84](#)
- [Example: Configuring Allowlist for UDP Flood Screens | 84](#)
- [Understanding Allowlist for All Screen Options | 88](#)
- [Understanding SYN Cookie Protection | 89](#)
- [Detecting and Protecting Your Network Against SYN Flood Attacks by Enabling SYN Cookie Protection | 92](#)
- [Understanding ICMP Flood Attacks | 96](#)
- [Protecting Your Network Against ICMP Flood Attacks by Enabling ICMP Flood Protection | 98](#)
- [Understanding UDP Flood Attacks | 101](#)
- [Protecting Your Network Against UDP Flood Attacks by Enabling UDP Flood Protection | 102](#)
- [Understanding Land Attacks | 106](#)
- [Protecting Your Network Against Land Attacks by Enabling Land Attack Protection | 107](#)

A network attack consists of three major stages. In the first stage, the attacker performs reconnaissance on the target network. This reconnaissance might consist of many different kinds of network probes. For more information, see the following topics:

Network DoS Attacks Overview

A denial-of-service (DoS) attack directed against one or more network resources floods the target with an overwhelming number of SYN, ICMP, or UDP packets or with an overwhelming number of SYN fragments.

Depending on the attackers' purpose and the extent and success of previous intelligence gathering efforts, the attackers might single out a specific host, such as a device or server or they might aim at random hosts across the targeted network. Either approach has the potential of upsetting service to a single host or to the entire network, depending on how critical the role of the victim is to the rest of the network.

Understanding SYN Flood Attacks

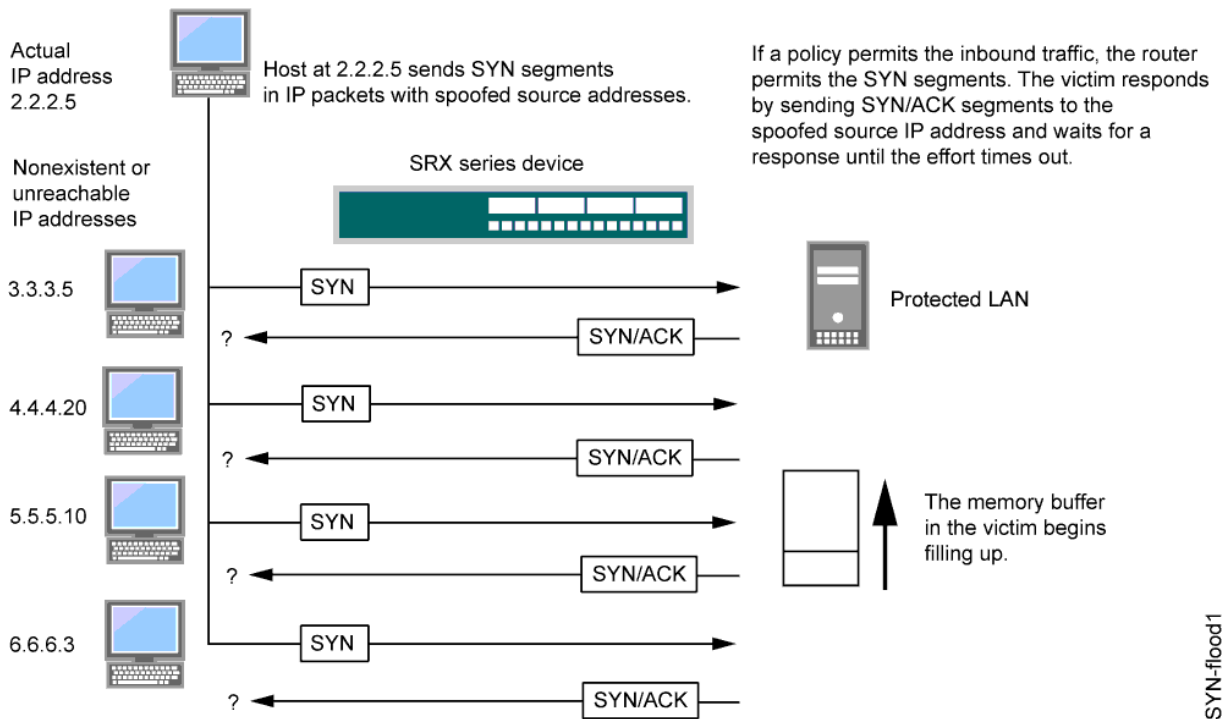
IN THIS SECTION

- SYN Flood Protection | 65
- SYN Flood Options | 66

A SYN flood occurs when a host becomes so overwhelmed by SYN segments initiating incomplete connection requests that it can no longer process legitimate connection requests.

Two hosts establish a TCP connection with a triple exchange of packets known as a *three-way handshake*: A sends a SYN segment to B; B responds with a SYN/ACK segment; and A responds with an ACK segment. A SYN flood attack inundates a site with SYN segments containing forged (spoofed) IP source addresses with nonexistent or unreachable addresses. B responds with SYN/ACK segments to these addresses and then waits for responding ACK segments. Because the SYN/ACK segments are sent to nonexistent or unreachable IP addresses, they never elicit responses and eventually time out. See [Figure 3 on page 65](#).

Figure 3: SYN Flood Attack



By flooding a host with incomplete TCP connections, the attacker eventually fills the memory buffer of the victim. Once this buffer is full, the host can no longer process new TCP connection requests. The flood might even damage the victim's operating system. Either way, the attack disables the victim and its normal operations.

This topic includes the following sections:

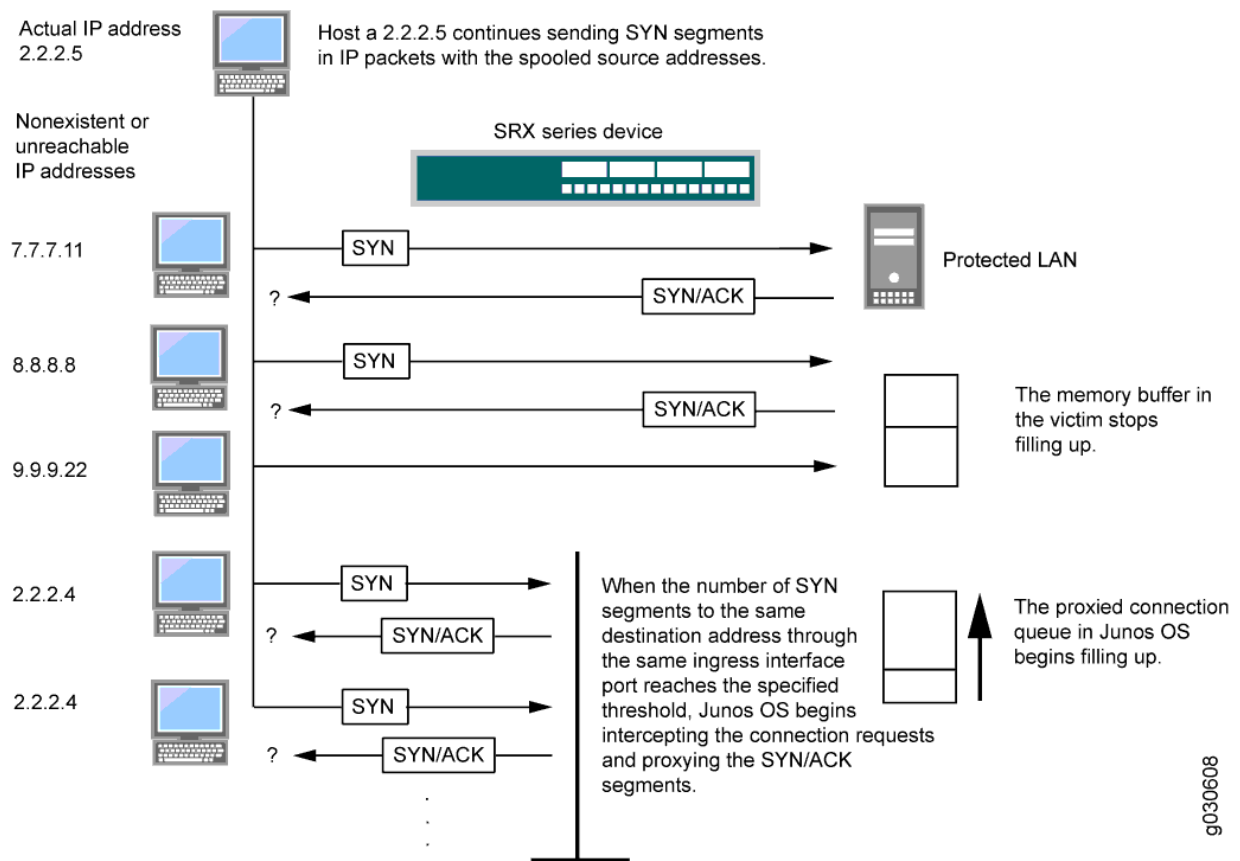
SYN Flood Protection

Junos OS can impose a limit on the number of SYN segments permitted to pass through the firewall per second. You can base the attack threshold on the destination address and ingress interface port, the destination address only, or the source address only. When the number of SYN segments per second exceeds the set threshold, Junos OS will either start proxying incoming SYN segments, replying with SYN/ACK segments and storing the incomplete connection requests in a connection queue, or it will drop the packets.

SYN proxying only happens when a destination address and ingress interface port attack threshold is exceeded. If a destination address or source address threshold is exceeded, additional packets are simply dropped.

In [Figure 4 on page 66](#), the SYN attack threshold for a destination address and ingress interface port has been exceeded and Junos OS has started proxying incoming SYN segments. The incomplete connection requests remain in the queue until the connection is completed or the request times out.

Figure 4: Proxying SYN Segments



SYN Flood Options

You can set the following parameters for proxying uncompleted TCP connection requests:

- **Attack Threshold**—This option allows you to set the number of SYN segments (that is, TCP segments with the SYN flag set) to the same destination address per second required to activate the SYN proxying mechanism. Although you can set the threshold to any number, you need to know the normal traffic patterns at your site to set an appropriate threshold for it. For example, if it is an e-business site that normally gets 20,000 SYN segments per second, you might want to set the threshold to 30,000 per second. If a smaller site normally gets 20 SYN segments per second, you might consider setting the threshold to 40.
- **Alarm Threshold**—This option allows you to set the number of proxied, half-complete TCP connection requests per second after which Junos OS enters an alarm in the event log. The value you set for an alarm threshold triggers an alarm when the number of proxied, half-completed connection requests to the same destination address per second exceeds that value. For example, if you set the SYN attack threshold at 2000 SYN segments per second and the alarm at 1000, then a total of 3000

SYN segments to the same destination address per second is required to trigger an alarm entry in the log.

For each SYN segment to the same destination address in excess of the alarm threshold, the attack detection module generates a message. At the end of the second, the logging module compresses all similar messages into a single log entry that indicates how many SYN segments to the same destination address and port number arrived after exceeding the alarm threshold. If the attack persists beyond the first second, the event log enters an alarm every second until the attack stops.

- **Source Threshold**—This option allows you to specify the number of SYN segments received per second from a single source IP address—regardless of the destination IP address—before Junos OS begins dropping connection requests from that source.

Tracking a SYN flood by source address uses different detection parameters from tracking a SYN flood by destination address. When you set a SYN attack threshold and a source threshold, you put both the basic SYN flood protection mechanism and the source-based SYN flood tracking mechanism in effect.

- **Destination Threshold**—This option allows you to specify the number of SYN segments received per second for a single destination IP address before Junos OS begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based on destination IP address only—regardless of the destination port number.

When you set a SYN attack threshold and a destination threshold, you put both the basic SYN flood protection mechanism and the destination-based SYN flood tracking mechanism in effect.

Consider a case where Junos OS has policies permitting FTP requests and HTTP requests to the same IP address. If the SYN flood attack threshold is 1000 packets per second (pps) and an attacker sends 999 FTP packets and 999 HTTP pps, Junos OS treats both FTP and HTTP packets with the same destination address as members of a single set and rejects the 1001st packet—FTP or HTTP—to that destination.

- **Timeout**—This option allows you to set the maximum length of time before a half-completed connection is dropped from the queue. The default is 20 seconds, and you can set the timeout from 1–50 seconds. You might try decreasing the timeout value to a shorter length until you begin to see any dropped connections during normal traffic conditions. Twenty seconds is a very conservative timeout for a three-way handshake ACK response.



NOTE: Junos OS supports SYN flood protection for both IPv4 and IPv6 traffic.

Protecting Your Network Against SYN Flood Attacks by Enabling SYN Flood Protection

IN THIS SECTION

- [Requirements | 68](#)
- [Overview | 68](#)
- [Configuration | 68](#)
- [Verification | 70](#)

This example shows how to protect your network against SYN flood attacks by enabling SYN flood protection.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable the zone-syn-flood protection screen option and set the timeout value to 20. You also specify the zone where the flood might originate.

Configuration

IN THIS SECTION

- [Procedure | 69](#)

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 screen ids-option zone-syn-flood tcp syn-flood source-threshold 10000
set security screen ids-option zone-syn-flood tcp syn-flood destination-threshold 10000
set security zones security-zone untrust screen zone-syn-flood
```

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 enable SYN flood protection:

1. Specify the screen object name.

```
[edit]
user@host# set security screen ids-option zone-syn-flood tcp syn-flood source-threshold 10000
user@host# set security screen ids-option zone-syn-flood tcp syn-flood destination-threshold
10000
```

2. Set the security zone for the zone screen.

```
[edit]
user@host# set security zones security-zone untrust screen zone-syn-flood
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zones` 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 screen
```



```
ids-option zone-syn-flood {  
    tcp {  
        syn-flood {  
            source-threshold 10000;  
            destination-threshold 10000;  
            timeout 20;  
        }  
    }  
}
```

```
[edit]  
user@host# show security zones  
    security-zone untrust {  
        screen zone-syn-flood;  
        interfaces {  
            ge-0/0/1.0;  
        }  
    }
```

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

Verification

IN THIS SECTION

- [Verifying SYN Flood Protection | 70](#)

Verifying SYN Flood Protection

Purpose

Verify SYN flood protection.

Action

Enter the `show security screen ids-option zone-syn-flood` and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option zone-syn-flood
node0:
-----
Screen object status:
Name                                     Value
TCP SYN flood attack threshold          200
TCP SYN flood alarm threshold            512
TCP SYN flood source threshold           10000
TCP SYN flood destination threshold       10000
TCP SYN flood timeout                    20

user@host> show security zones
Security zone: untrust
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: zone-syn-flood
  Interfaces bound: 1
  Interfaces:
    ge-0/0/1.0
```

Meaning

The sample output shows that SYN flood protection is enabled with source and destination threshold.

Example: Enabling SYN Flood Protection for Webserver in the DMZ

IN THIS SECTION

- [Requirements | 72](#)
- [Overview | 72](#)
- [Configuration | 75](#)

● Verification | 80

This example shows how to enable SYN flood protection for webserver in the DMZ.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

IN THIS SECTION


● Topology | 75

This example shows how to protect four webserver in the DMZ from SYN flood attacks originating in the external zone, by enabling the SYN flood protection screen option for the external zone. See [Figure 5 on page 73](#).



NOTE: We recommend that you augment the SYN flood protection that Junos OS provides with device-level SYN flood protection on each webserver. In this example, the webserver are running UNIX, which also provides some SYN flood defenses, such as adjusting the length of the connection request queue and changing the timeout period for incomplete connection requests.

Figure 5: Device-Level SYN Flood Protection



To configure the SYN flood protection parameters with appropriate values for your network, you must first establish a baseline of typical traffic flows. For example, for one week, you run a sniffer on ethernet3—the interface bound to zone_external—to monitor the number of new TCP connection requests arriving every second for the four web servers in the DMZ. Your analysis of the data accumulated from one week of monitoring produces the following statistics:

- Average number of new connection requests per server: 250 per second
- Average peak number of new connection requests per server: 500 per second



NOTE: A sniffer is a network-analyzing device that captures packets on the network segment to which you attach it. Most sniffers allow you to define filters to collect only the type of traffic that interests you. Later, you can view and evaluate the accumulated information. In this example, you want the sniffer to collect all TCP packets with the SYN flag set arriving at ethernet3 and destined for one of the four web servers in the DMZ. You might want to continue running the sniffer at regular intervals to see whether there are traffic patterns based on the time of day, day of the week, time of the month, or season of the year. For example, in some organizations, traffic might increase dramatically during a critical event. Significant changes probably warrant adjusting the various thresholds.

Based on this information, you set the following SYN flood protection parameters for zone_external as shown in [Table 8 on page 74](#).

Table 8: SYN Flood Protection Parameters

Parameter	Value	Reason for Each Value
Attack threshold	625 pps	This is 25% higher than the average peak number of new connection requests per second per server, which is unusual for this network environment. When the number of SYN packets per second for any one of the four web servers exceeds this number, the device begins proxying new connection requests to that server. (In other words, beginning with the 626th SYN packet to the same destination address in one second, the device begins proxying connection requests to that address.)
Alarm threshold	250 pps	When the device proxies 251 new connection requests in one second, it makes an alarm entry in the event log. By setting the alarm threshold somewhat higher than the attack threshold, you can avoid alarm entries for traffic spikes that only slightly exceed the attack threshold.

Table 8: SYN Flood Protection Parameters *(Continued)*

Parameter	Value	Reason for Each Value
Source threshold	25 pps	<p>When you set a source threshold, the device tracks the source IP address of SYN packets, regardless of the destination address. (Note that this source-based tracking is separate from the tracking of SYN packets based on destination address, which constitutes the basic SYN flood protection mechanism.)</p> <p>In the one week of monitoring activity, you observed that no more than 1/25 of new connection requests for all servers came from any one source within a one-second interval. Therefore, connection requests exceeding this threshold are unusual and provide sufficient cause for the device to execute its proxying mechanism. (Note that 25 pps is 1/25 of the attack threshold, which is 625 pps.)</p> <p>If the device tracks 25 SYN packets from the same source IP address, then, beginning with the 26th packet, it rejects all further SYN packets from that source for the remainder of that second and for the next second as well.</p>
Destination threshold	4000 pps	<p>When you set a destination threshold, the device runs a separate tracking of only the destination IP address, regardless of the destination port number. Because the four web servers receive only HTTP traffic (destination port 80)—no traffic to any other destination port number reaches them—setting another destination threshold offers no additional advantage.</p>
Timeout	20 seconds	<p>The default value of 20 seconds is a reasonable length of time to hold incomplete connection requests.</p>

Topology

Configuration

IN THIS SECTION

- Procedure | 76

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 interfaces ge-0/0/0 unit 0 family inet address 10.2.2.1/24
set interfaces fe-1/0/0 unit 0 family inet address 10.1.1.1/24
set security zones security-zone zone_dmz interfaces ge-0/0/0.0
set security zones security-zone zone_external interfaces fe-1/0/0.0
set security zones security-zone zone_dmz address-book address ws1 10.2.2.10/32
set security zones security-zone zone_dmz address-book address ws2 10.2.2.20/32
set security zones security-zone zone_dmz address-book address ws3 10.2.2.30/32
set security zones security-zone zone_dmz address-book address ws4 10.2.2.40/32
set security zones security-zone zone_dmz address-book address-set web_servers address ws1
set security zones security-zone zone_dmz address-book address-set web_servers address ws2
set security zones security-zone zone_dmz address-book address-set web_servers address ws3
set security zones security-zone zone_dmz address-book address-set web_servers address ws4
set security policies from-zone zone_external to-zone zone_dmz policy id_1 match source-address
any destination-address web_servers application junos-http
set security policies from-zone zone_external to-zone zone_dmz policy id_1 then permit
set security screen ids-option zone_external-syn-flood tcp syn-flood alarm-threshold 250 attack-
threshold 625 source-threshold 25 timeout 20
set security zones security-zone zone_external screen zone_external-syn-flood
```

Step-by-Step Procedure

To configure SYN flood protection parameters:

1. Set interfaces.

```
[edit]
user@host# set interfaces ge-0/0/0 unit 0 family inet address 10.2.2.1/24
user@host# set interfaces fe-1/0/0 unit 0 family inet address 10.1.1.1/24
user@host# set security zones security-zone zone_dmz interfaces ge-0/0/0.0
user@host# set security zones security-zone zone_external interfaces fe-1/0/0.0
```

2. Define addresses.

```
[edit]
user@host# set security zones security-zone zone_dmz address-book address ws1 10.2.2.10/32
user@host# set security zones security-zone zone_dmz address-book address ws2 10.2.2.20/32
user@host# set security zones security-zone zone_dmz address-book address ws3 10.2.2.30/32
user@host# set security zones security-zone zone_dmz address-book address ws4 10.2.2.40/32
user@host# set security zones security-zone zone_dmz address-book address-set web_servers
address ws1
user@host# set security zones security-zone zone_dmz address-book address-set web_servers
address ws2
user@host# set security zones security-zone zone_dmz address-book address-set web_servers
address ws3
user@host# set security zones security-zone zone_dmz address-book address-set web_servers
address ws4
```

3. Configure the policy.

```
[edit]
user@host# set security policies from-zone zone_external to-zone zone_dmz policy id_1 match
source-address any
user@host# set security policies from-zone zone_external to-zone zone_dmz policy id_1 match
destination-address web_servers
user@host# set security policies from-zone zone_external to-zone zone_dmz policy id_1 match
application junos-http
user@host# set security policies from-zone zone_external to-zone zone_dmz policy id_1 then
permit
```

4. Configure the screen options.

```
[edit]
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood alarm-
threshold 250
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood attack-
threshold 625
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood source-
threshold 25
user@host# set security screen ids-option zone_external-syn-flood tcp syn-flood timeout 20
user@host# set security zones security-zone zone_external screen zone_external-syn-flood
```


Results

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

For brevity, this `show` command output includes only the configuration that is relevant to this example. Any other configuration on the system has been replaced with ellipses (...).

```
[edit]
user@host# show interfaces
ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.2.2.1/24;
    }
  }
}
fe-1/0/0 {
  unit 0 {
    family inet {
      address 10.1.1.1/24;
    }
  }
}
...
[edit]
user@host# show security zones
...
  security-zone zone_dmz {
address-book {
address ws1 10.2.2.10/32;
  address ws2 10.2.2.20/32;
  address ws3 10.2.2.30/32;
  address ws4 10.2.2.40/32;
  address-set web_servers {
    address ws1;
    address ws2;
    address ws3;
    address ws4;
  }
}
```

```

interfaces {
    ge-0/0/0.0;
}
}
security-zone zone_external {
    screen zone_external-syn-flood;
    interfaces {
        fe-1/0/0.0;
    }
}
}
[edit]
user@host# show security policies
from-zone zone_external to-zone zone_dmz {
    policy id_1 {
    match {
    source-address any;
        destination-address web_servers;
        application junos-http;
    }
    then {
    permit;
        }
    }
}
[edit]
user@host# show security screen
...
ids-option zone_external-syn-flood {
    tcp {
    syn-flood {
    alarm-threshold 250;
        attack-threshold 625;
        source-threshold 25;
        timeout 20;
    }
    }
}
}

```

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

Verification

IN THIS SECTION

- [Verifying SYN Flood Protection for Webservers in the DMZ | 80](#)

Verifying SYN Flood Protection for Webservers in the DMZ

Purpose

Verify SYN flood protection for web servers in the DMZ.

Action

From operational mode, enter the `show interfaces`, `show security zones`, `show security policies`, and `show security screen ids-option zone_external-syn-flood` commands.

Understanding Allowlists for SYN Flood Screens

Junos OS provides the administrative option to configure a allowlist of trusted IP addresses to which the SYN flood screen will not reply with a SYN/ACK. Instead, the SYN packets from the source addresses or to the destination addresses in the list are allowed to bypass the SYN cookie and SYN proxy mechanisms. This feature is needed when you have a service in your network that cannot tolerate proxied SYN/ACK replies under any condition, including a SYN flood event.

Both IP version 4 (IPv4) and IP version 6 (IPv6) allowlists are supported. Addresses in a allowlist should be all IPv4 or all IPv6. In each allowlist, there can be up to 32 IP address prefixes. You can specify multiple addresses or address prefixes as a sequence of addresses separated by spaces and enclosed in square brackets.

A allowlist can cause high CPU usage on a central point depending on the traffic level. For example, when no screen is enabled, the connections per second (cps) is 492K; when the screen is enabled and the allowlist is disabled, the cps is 373K; and when both the screen and the allowlist are enabled, the cps is 194K. After enabling the allowlist, the cps drops by 40 percent.

Example: Configuring Allowlists for SYN Flood Screens

IN THIS SECTION

- Requirements | 81
- Overview | 81
- Configuration | 81
- Verification | 83

This example shows how to configure allowlists of IP addresses to be exempted from the SYN cookie and SYN proxy mechanisms that occur during the SYN flood screen protection process.

Requirements

Before you begin, configure a security screen and enable the screen in the security zone. See ["Example: Enabling SYN Flood Protection for Webservers in the DMZ" on page 71](#).

Overview

In this example, you configure allowlists named `wlipv4` and `wlipv6`. All addresses are IP version 4 (IPv4) for `wlipv4`, and all addresses are IP version 6 (IPv6) for `wlipv6`. Both allowlists include destination and source IP addresses.

Multiple addresses or address prefixes can be configured as a sequence of addresses separated by spaces and enclosed in square brackets, as shown in the configuration of the destination addresses for `wlipv4`.

Configuration

IN THIS SECTION

- Procedure | 82

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 screen ids-option js1 tcp syn-flood white-list wlipv4 source-address 1.1.1.0/24
set security screen ids-option js1 tcp syn-flood white-list wlipv4 destination-address
2.2.2.2/32
set security screen ids-option js1 tcp syn-flood white-list wlipv4 destination-address
3.3.3.3/32
set security screen ids-option js1 tcp syn-flood white-list wlipv4 destination-address
4.4.4.4/32
set security screen ids-option js1 tcp syn-flood white-list wlipv6 source-address 2001::1/64
set security screen ids-option js1 tcp syn-flood white-list wlipv6 destination-address 2002::1/64
```

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 *drop-profiles*.

To configure the allowlists:

1. Specify the name of the allowlist and the IP addresses to be exempted from the SYN/ACK.

```
[edit security screen ids-option js1 tcp syn-flood]
user@host# set white-list wlipv4 source-address 1.1.1.0/24
user@host# set white-list wlipv4 destination-address [2.2.2.2 3.3.3.3 4.4.4.4]
user@host# set white-list wlipv6 source-address 2001::1/64
user@host# set white-list wlipv6 destination-address 2002::1/64
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` 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 screen
ids-option js1 {
  tcp {
    syn-flood {
      white-list wlipv4 {
        source-address 1.1.1.0/24;
        destination-address [2.2.2.2/32 3.3.3.3/32 4.4.4.4/32];
      }
      white-list wlipv6 {
        source-address 2001::1/64;
        destination-address 2002::1/64;
      }
    }
  }
}
```

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

Verification

IN THIS SECTION

- [Verifying Whitelist Configuration | 83](#)

Verifying Whitelist Configuration

Purpose

Verify that the allowlist is configured properly.

Action

From operational mode, enter the `show security screen ids-option` command.

Understanding Allowlist for UDP Flood Screens

Junos OS provides the administrative option to configure a allowlist of trusted IP addresses on UDP flood. When UDP flood is enabled, all the UDP packets that are above the threshold value will be dropped. Some of these packets are valid and should not be dropped from the traffic. When you configure allowlist on UDP flood screen, only the source addresses in the list are allowed to bypass the UDP flood detection. This feature is needed when all traffic from addresses in the allowlist groups should bypass UDP flood check.

Both IPv4 and IPv6 allowlists are supported. Addresses in a allowlist should be all IPv4 or all IPv6. In each allowlist, there can be up to 32 IP address prefixes. You can specify multiple addresses or address prefixes as a sequence of addresses separated by spaces and enclosed in square brackets. You can configure single address or subnet address.



NOTE: UDP flood screen allowlist is not supported on SRX5400, SRX5600, and SRX5800 devices.

Example: Configuring Allowlist for UDP Flood Screens

IN THIS SECTION

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

This example shows how to configure allowlists of IP addresses to be exempted from UDP flood detection that occur during the UDP flood screen protection process.

Requirements

Before you begin, configure a security screen and enable the screen in the security zone.

Overview

In this example, you configure allowlists named `wlipv4` and `wlipv6`. All addresses are IPv4 for `wlipv4`, and all addresses are IPv6 for `wlipv6`. Both allowlists include destination and source IP addresses.

Multiple addresses or address prefixes can be configured as a sequence of addresses separated by spaces and enclosed in square brackets, as shown in the configuration of the destination addresses for `wlipv4` and `wlipv6`.

Configuration

IN THIS SECTION

- [Procedure | 85](#)

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 screen white-list wlipv4 address 198.51.100.10/24
set security screen white-list wlipv4 address 198.51.100.11/24
set security screen white-list wlipv4 address 198.51.100.12/24
set security screen white-list wlipv4 address 198.51.100.13/24
set security screen white-list wlipv6 address 2001:db8::1/32
set security screen white-list wlipv6 address 2001:db8::2/32
set security screen white-list wlipv6 address [2001:db8::3/32]
set security screen white-list wlipv6 address [2001:db8::4/32]
set security screen ids-options jscreen udp flood white-list wlipv4
set security screen ids-options jscreen udp flood white-list wlipv6
```


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 the allowlists:

1. Specify the name of the allowlist and the IPv4 addresses to bypass UDP flood detection.

```
[edit security screen]
user@host# set white-list wlipv4 address 198.51.100.10/32
user@host# set white-list wlipv4 address 198.51.100.11/32
user@host# set white-list wlipv4 address 198.51.100.12/32
user@host# set white-list wlipv4 address 198.51.100.13/32
```

2. Specify the name of the allowlist and the IPv6 addresses to bypass UDP flood detection.

```
[edit security screen]
user@host# set white-list wlipv6 address 2001:db8::1/32
user@host# set white-list wlipv6 address 2001:db8::2/32
user@host# set white-list wlipv6 address 2001:db8::3/32
user@host# set white-list wlipv6 address 2001:db8::4/32
```

3. Set the UDP flood allowlist option.

```
[edit security screen]
user@host# set ids-option jscreen udp flood white-list wlipv4
user@host# set ids-option jscreen udp flood white-list wlipv6
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` 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 screen
ids-option jscreen {
    udp {
```

```

        flood {
            white-list [ wlipv4 wlipv6 ];
        }
    }
}
white-list wlipv4 {
    address [ 198.51.100.11/32 198.51.100.12/32 198.51.100.13/32 198.51.100.14/32 ];
}
white-list wlipv6 {
    address [ 2001:db8::1/32 2001:db8::2/32 2001:db8::3/32 2001:db8::4/32 ];
}

```

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

Verification

IN THIS SECTION

- [Verifying Whitelist Configuration | 87](#)

Verifying Whitelist Configuration

Purpose

Verify that the allowlist is configured properly.

Action

From operational mode, enter the `show security screen white-list wlipv4` and `show security screen ids-option jscreen` command.

```
user@host> show security screen white-list wlipv4
```

```
Screen white list:
```

```

198.51.100.10/32
198.51.100.11/32
198.51.100.12/32
198.51.100.13/32

```

```
user@host> show security screen ids-option jscreen
```

Name	Value
.....	
UDP flood threshold	##
UDP flood white-list	wlipv4
UDP flood white-list	wlipv6

Understanding Allowlist for All Screen Options

IN THIS SECTION

- [Benefits | 88](#)

Junos OS provides the administrative option to configure allowlist for all IP screen options in a security zone. When screen is enabled in a security zone, all IP packets exceeding the threshold value are dropped. Some of these packets are valid from specific sources and should not be dropped from the traffic. When you configure allowlist at a zone level, all the IP addresses from the specific sources are allowed to bypass the attack detection check.

This feature is needed when all IP addresses from a specific source should bypass the attack detection check.

Both IPv4 and IPv6 allowlists are supported. Addresses in a allowlist should be all IPv4 or all IPv6. In each allowlist, there can be up to 32 IP address prefixes. You can specify multiple addresses or address prefixes as a sequence of addresses separated by spaces and enclosed in square brackets. You can configure a single address or a subnet address.

Benefits

- Global IP allowlist bypasses the IP packet screening check to allow all the IP packets from specific sources.

Understanding SYN Cookie Protection

IN THIS SECTION

- [SYN Cookie Options | 90](#)

SYN cookie is a stateless SYN proxy mechanism you can use in conjunction with other defenses against a SYN flood attack.

As with traditional SYN proxying, SYN cookie is activated when the SYN flood attack threshold is exceeded. However, because SYN cookie is stateless, it does not set up a session or policy and route lookups upon receipt of a SYN segment, and it maintains no connection request queues. This dramatically reduces CPU and memory usage and is the primary advantage of using SYN cookie over the traditional SYN proxying mechanism.

When SYN cookie is enabled on Junos OS and becomes the TCP-negotiating proxy for the destination server, it replies to each incoming SYN segment with a SYN/ACK containing an encrypted cookie as its initial sequence number (ISN). The cookie is an MD5 hash of the original source address, destination address, and ISN from the original SYN packet. After sending the cookie, Junos OS drops the original SYN packet and deletes the calculated cookie from memory. If there is no response to the packet containing the cookie, the attack is noted as an active SYN attack and is effectively stopped.

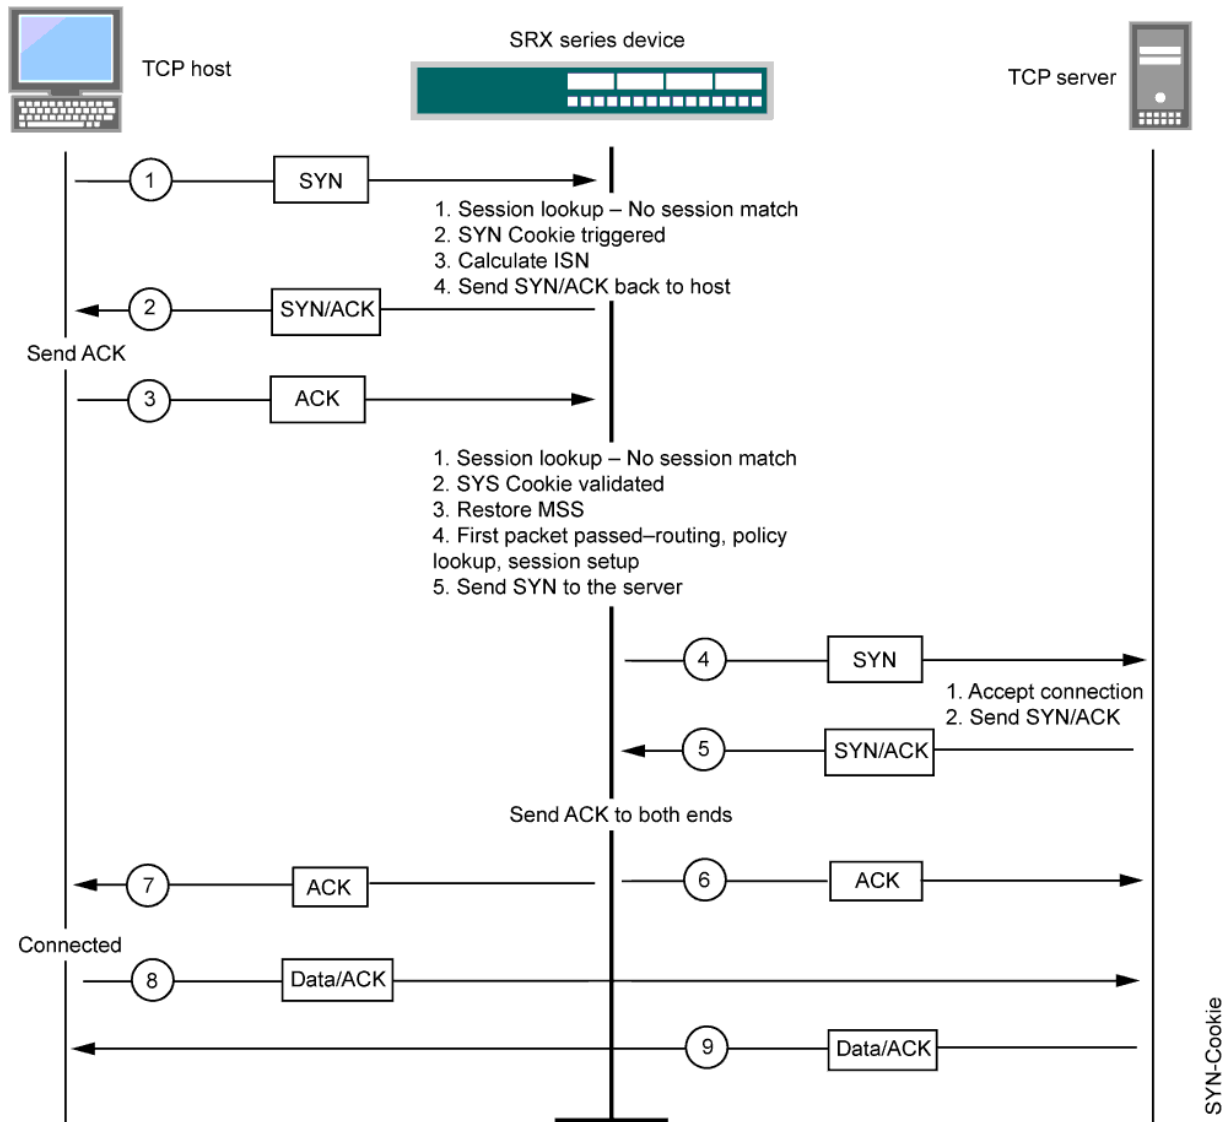
If the initiating host responds with a TCP packet containing the cookie +1 in the TCP ACK field, Junos OS extracts the cookie, subtracts 1 from the value, and recomputes the cookie to validate that it is a legitimate ACK. If it is legitimate, Junos OS starts the TCP proxy process by setting up a session and sending a SYN to the server containing the source information from the original SYN. When Junos OS receives a SYN/ACK from the server, it sends ACKs to the server and to the initiation host. At this point the connection is established and the host and server are able to communicate directly.



NOTE: The use of SYN cookie or SYN proxy enables the SRX Series Firewall to protect the TCP servers behind it from SYN flood attacks in IPv6 flows.

[Figure 6 on page 90](#) shows how a connection is established between an initiating host and a server when SYN cookie is active on Junos OS.

Figure 6: Establishing a Connection with SYN Cookie Active



SYN Cookie Options

You can set the following parameters for incomplete TCP proxy connection requests:

- **Attack Threshold**—This option allows you to set the number of SYN segments (that is, TCP segments with the SYN flag set) to the same destination address per second required to activate the SYN proxy mechanism. Although you can set the threshold to any number, you need to know the normal traffic patterns at your site to set an appropriate threshold for it. For example, for an e-business site that normally gets 2000 SYN segments per second, you might want to set the threshold to 30,000 SYN segments per second. The valid threshold range is 1 to 1,000,000. For a smaller site that normally

gets 20 SYN segments per second, you might consider setting the threshold to 40 SYN segments per second.

- **Alarm Threshold**—This option allows you to set the number of proxied, half-complete TCP connection requests per second after which Junos OS enters an alarm in the event log. The alarm threshold value you set triggers an alarm when the number of proxied, half-completed connection requests to the same destination address per second exceeds that value. For example, if you set the SYN attack threshold at 2000 SYN segments per second and the alarm at 1000, then a total of 3001 SYN segments to the same destination address per second is required to trigger an alarm entry in the log. The valid threshold range is 1 to 1,000,000 and the default alarm threshold value is 512.
- **Source Threshold**—This option allows you to specify the number of SYN segments received per second from a single source IP address—regardless of the destination IP address—before Junos OS begins dropping connection requests from that source.

When you set a SYN attack threshold and a source threshold, you put both the basic SYN flood protection mechanism and the source-based SYN flood tracking mechanism in effect. The valid threshold range is 4 to 1,000,000 and the default alarm threshold value is 4000.

- **Destination Threshold**—This option allows you to specify the number of SYN segments received per second for a single destination IP address before Junos OS begins dropping connection requests to that destination. If a protected host runs multiple services, you might want to set a threshold based on destination IP address only—regardless of the destination port number. The valid threshold range is 4 to 1,000,000 and the default alarm threshold value is 4000.

When you set a SYN attack threshold and a destination threshold, you put both the basic SYN flood protection mechanism and the destination-based SYN flood tracking mechanism in effect.

- **Timeout**—This option allows you to set the maximum length of time before a half-completed connection is dropped from the queue. The default is 20 seconds, and you can set the timeout from 0 to 50 seconds. You might try decreasing the timeout value to a shorter length until you begin to see dropped connections during normal traffic conditions.

When either a source or destination threshold is not configured, the system will use the default threshold value. The default source and destination threshold value is 4000 pps.

Detecting and Protecting Your Network Against SYN Flood Attacks by Enabling SYN Cookie Protection

IN THIS SECTION

- [Requirements | 92](#)
- [Overview | 92](#)
- [Configuration | 92](#)
- [Verification | 94](#)

This example shows how to detect and protect your network against SYN flood attacks by enabling the SYN cookie protection.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you set the external-syn-flood timeout value to 20 and set the security zone for external screen to external-syn-flood. Also, you set the protection mode to syn-cookie.



NOTE: The SYN cookie feature can detect and protect only against spoofed SYN flood attacks, minimizing the negative impact on hosts that are secured by Junos OS. If an attacker uses a legitimate IP source address, rather than a spoofed IP source, then the SYN cookie mechanism does not stop the attack.

Configuration

IN THIS SECTION

- [Procedure | 93](#)

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 screen ids-option external-syn-flood tcp syn-flood timeout 20
set security zones security-zone external screen external-syn-flood
set security flow syn-flood-protection-mode syn-cookie
```

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 enable the SYN cookie protection:

1. Specify the external-syn-flood timeout value.

```
[edit]
user@host# set security screen ids-option external-syn-flood tcp syn-flood timeout 20
```

2. Set the security-zone for external screen.

```
[edit]
user@host# set security zones security-zone external screen external-syn-flood
```

3. Set the protection mode.

```
[edit]
user@host# set security flow syn-flood-protection-mode syn-cookie
```


Results

From configuration mode, confirm your configuration by entering the `show security screen`, `show security zones`, and `show security flow` 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 screen
  ids-option external-syn-flood {
    tcp {
      syn-flood {
        timeout 20;
      }
    }
  }
```

```
[edit]
user@host# show security zones
security-zone zone {
  screen external-syn-flood;
}
[edit]
user@host# show security flow
  syn-flood-protection-mode syn-cookie;
```

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

Verification

IN THIS SECTION

- [Verifying SYN Cookie Protection | 95](#)

Verifying SYN Cookie Protection

Purpose

Verifying SYN cookie protection.

Action

Enter the `show security screen ids-option external-syn-flood` and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option external-syn-flood
node0:
-----
Screen object status:
Name                                     Value
TCP SYN flood attack threshold          200
TCP SYN flood alarm threshold            512
TCP SYN flood source threshold           4000
TCP SYN flood destination threshold       4000
TCP SYN flood timeout                    20

user@host> show security zones
Security zone: external
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: external-syn-flood
  Interfaces bound: 0
  Interfaces:
```

Meaning

The sample output shows that SYN cookie protection is enabled with a source and destination threshold.

Understanding ICMP Flood Attacks

An ICMP flood typically occurs when ICMP echo requests overload the target of the attack with so many requests that the target expends all its resources responding until it can no longer process valid network traffic.



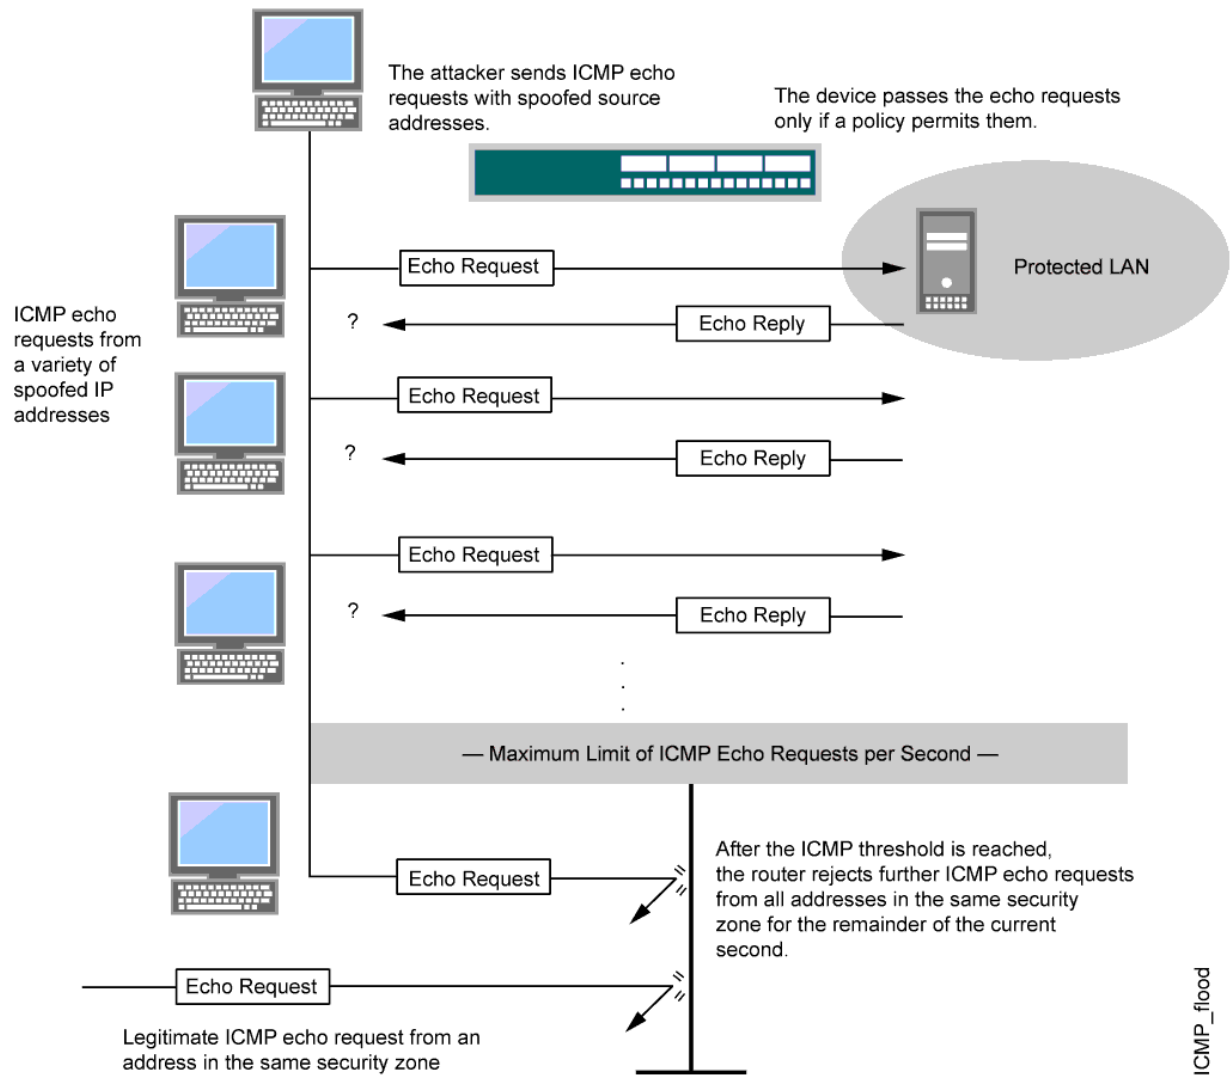
NOTE: ICMP messages generated in flow mode are limited to 12 messages every 10 seconds. This rate limit is calculated on a per-CPU basis. Once the threshold is reached, no further acknowledgement messages are sent to the device.

When enabling the ICMP flood protection feature, you can set a threshold that, once exceeded, invokes the ICMP flood attack protection feature. (The default threshold value is 1000 packets per second.) If the threshold is exceeded, Junos OS ignores further ICMP echo requests for the remainder of that second plus the next second as well. See [Figure 7 on page 97](#).



NOTE: An ICMP flood can consist of any type of ICMP message. Therefore, Junos OS monitors all ICMP message types, not just echo requests.

Figure 7: ICMP Flooding



NOTE: Junos OS supports ICMP flood protection for both IPv4 and IPv6 traffic.

Protecting Your Network Against ICMP Flood Attacks by Enabling ICMP Flood Protection

IN THIS SECTION

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

This example shows how to protect your network against ICMP flood attacks by enabling ICMP flood protection.

Requirements

No special configuration beyond device initialization is required before enabling ICMP flood protection.

Overview

In this example, you enable ICMP flood protection. The value unit is ICMP packets per second, or pps. The default value is 1000 pps. You specify the zone where a flood might originate.

Configuration

IN THIS SECTION

- [Procedure | 99](#)

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 screen ids-option 1000-icmp-flood icmp flood threshold 1000
set security zones security-zone zone screen 1000-icmp-flood
```

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 enable ICMP flood protection:

1. Specify the ICMP flood threshold value.

```
[edit]
user@host# set security screen ids-option 1000-icmp-flood icmp flood threshold 1000
```

2. Set the security zone for zone screen.

```
[edit]
user@host# set security zones security-zone zone screen 1000-icmp-flood
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zones` 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 screen
  ids-option 1000-icmp-flood {
    icmp {
      flood threshold 1000;
```

```
    }
}
```

```
[edit]
user@host# show security zones
    security-zone zone {
        screen 1000-icmp-flood;
    }
```

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

Verification

IN THIS SECTION

- [Verifying ICMP Flood Protection | 100](#)

Verifying ICMP Flood Protection

Purpose

Verify ICMP flood protection

Action

Enter the `show security screen ids-option 1000-icmp-flood` and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option 1000-icmp-flood
node0:
-----
Screen object status:
Name                               Value
  ICMP flood threshold             1000

user@host> show security zones
Security zone: zone
```

```
Send reset for non-SYN session TCP packets: Off
Policy configurable: Yes
Screen: 1000-icmp-flood
Interfaces bound: 0
Interfaces:
```

Meaning

The sample output shows that ICMP flood protection is enabled and threshold is set.

Understanding UDP Flood Attacks

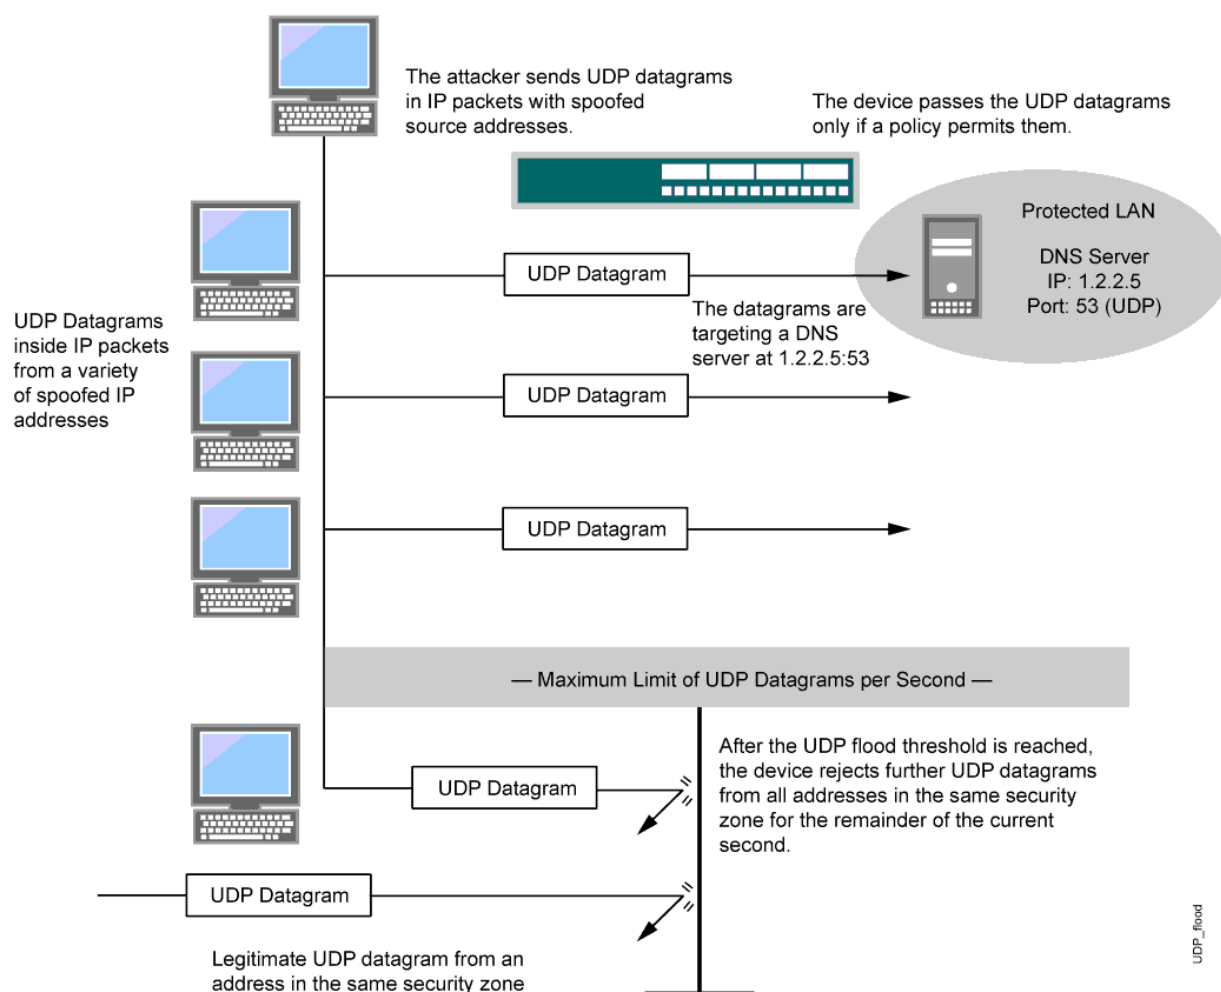
Similar to an ICMP flood, a UDP flood occurs when an attacker sends IP packets containing UDP datagrams with the purpose of slowing down the victim to the point that the victim can no longer handle valid connections.

After enabling the UDP flood protection feature, you can set a threshold that, once exceeded, invokes the UDP flood attack protection feature. (The default threshold value is 1000 packets per second, or pps.) If the number of UDP datagrams from one or more sources to a single destination exceeds this threshold, Junos OS ignores further UDP datagrams to that destination for the remainder of that second plus the next second as well. See [Figure 8 on page 102](#).



NOTE: The SRX5400, SRX5600, and SRX5800 devices do not drop the packet in the next second.

Figure 8: UDP Flooding



NOTE: Junos OS supports UDP flood protection for IPV4 and IPV6 packets.

Protecting Your Network Against UDP Flood Attacks by Enabling UDP Flood Protection

IN THIS SECTION

● Requirements | 103

- [Overview | 103](#)
- [Configuration | 103](#)
- [Verification | 105](#)

This example shows how to protect your network against UDP flood attacks by enabling UDP flood protection.

Requirements

No special configuration beyond device initialization is required before enabling UDP flood protection.

Overview

In this example, you enable UDP flood protection. The value unit is UDP packets per second, or pps. The default value is 1000 pps. You specify the zone where a flood might originate.

Configuration

IN THIS SECTION

- [Procedure | 103](#)

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 screen ids-option 1000-udp-flood udp flood threshold 1000
set security zones security-zone external screen 1000-udp-flood
```

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 *drop-profiles*. To enable UDP flood protection:

1. Specify the UDP flood threshold value.

```
[edit]
user@host# set security screen ids-option 1000-udp-flood udp flood threshold 1000
```

2. Set the security zone for external screen.

```
[edit]
user@host# set security zones security-zone external screen 1000-udp-flood
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zones` 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 screen
  ids-option 1000-udp-flood {
    udp {
      flood threshold 1000;
    }
  }
```

```
[edit]
user@host# show security zones
  security-zone external {
    screen 1000-udp-flood;
  }
```

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

Verification

IN THIS SECTION

- [Verifying UDP Flood Protection | 105](#)

Verifying UDP Flood Protection

Purpose

Verify UDP flood protection.

Action

Enter the `show security screen ids-option 1000-udp-flood` and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option 1000-udp-flood
node0:
-----
Screen object status:
Name                               Value
UDP flood threshold                1000

user@host> show security zones
Security zone: external
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 1000-udp-flood
  Interfaces bound: 0
  Interfaces:
```

Meaning


The sample output shows that UDP flood protection is enabled and threshold is set.

Understanding Land Attacks

Combining a SYN attack with IP spoofing, a land attack occurs when an attacker sends spoofed SYN packets containing the IP address of the victim as both the destination and the source IP address.

The receiving system responds by sending the SYN-ACK packet to itself, creating an empty connection that lasts until the idle timeout value is reached. Flooding a system with such empty connections can overwhelm the system, causing a denial of service (DoS). See [Figure 9 on page 106](#).

Figure 9: Land Attack



When you enable the screen option to block land attacks, Junos OS combines elements of the SYN flood defense and IP spoofing protection to detect and block any attempts of this nature.



NOTE: Junos OS supports land attack protection for both IPv4 and IPv6 packets.

Protecting Your Network Against Land Attacks by Enabling Land Attack Protection

IN THIS SECTION

- [Requirements | 107](#)
- [Overview | 107](#)
- [Configuration | 107](#)
- [Verification | 109](#)

This example shows how to protect your network against attacks by enabling land attack protection.

Requirements

No special configuration beyond device initialization is required before enabling land attack protection.

Overview

This example shows how to enable protection against a land attack. In this example, you set the security screen object name as land and set the security zone as zone.

Configuration

IN THIS SECTION

- [Procedure | 108](#)

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 screen ids-option land tcp land
set security zones security-zone zone screen land
```

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 enable protection against a land attack:

1. Specify the screen object name.

```
[edit]
user@host# set security screen ids-option land tcp land
```

2. Set the security zone.

```
[edit]
user@host# set security zones security-zone zone screen land
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` and `show security zones` 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 screen
  ids-option land {
    tcp {
      land;
```

```
    }
}
```

```
[edit]
user@host# show security zones
    security-zone zone {
        screen land;
    }
```

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

Verification

IN THIS SECTION

- [Verifying Protection Against a Land Attack | 109](#)

Verifying Protection Against a Land Attack

Purpose

Verify protection against a land attack.

Action

Enter the `show security screen ids-option land` and `show security zones` commands from operational mode.

```
user@host> show security screen ids-option land
node0:
-----
Screen object status:
Name                               Value
TCP land attack                   enabled

user@host> show security zones
Security zone: zone
Send reset for non-SYN session TCP packets: Off
```



```
Policy configurable: Yes
Screen: land
Interfaces bound: 0
Interfaces:
```

Meaning

The sample output shows that protection against a land attack is enabled.

RELATED DOCUMENTATION

| [DoS Attack Overview](#) | 47

OS-Specific DoS Attack

IN THIS SECTION

- [OS-Specific DoS Attacks Overview](#) | 111
- [Understanding Ping of Death Attacks](#) | 111
- [Example: Protecting Against a Ping of Death Attack](#) | 112
- [Understanding Teardrop Attacks](#) | 114
- [Understanding WinNuke Attacks](#) | 115
- [Example: Protecting Against a WinNuke Attack](#) | 116

OS-specific DoS attack focuses on one-packet or two-packet kills. These attacks include the Ping of Death attack, the Teardrop attack, and the WinNuke attack. The Junos OS has the capability to mitigate these attacks. For more information, see the following topics:

OS-Specific DoS Attacks Overview

If an attacker not only identifies the IP address and responsive port numbers of an active host but also its operating system (OS), instead of resorting to brute-force attacks, the attacker can launch more elegant attacks that can produce one-packet or two-packet “kills.”

OS-specific denial-of-service (DoS) attacks, including ping of death attacks, teardrop attacks, and WinNuke attacks, can cripple a system with minimal effort. If Junos OS is protecting hosts susceptible to these attacks, you can configure Junos OS to detect these attacks and block them before they reach their target.

Understanding Ping of Death Attacks

OS-specific DoS attacks, such as ping of death attacks, can cripple a system with minimal effort.

The maximum allowable IP packet size is 65,535 bytes, including the packet header, which is typically 20 bytes. An ICMP echo request is an IP packet with a pseudo header, which is 8 bytes. Therefore, the maximum allowable size of the data area of an ICMP echo request is 65,507 bytes ($65,535 - 20 - 8 = 65,507$).


However, many ping implementations allow the user to specify a packet size larger than 65,507 bytes. A grossly oversized ICMP packet can trigger a range of adverse system reactions such as denial of service (DoS), crashing, freezing, and rebooting.

When you enable the ping of death screen option, Junos OS detects and rejects such oversized and irregular packet sizes even when the attacker hides the total packet size by fragmenting it. See [Figure 10 on page 112](#).



NOTE: For information about IP specifications, see RFC 791, *Internet Protocol*. For information about ICMP specifications, see RFC 792, *Internet Control Message Protocol*. For information about ping of death attacks, see <http://www.insecure.org/spl0its/ping-of-death.html>.

Figure 10: Ping of Death



The size of this packet is 65,538 bytes. It exceeds the size limit prescribed by RFC 791, *Internet Protocol*, which is 65,535 bytes. As the packet is transmitted, it becomes broken into numerous fragments. The reassembly process might cause the receiving system to crash.



NOTE: Junos OS supports ping of death protection for both IPv4 and IPv6 packets.

Example: Protecting Against a Ping of Death Attack

IN THIS SECTION

- [Requirements | 112](#)
- [Overview | 112](#)
- [Configuration | 113](#)
- [Verification | 113](#)

This example shows how to protect against a ping-of-death attack.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable protection against a ping-of-death attack and specify the zone where the attack originates.

Configuration

IN THIS SECTION

- Procedure | [113](#)

Procedure

Step-by-Step Procedure

To enable protection against a ping of death:

1. Specify the screen object name.

```
[edit]  
user@host# set security screen ids-option ping-death icmp ping-death
```

2. Set the security zone for zone screen.

```
[edit]  
user@host# set security zones security-zone zone screen ping-death
```

3. If you are done configuring the device, commit the configuration.

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

Verification

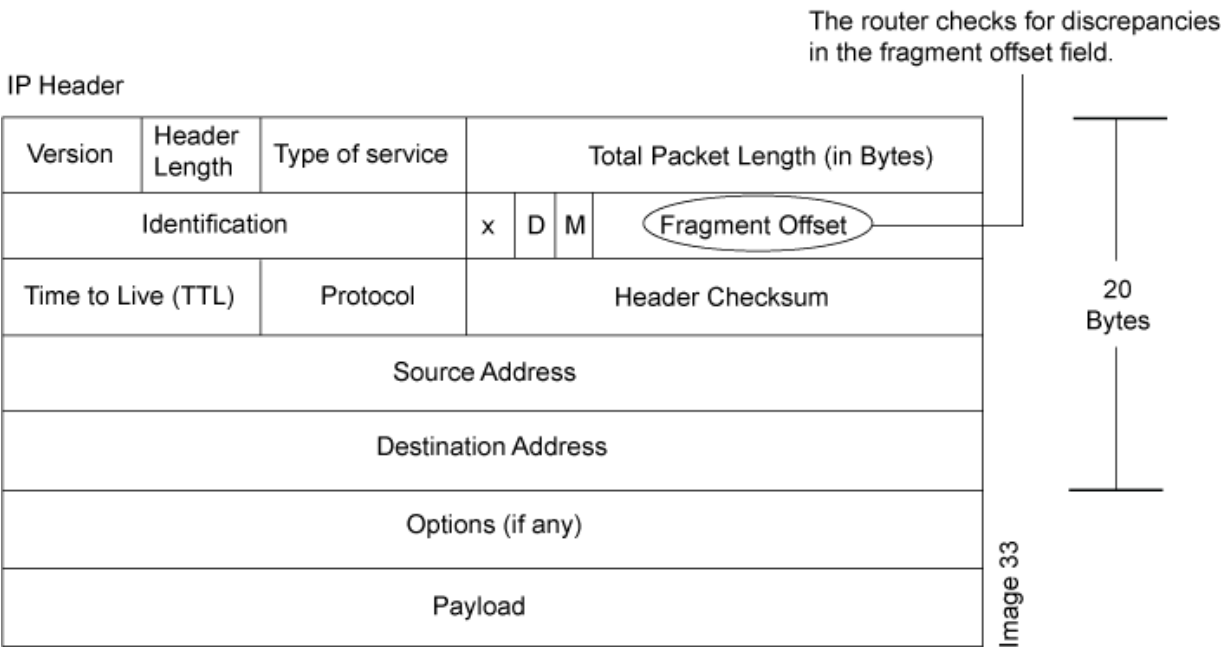
To verify the configuration is working properly, enter the `show security screen ids-option ping-death` and `show security zones` commands in operational mode.

Understanding Teardrop Attacks

OS-specific denial-of-service (DoS) attacks, such as teardrop attacks, can cripple a system with minimal effort.

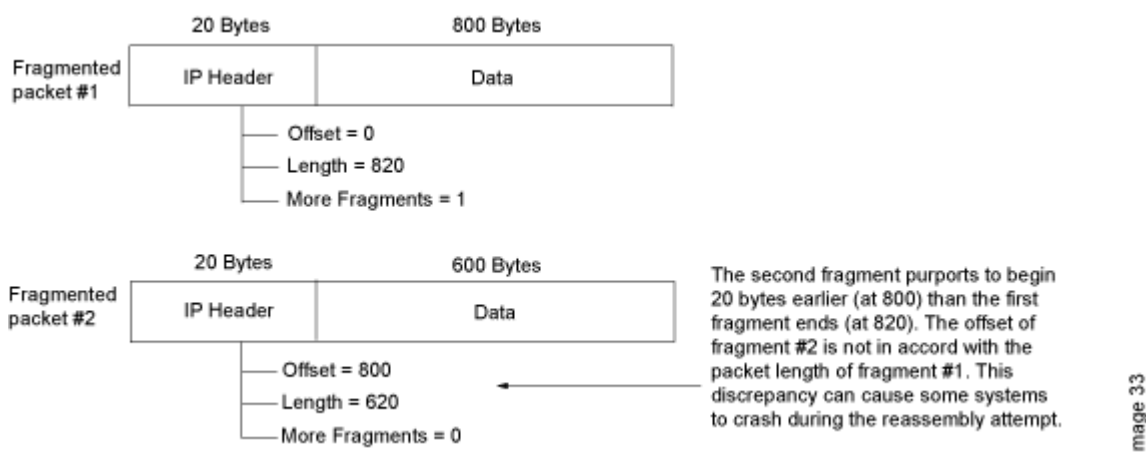
Teardrop attacks exploit the reassembly of fragmented IP packets. In the IP header, one of the fields is the fragment offset field, which indicates the starting position, or offset, of the data contained in a fragmented packet relative to the data of the original unfragmented packet. See [Figure 11 on page 114](#).

Figure 11: Teardrop Attacks



When the sum of the offset and size of one fragmented packet differ from that of the next fragmented packet, the packets overlap, and the server attempting to reassemble the packet can crash, especially if it is running an older OS that has this vulnerability. See [Figure 12 on page 115](#).

Figure 12: Fragment Discrepancy



After you enable the teardrop attack screen option, whenever Junos OS detects this discrepancy in a fragmented packet, it drops it.



NOTE: Junos OS supports teardrop attack prevention for both IPv4 and IPv6 packets.


Understanding WinNuke Attacks

OS-specific denial-of-service (DoS) attacks, such as WinNuke attacks, can cripple a system with minimal effort.

WinNuke is a DoS attack targeting any computer on the Internet running Windows. The attacker sends a TCP segment—usually to NetBIOS port 139 with the urgent (URG) flag set—to a host with an established connection (see [Figure 13 on page 116](#)). This introduces a NetBIOS fragment overlap, which causes many machines running Windows to crash. After the attacked machine is rebooted, the following message appears, indicating that an attack has occurred:

```
An exception OE has occurred at 0028:[address] in VxD MSTCP(01) +
000041AE. This was called from 0028:[address] in VxD NDIS(01) +
00008660. It may be possible to continue normally.
Press any key to attempt to continue.
Press CTRL+ALT+DEL to restart your computer. You will lose any unsaved information in all
applications.
Press any key to continue.
```

Figure 13: WinNuke Attack Indicators



If you enable the WinNuke attack defense screen option, Junos OS scans any incoming Microsoft NetBIOS session service (port 139) packets. If Junos OS observes that the URG flag is set in one of those packets, it unsets the URG flag, clears the URG pointer, forwards the modified packet, and makes an entry in the event log noting that it has blocked an attempted WinNuke attack.



NOTE: Junos OS supports WinNuke attack protection for both IPv4 and IPv6 traffic.

Example: Protecting Against a WinNuke Attack

IN THIS SECTION

- [Requirements | 117](#)
- [Overview | 117](#)
- [Configuration | 117](#)

- [Verification | 118](#)

This example shows how to protect against a WinNuke attack.

Requirements

No special configuration beyond device initialization is required before configuring this feature.

Overview

In this example, you enable protection against a WinNuke attack and specify the zone where the attack originates.

Configuration

IN THIS SECTION

- [Procedure | 117](#)

Procedure

Step-by-Step Procedure

To enable protection against WinNuke attack:

1. Specify the screen name.

```
[edit]
user@host# set security screen ids-option winnuke tcp winnuke
```

2. Associate the screen with a security zone.

```
[edit]
user@host# set security zones security-zone zone screen winnuke
```


3. If you are done configuring the device, commit the configuration.

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

Verification

To verify the configuration is working properly, enter the `show security screen ids-option winnuke` and `show security zones` commands in operational mode.

RELATED DOCUMENTATION

[DoS Attack Overview](#) | 47

[Firewall DoS Attacks](#) | 49

3

CHAPTER

Suspicious Packets

[Suspicious Packet Attributes Overview](#) | 120

[ICMP and SYN Fragment Attacks](#) | 120

[IP Packet Protection](#) | 130

Suspicious Packet Attributes Overview

Attackers can craft packets to perform reconnaissance or launch denial-of-service (DoS) attacks. Sometimes it is unclear what the intent of a crafted packet is, but the very fact that it is crafted suggests that it is being put to some kind of insidious use.

The following topics describe screen options that block suspicious packets that might contain hidden threats:

- ["Understanding ICMP Fragment Protection" on page 121](#)
- ["Understanding Large ICMP Packet Protection" on page 124](#)
- ["Understanding Bad IP Option Protection" on page 134](#)
- ["Understanding Unknown Protocol Protection" on page 137](#)
- ["Understanding IP Packet Fragment Protection" on page 130](#)
- ["Understanding SYN Fragment Protection" on page 127](#)

ICMP and SYN Fragment Attacks

IN THIS SECTION

- [Understanding ICMP Fragment Protection | 121](#)
- [Example: Blocking Fragmented ICMP Packets | 122](#)
- [Understanding Large ICMP Packet Protection | 124](#)
- [Example: Blocking Large ICMP Packets | 125](#)
- [Understanding SYN Fragment Protection | 127](#)
- [Example: Dropping IP Packets Containing SYN Fragments | 128](#)

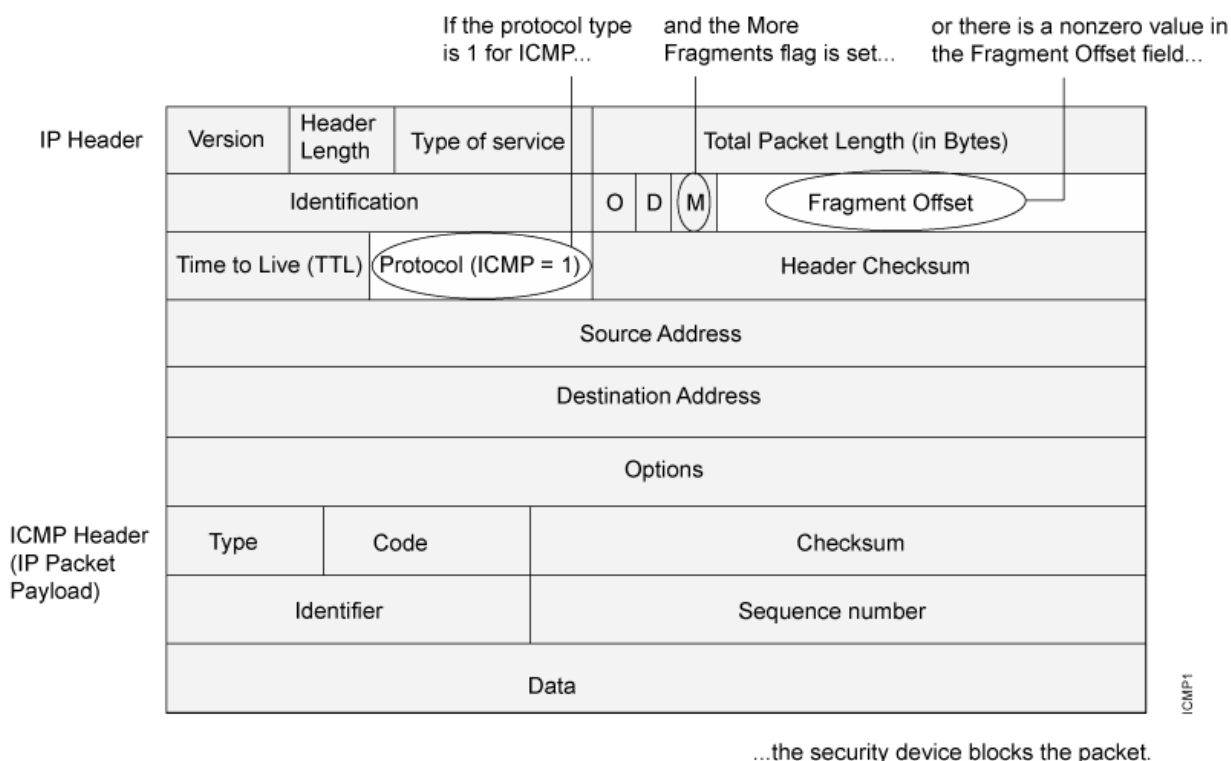
An ICMP flood typically occurs when ICMP echo request messages overload the victim, causing resources to stop responding to valid traffic. A fragmented SYN packet is anomalous, and as such, it is suspect. When a victim receives these packets, the results can range from processing packets incorrectly to crashing the entire system. For more information, see the following topics:

Understanding ICMP Fragment Protection

Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate reason for ICMP packets to be fragmented. If an ICMP packet is so large that it must be fragmented, something is amiss.

When you enable the ICMP fragment protection screen option, Junos OS blocks any ICMP packet that has the More Fragments flag set or that has an offset value indicated in the offset field. See [Figure 14 on page 121](#).

Figure 14: Blocking ICMP Fragments



NOTE: Junos OS supports ICMP fragment protection for ICMPv6 packets.

Example: Blocking Fragmented ICMP Packets

IN THIS SECTION

- Requirements | 122
- Overview | 122
- Configuration | 123
- Verification | 123

This example shows how to block fragmented ICMP packets.

Requirements

Before you begin, Understand ICMP fragment protection. See ["Suspicious Packet Attributes Overview" on page 120](#).

Overview

IN THIS SECTION

- Topology | 122

When you enable the ICMP fragment protection screen option, Junos OS blocks any ICMP packet that has the more fragments flag set or that has an offset value indicated in the offset field.

In this example, you configure the ICMP fragment screen to block fragmented ICMP packets originating from the zone1 security zone.

Topology

Configuration

IN THIS SECTION

- Procedure | [123](#)

Procedure

Step-by-Step Procedure

To block fragmented ICMP packets:

1. Configure the screen.

```
[edit]  
user@host# set security screen ids-option icmp-fragment icmp fragment
```

2. Configure a security zone.

```
[edit]  
user@host# set security zones security-zone zone1 screen icmp-fragment
```

3. If you are done configuring the device, commit the configuration.

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

Verification

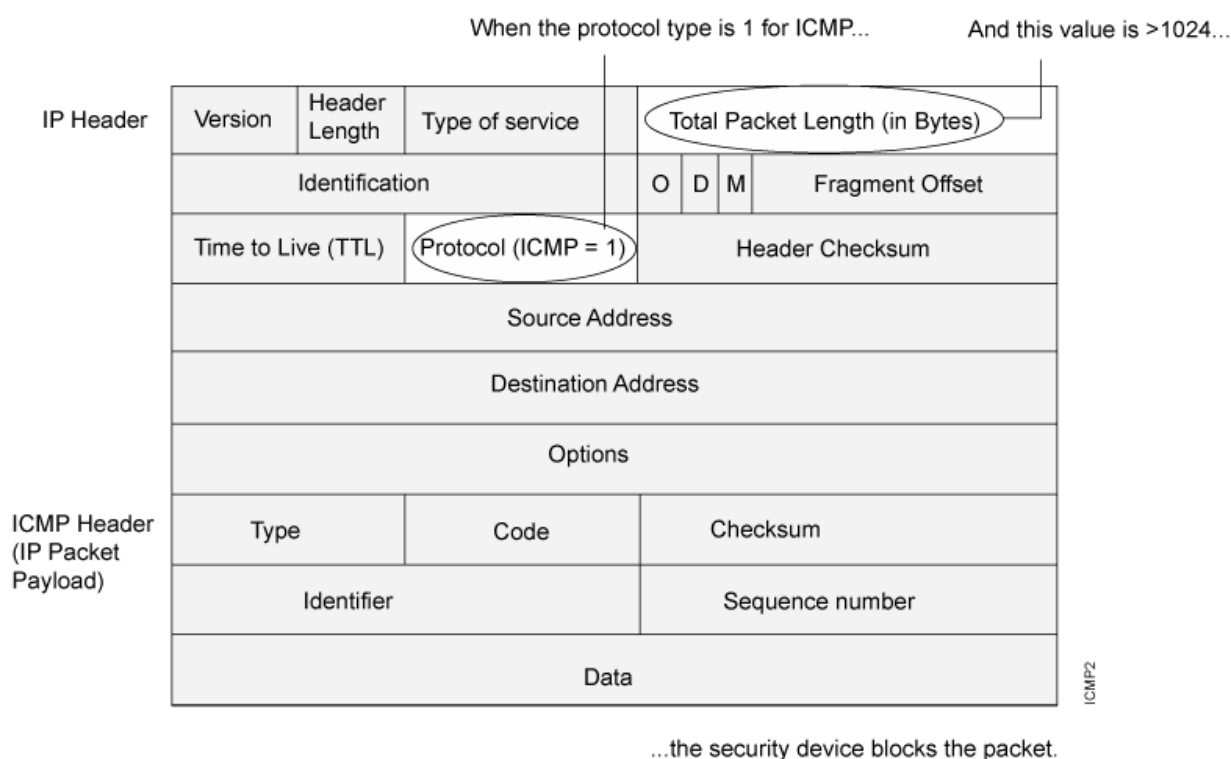
To verify the configuration is working properly, enter the `show security screen statistics zone zone-name` command.

Understanding Large ICMP Packet Protection

Internet Control Message Protocol (ICMP) provides error reporting and network probe capabilities. Because ICMP packets contain very short messages, there is no legitimate reason for large ICMP packets. If an ICMP packet is unusually large, something is amiss.

See [Figure 15 on page 124](#).

Figure 15: Blocking Large ICMP Packets



When you enable the large size ICMP packet protection screen option, Junos OS drops ICMP packets with a length greater than 1024 bytes.



NOTE: Junos OS supports large ICMP packet protection for both ICMP and ICMPv6 packets.

Example: Blocking Large ICMP Packets

IN THIS SECTION

- Requirements | 125
- Overview | 125
- Configuration | 126
- Verification | 126

This example shows how to block large ICMP packets.

Requirements

Before you begin, Understand large ICMP packet protection. See ["Suspicious Packet Attributes Overview" on page 120](#).

Overview

IN THIS SECTION

- Topology | 125

When you enable the large ICMP packet protection screen option, Junos OS drops ICMP packets that are larger than 1024 bytes.

In this example, you configure the ICMP large screen to block large ICMP packets originating from the zone1 security zone.

Topology

Configuration

IN THIS SECTION

- Procedure | 126

Procedure

Step-by-Step Procedure

To block large ICMP packets:

1. Configure the screen.

```
[edit]  
user@host# set security screen ids-option icmp-large icmp large
```

2. Configure a security zone.

```
[edit]  
user@host# set security zones security-zone zone1 screen icmp-large
```

3. If you are done configuring the device, commit the configuration.

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

Verification

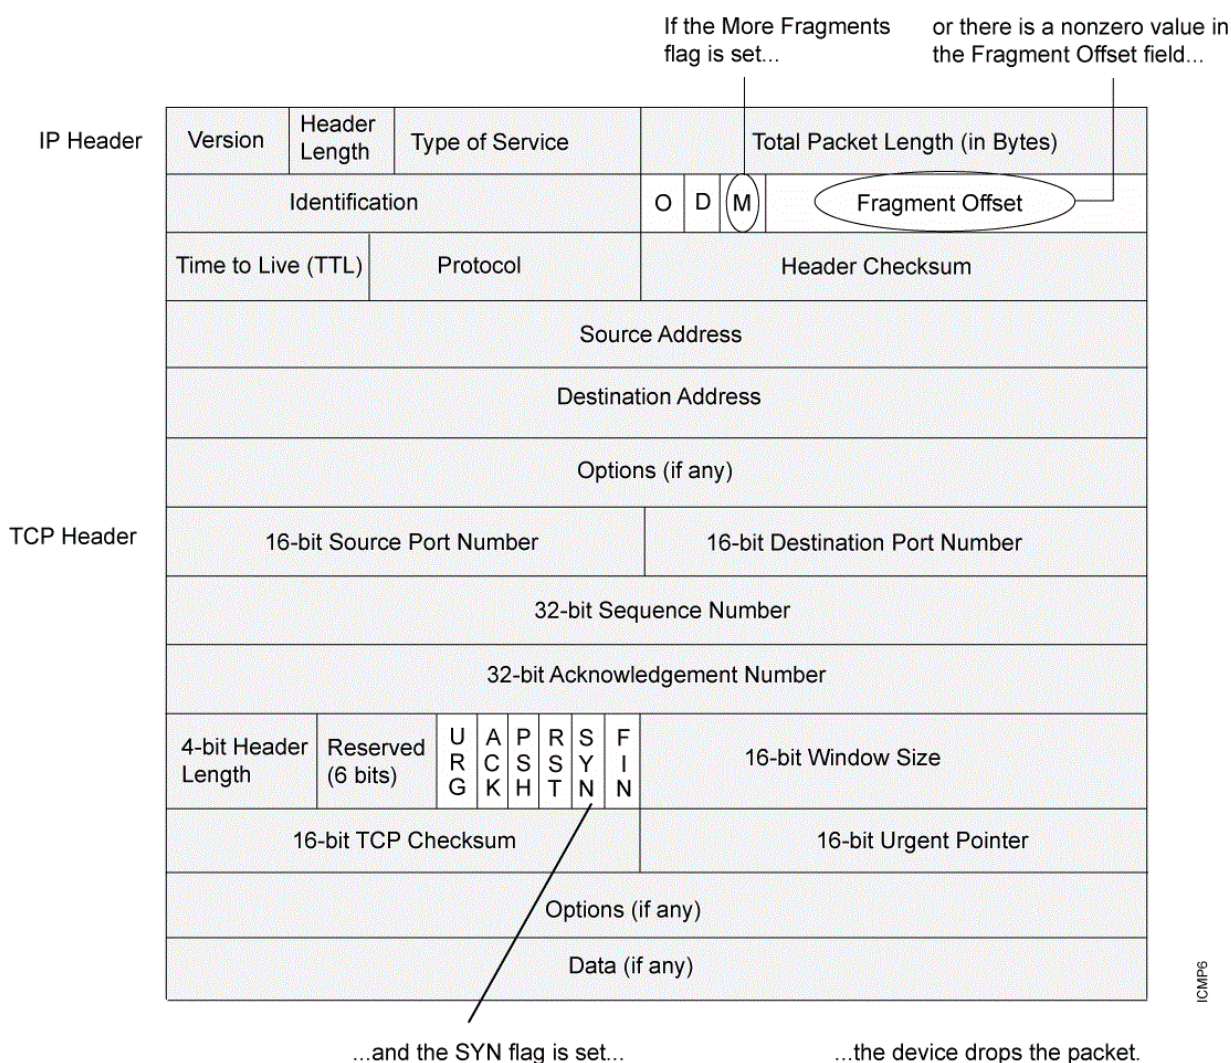
To verify the configuration is working properly, enter the `show security screen statistics zone zone-name` command.

Understanding SYN Fragment Protection

The IP encapsulates a TCP SYN segment in the IP packet that initiates a TCP connection. Because the purpose of this packet is to initiate a connection and invoke a SYN/ACK segment in response, the SYN segment typically does not contain any data. Because the IP packet is small, there is no legitimate reason for it to be fragmented.

A fragmented SYN packet is anomalous, and, as such, it is suspect. To be cautious, block such unknown elements from entering your protected network. See [Figure 16 on page 127](#).

Figure 16: SYN Fragments



When you enable the SYN fragment detection screen option, Junos OS detects packets when the IP header indicates that the packet has been fragmented and the SYN flag is set in the TCP header. Junos OS records the event in the screen counters list for the ingress interface.



NOTE: Junos OS supports SYN fragment protection for both IPv4 and IPv6 packets.

Example: Dropping IP Packets Containing SYN Fragments

IN THIS SECTION

- Requirements | 128
- Overview | 128
- Configuration | 129
- Verification | 129

This example shows how to drop IP packets containing SYN fragments.

Requirements

Before you begin, Understand IP packet fragment protection. See "[Suspicious Packet Attributes Overview](#)" on page 120.

Overview

IN THIS SECTION

- Topology | 129

When you enable the SYN fragment detection screen option, Junos OS detects packets when the IP header indicates that the packet has been fragmented and the SYN flag is set in the TCP header. Also, Junos OS records the event in the screen counters list for the ingress interface.

In this example, you configure the SYN fragment screen to drop fragmented SYN packets originating from the zone1 security zone.

Topology

Configuration

IN THIS SECTION

- [Procedure | 129](#)

Procedure

Step-by-Step Procedure

To drop IP packets containing SYN fragments:

1. Configure the screen.

```
[edit]
user@host# set security screen ids-option syn-frag tcp syn-frag
```

2. Configure the security zone.

```
[edit]
user@host# set security zones security-zone zone1 screen syn-frag
```

3. If you are done configuring the device, commit the configuration.

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

Verification

To verify the configuration is working properly, enter the `show security screen statistics zone zone-name` command.

RELATED DOCUMENTATION

| [IP Packet Protection](#) | 130

IP Packet Protection

IN THIS SECTION

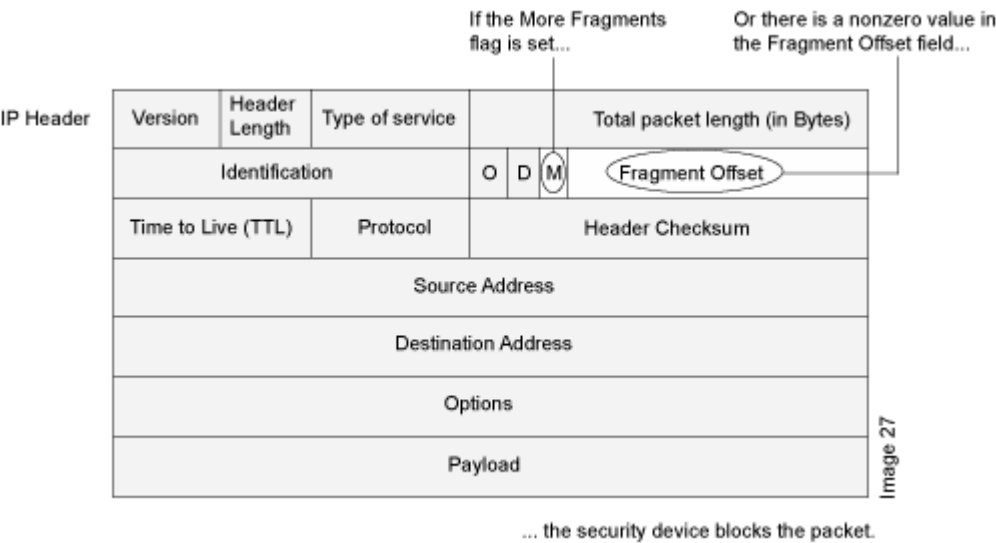
- [Understanding IP Packet Fragment Protection](#) | 130
- [Example: Dropping Fragmented IP Packets](#) | 132
- [Understanding Bad IP Option Protection](#) | 134
- [Example: Blocking IP Packets with Incorrectly Formatted Options](#) | 135
- [Understanding Unknown Protocol Protection](#) | 137
- [Example: Dropping Packets Using an Unknown Protocol](#) | 138
- [Understanding Allowlists for IP Block Fragment Screen](#) | 140

Some attackers can abuse the IP option fields, the original intent of which was (and still is) to provide special routing controls, diagnostic tools, and security. By misconfiguring these options, attackers produce either incomplete or malformed fields within a packet. Attackers can use these malformed packets to compromise hosts on the network. For more information, see the following topics:

Understanding IP Packet Fragment Protection

As packets traverse different networks, it is sometimes necessary to break a packet into smaller pieces (fragments) based upon the maximum transmission unit (MTU) of each network. IP fragments might contain an attacker's attempt to exploit the vulnerabilities in the packet reassembly code of specific IP stack implementations. When the victim receives these packets, the results can range from processing the packets incorrectly to crashing the entire system. See [Figure 17 on page 131](#).

Figure 17: IP Packet Fragments




When you enable Junos OS to deny IP fragments on a security zone, it blocks all IP packet fragments that it receives at interfaces bound to that zone.

NOTE: Junos OS supports IP fragment protection for both IPv4 and IPv6 packets.

In IPv6 packets, fragment information is not present in the IPv6 header. The fragment information is present in the fragment extension header, which is responsible for IPv6 fragmentation and reassembly. The source node inserts the fragment extension header between the IPv6 header and the payload header if fragmentation is required. See [Figure 18 on page 131](#).

Figure 18: IPv6 Packet



The general format of the fragment extension header is shown in [Figure 19 on page 132](#).

Figure 19: Fragment Extension Header

FRAGMENT EXTENSION HEADER FORMAT																																
Offsets	Octet	0						1							2							3										
Octet	Bit	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0	0	Next Header						Reserved							Fragment Offset										Res		M					
4	32	Identification																														

Example: Dropping Fragmented IP Packets

IN THIS SECTION

- Requirements | 132
- Overview | 132
- Configuration | 133
- Verification | 133

This example shows how to drop fragmented IP packets.

Requirements

Before you begin, Understand IP packet fragment protection. See "[Suspicious Packet Attributes Overview](#)" on page 120.

Overview

IN THIS SECTION

- Topology | 133

When this feature is enabled, Junos OS denies IP fragments on a security zone and blocks all IP packet fragments that are received at interfaces bound to that zone.

In this example, you configure the block fragment screen to drop fragmented IP packets originating from the zone1 security zone.

Topology

Configuration

IN THIS SECTION

- [Procedure | 133](#)

Procedure

Step-by-Step Procedure

To drop fragmented IP packets:

1. Configure the screen.

```
[edit]
user@host# set security screen ids-option block-frag ip block-frag
```

2. Configure the security zone.

```
[edit]
user@host# set security zones security-zone zone1 screen block-frag
```

3. If you are done configuring the device, commit the configuration.

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

Verification


To verify the configuration is working properly, enter the `show security screen statistics zone zone-name` command.

Understanding Bad IP Option Protection

The IP standard RFC 791, *Internet Protocol*, specifies a set of eight options that provide special routing controls, diagnostic tools, and security. Although the original, intended uses for these options served worthy ends, people have figured out ways to twist these options to accomplish less commendable objectives.

Either intentionally or accidentally, attackers sometimes configure IP options incorrectly, producing either incomplete or malformed fields. Regardless of the intentions of the person who crafted the packet, the incorrect formatting is anomalous and potentially harmful to the intended recipient. See [Figure 20 on page 134](#).

Figure 20: Incorrectly Formatted IP Options



When you enable the bad IP option protection screen option, Junos OS blocks packets when any IP option in the IP packet header is incorrectly formatted. Additionally, Junos OS records the event in the event log.



NOTE: Junos OS supports bad IP option protection for both IPv4 and IPv6 packets.

Example: Blocking IP Packets with Incorrectly Formatted Options

IN THIS SECTION

- Requirements | 135
- Overview | 135
- Configuration | 136
- Verification | 136

This example shows how to block large ICMP packets with incorrectly formatted options.

Requirements

Before you begin, Understand bad IP option protection. See ["Suspicious Packet Attributes Overview"](#) on page 120.

Overview

IN THIS SECTION

- Topology | 135

When you enable the bad IP option protection screen option, Junos OS blocks packets when any IP option in the IP packet header is incorrectly formatted. Additionally, Junos OS records the event in the event log.

In this example, you configure the IP bad option screen to block large ICMP packets originating from the zone1 security zone.

Topology

Configuration

IN THIS SECTION

- [Procedure](#) | 136

Procedure

Step-by-Step Procedure

To detect and block IP packets with incorrectly formatted IP options:

1. Configure the screen.

```
[edit]
user@host# set security screen ids-option ip-bad-option ip bad-option
```



NOTE: Currently this screen option is applicable only to IPv4.

2. Configure a security zone.

```
[edit]
user@host# set security zones security-zone zone1 screen ip-bad-option
```

3. If you are done configuring the device, commit the configuration.

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

Verification

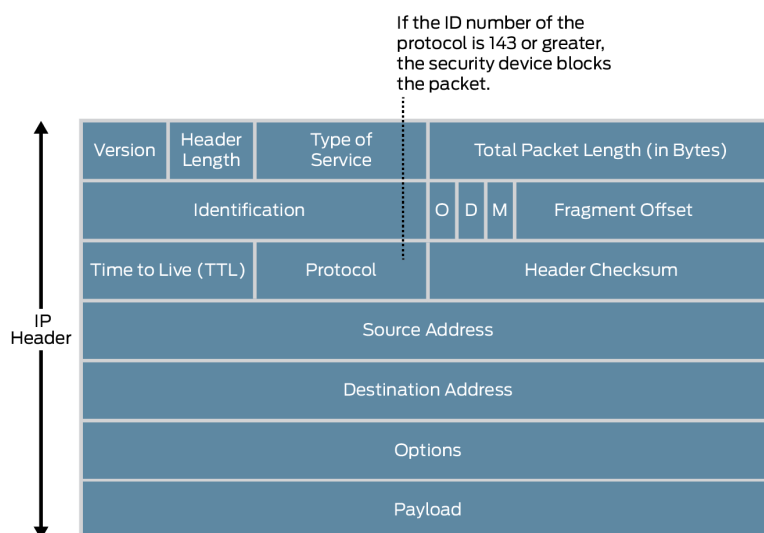
To verify the configuration is working properly, enter the `show security screen statistics zone zone-name` command.

Understanding Unknown Protocol Protection

Based on the latest IANA protocol numbers document, the protocol types with ID numbers of 143 or greater are reserved and undefined at this time. Precisely because these protocols are undefined, there is no way to know in advance if a particular unknown protocol is benign or malicious.

Unless your network makes use of a nonstandard protocol with an ID number of 143 or greater, a cautious stance is to block such unknown elements from entering your protected network. See [Figure 21 on page 137](#).

Figure 21: Unknown Protocols



When you enable the unknown protocol protection screen option, Junos OS drops packets when the protocol field contains a protocol ID number of 143 or greater by default.



NOTE: When you enable the unknown protocol protection screen option for IPv6 protocol, Junos OS drops packets when the protocol field contains a protocol ID number of 143 or greater by default.

Example: Dropping Packets Using an Unknown Protocol

IN THIS SECTION

- Requirements | 138
- Overview | 138
- Configuration | 139
- Verification | 139

This example shows how to drop packets using an unknown protocol.

Requirements

Before you begin, Understand unknown protocol protection. See ["Suspicious Packet Attributes Overview" on page 120](#).

Overview

IN THIS SECTION

- Topology | 138

When you enable the unknown protocol protection screen option, Junos OS drops packets when the protocol field contains a protocol ID number of 137 or greater by default.

In this example, you configure the unknown protocol screen to block packets with an unknown protocol originating from the zone1 security zone.

Topology

Configuration

IN THIS SECTION

- [Procedure | 139](#)

Procedure

Step-by-Step Procedure

To drop packets that use an unknown protocol:

1. Configure the unknown protocol screen.

```
[edit]
user@host# set security screen ids-option unknown-protocol ip unknown-protocol
```

2. Configure a security zone.

```
[edit]
user@host# set security zones security-zone zone1 screen unknown-protocol
```

3. If you are done configuring the device, commit the configuration.

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

Verification

To verify the configuration is working properly, enter the `show security screen statistics zone zone-name` command.

Understanding Allowlists for IP Block Fragment Screen

IN THIS SECTION

- [Benefits of IP Block Fragment Allowlist | 140](#)

Junos OS provides the administrative option to configure an allowlist of trusted IP addresses on IP block fragment screen. When you enable IP block fragmentation in a zone, Junos OS denies IP fragments and blocks all IP packet fragments. All the fragmented IP packets will be dropped. To avoid these packets dropping and instead allow these packets to bypass the IP block fragmentation check, you must configure IP block fragment allowlist.

When you configure allowlist on IP block fragment screen, the traffic from source addresses in the allowlist groups bypasses the IP block fragmentation check. IP block fragment allowlist supports both IPv4 and IPv6 addresses and in each allowlist, there can be up to 32 IP address prefixes. You can configure single address or subnet address.

Benefits of IP Block Fragment Allowlist

- IP block fragment allowlist bypasses the IP block fragmentation check to allow fragmented IP packets from specific sources.

RELATED DOCUMENTATION

| [Suspicious Packet Attributes Overview | 120](#)

4

CHAPTER

Network Reconnaissance

Reconnaissance Deterrence Overview | 142

IP Address Sweep and Port Scan | 142

Operating System Identification Probes | 162

Attacker Evasion Techniques | 176

Reconnaissance Deterrence Overview

Attackers can better plan their attack when they first know the layout of the targeted network (which IP addresses have active hosts), the possible entry points (which port numbers are active on the active hosts), and the constitution of their victims (which operating system the active hosts are running). To gain this information, attackers must perform reconnaissance.

Juniper Networks provides several screen options for deterring attackers' reconnaissance efforts and thereby hindering them from obtaining valuable information about the protected network and network resources.

RELATED DOCUMENTATION

[Operating System Identification Probes | 162](#)

[Attacker Evasion Techniques | 176](#)

IP Address Sweep and Port Scan

IN THIS SECTION

- [Understanding Network Reconnaissance Using IP Options | 143](#)
- [Example: Detecting Packets That Use IP Screen Options for Reconnaissance | 147](#)
- [Understanding IP Address Sweeps | 151](#)
- [Example: Blocking IP Address Sweeps | 152](#)
- [Understanding TCP Port Scanning | 155](#)
- [Understanding UDP Port Scanning | 156](#)
- [Enhancing Traffic Management by Blocking Port Scans | 158](#)

An address sweep occurs when one source IP address sends a predefined number of ICMP packets to various hosts within a predefined interval of time. Port scanning occurs when one source IP address sends IP packets containing TCP SYN segments to a predefined number of different ports at the same destination IP address within a predefined time interval. For more information, see the following topics:

Understanding Network Reconnaissance Using IP Options

IN THIS SECTION

- [Uses for IP Packet Header Options | 143](#)
- [Screen Options for Detecting IP Options Used for Reconnaissance | 146](#)

The IP standard RFC 791, *Internet Protocol*, specifies a set of options for providing special routing controls, diagnostic tools, and security.

RFC 791 states that these options are “unnecessary for the most common communications” and, in reality, they rarely appear in IP packet headers. These options appear after the destination address in an IP packet header, as shown in [Figure 22 on page 143](#). When they do appear, they are frequently being put to some illegitimate use.

Figure 22: Routing Options

Version	Header	Type of Service	Total Packet Length (in Bytes)			
Identification			O	D	M	Fragment Offset
Time to Live (TTL)	Protocol		Header Checksum			
Source Address						
Destination Address						
Options						
Payload						

g030607

This topic contains the following sections:

Uses for IP Packet Header Options

[Table 9 on page 144](#) lists the IP options and their accompanying attributes.

Table 9: IP Options and Attributes

Type	Class	Number	Length	Intended Use	Nefarious Use
End of Options	0*	0	0	Indicates the end of one or more IP options.	None.
No Options	0	1	0	Indicates there are no IP options in the header.	None.
Security	0	2	11 bits	<p>Provides a way for hosts to send security, TCC (closed user group) parameters, and Handling Restriction Codes compatible with Department of Defense (DoD) requirements. (This option, as specified in RFC 791, <i>Internet Protocol</i>, and RFC 1038, <i>Revised IP Security Option</i>, is obsolete.)</p> <p>Currently, this screen option is applicable only to IPv4.</p>	Unknown. However, because it is obsolete, its presence in an IP header is suspect.
Loose Source Route	0	3	Varies	Specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.	Evasion. The attacker can use the specified routes to hide the true source of a packet or to gain access to a protected network.

Table 9: IP Options and Attributes *(Continued)*

Type	Class	Number	Length	Intended Use	Nefarious Use
Record Route	0	7	Varies	<p>Records the IP addresses of the network devices along the path that the IP packet travels. The destination machine can then extract and process the route information. (Due to the size limitation of 40 bytes for both the option and storage space, this can only record up to 9 IP addresses.)</p> <p>Currently, this screen option is applicable only to IPv4.</p>	Reconnaissance. If the destination host is a compromised machine in the attacker's control, he or she can glean information about the topology and addressing scheme of the network through which the packet passed.
Stream ID	0	8	4 bits	<p>(Obsolete) Provided a way for the 16-bit SATNET stream identifier to be carried through networks that did not support the stream concept.</p> <p>Currently, this screen option is applicable only to IPv4.</p>	Unknown. However, because it is obsolete, its presence in an IP header is suspect.
Strict Source Route	0	9	Varies	<p>Specifies the complete route list for a packet to take on its journey from source to destination. The last address in the list replaces the address in the destination field.</p> <p>Currently, this screen option is applicable only to IPv4.</p>	Evasion. An attacker can use the specified routes to hide the true source of a packet or to gain access to a protected network.

Table 9: IP Options and Attributes (Continued)

Type	Class	Number	Length	Intended Use	Nefarious Use
Timestamp	2**	4		<p>Records the time (in coordinated universal time [UTC]**) when each network device receives the packet during its trip from the point of origin to its destination. The network devices are identified by IP address.</p> <p>This option develops a list of IP addresses of the devices along the path of the packet and the duration of transmission between each one.</p> <p>Currently, this screen option is applicable only to IPv4.</p>	Reconnaissance. If the destination host is a compromised machine in the attacker's control, he or she can glean information about the topology and addressing scheme of the network through which the packet has passed.

* The class of options identified as 0 was designed to provide extra packet or network control.

** The class of options identified as 2 was designed for diagnostics, debugging, and measurement.

*** The timestamp uses the number of milliseconds since midnight UTC. UTC is also known as Greenwich Mean Time (GMT), which is the basis for the international time standard.

Screen Options for Detecting IP Options Used for Reconnaissance

The following screen options detect IP options that an attacker can use for reconnaissance or for some unknown but suspect purpose:

- **Record Route**—Junos OS detects packets where the IP option is 7 (record route) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- **Timestamp**—Junos OS detects packets where the IP option list includes option 4 (Internet timestamp) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.
- **Security**—Junos OS detects packets where the IP option is 2 (security) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.

- Stream ID—Junos OS detects packets where the IP option is 8 (stream ID) and records the event in the screen counters list for the ingress interface. Currently, this screen option is applicable only to IPv4.

If a packet with any of the previous IP options is received, Junos OS flags this as a network reconnaissance attack and records the event for the ingress interface.

Example: Detecting Packets That Use IP Screen Options for Reconnaissance

IN THIS SECTION

- Requirements | 147
- Overview | 147
- Configuration | 148
- Verification | 150

This example shows how to detect packets that use IP screen options for reconnaissance.

Requirements

Before you begin, understand how network reconnaissance works. See "[Understanding Network Reconnaissance Using IP Options](#)" on page 143.

Overview

IN THIS SECTION

- Topology | 148

RFC 791, *Internet Protocol*, specifies a set of options for providing special routing controls, diagnostic tools, and security. The screen options detect IP options that an attacker can use for reconnaissance, including record route, timestamp, security, and stream ID.

In this example, you configure an IP screen screen-1 and enable it in a security zone called zone-1.



NOTE: You can enable only one screen in one security zone.

Topology

Configuration

IN THIS SECTION

● Procedure | [148](#)

Procedure

CLI Quick Configuration

To quickly detect packets with the record route, timestamp, security, and stream ID IP screen options, copy the following commands and paste them into the CLI.

```
[edit]
set security screen ids-option screen-1 ip record-route-option
set security screen ids-option screen-1 ip timestamp-option
set security screen ids-option screen-1 ip security-option
set security screen ids-option screen-1 ip stream-option
set security zones security-zone zone-1 screen screen-1
```

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 detect packets that use IP screen options for reconnaissance:

1. Configure IP screen options.



NOTE: Currently, these screen options support IPv4 only.

```
[edit security screen]
user@host# set ids-option screen-1 ip record-route-option
user@host# set ids-option screen-1 ip timestamp-option
user@host# set ids-option screen-1 ip security-option
user@host# set ids-option screen-1 ip stream-option
```

2. Enable the screen in the security zone.

```
[edit security zones ]
user@host# set security-zone zone-1 screen screen-1
```

Results

From configuration mode, confirm your configuration by entering the `show security screen` 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 screen
  ids-option screen-1 {
    ip {
      record-route-option;
      timestamp-option;
      security-option;
      stream-option;
    }
  }
[edit]
[user@host]show security zones
  zones {
    security-zone zone-1 {
      screen screen-1;
    }
  }
```

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 150](#)
- [Verifying the Security Screen Configuration | 150](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:
```

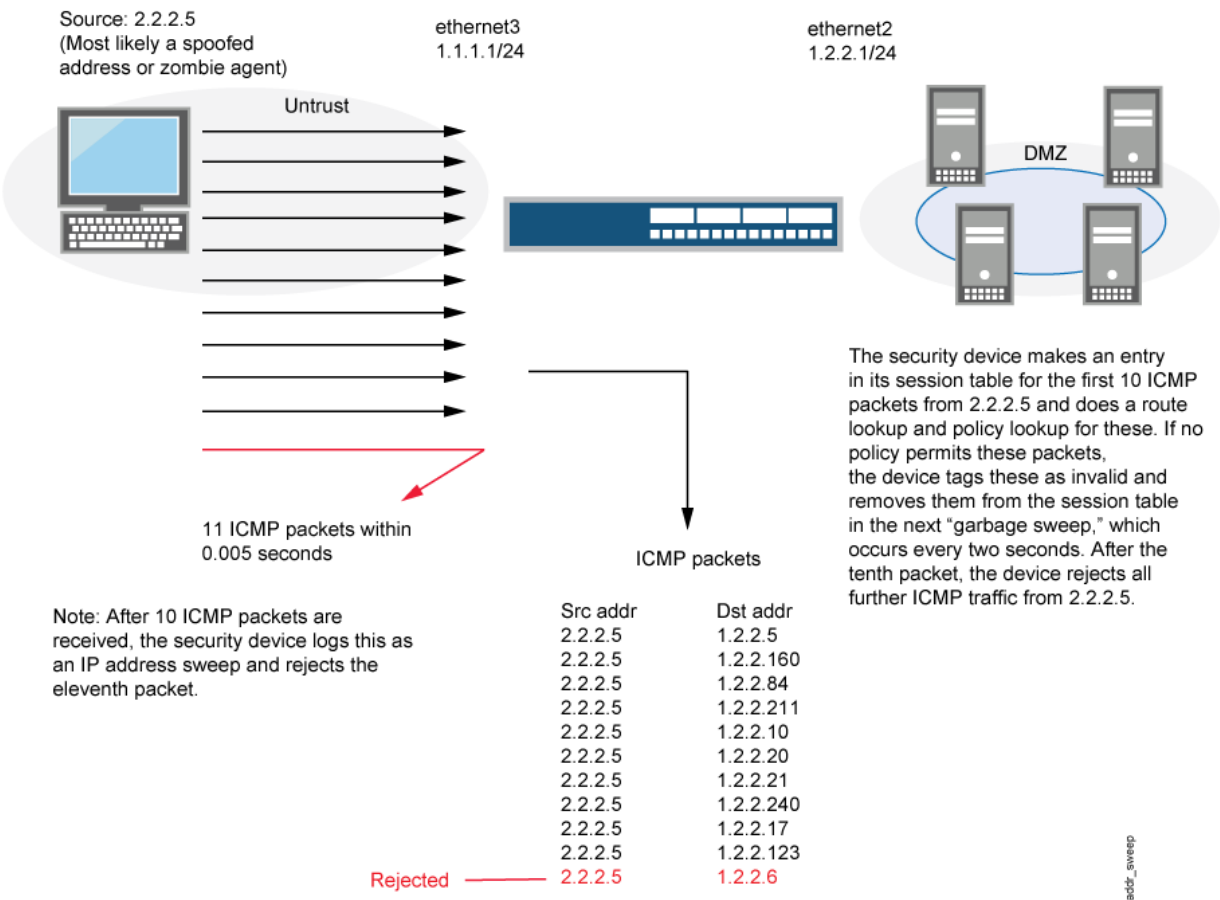
Name	Value
IP record route option	enabled
IP timestamp option	enabled
IP security option	enabled
IP stream option	enabled

Understanding IP Address Sweeps

An address sweep occurs when one source IP address sends a defined number of ICMP packets sent to different hosts within a defined interval (5000 microseconds is the default). The purpose of this attack is to send ICMP packets—typically echo requests—to various hosts in the hopes that at least one replies, thus uncovering an address to target.

Junos OS internally logs the number of ICMP packets to different addresses from one remote source. Using the default settings, if a remote host sends ICMP traffic to 10 addresses in 0.005 seconds (5000 microseconds), then the device flags this as an address sweep attack and rejects all further ICMP packets from that host for the remainder of the specified threshold time period. See [Figure 23 on page 152](#).

Figure 23: Address Sweep



Consider enabling this screen option for a security zone only if there is a policy permitting ICMP traffic from that zone. Otherwise, you do not need to enable the screen option. The lack of such a policy denies all ICMP traffic from that zone, precluding an attacker from successfully performing an IP address sweep anyway.



NOTE: Junos OS supports this screen option for ICMPv6 traffic also.

Example: Blocking IP Address Sweeps

IN THIS SECTION

Requirements | 153

- Overview | 153
- Configuration | 153
- Verification | 154

This example describes how to configure a screen to block an IP address sweep originating from a security zone.

Requirements

Before you begin:

- Understand how IP address sweeps work. See ["Understanding IP Address Sweeps" on page 151](#).
- Configure security zones. See *Security Zones Overview*.

Overview

IN THIS SECTION

- Topology | 153

You need to enable a screen for a security zone if you have configured a policy that permits ICMP traffic from that zone. If you have not configured such a policy, then your system denies all ICMP traffic from that zone, and the attacker cannot perform an IP address sweep successfully anyway.

In this example you configure a 5000-ip-sweep screen to block IP address sweeps originating in the zone-1 security zone.

Topology

Configuration

IN THIS SECTION

- Procedure | 154

Procedure

Step-by-Step Procedure

To configure a screen to block IP address sweeps:

1. Configure a screen.

```
[edit]
user@host# set security screen ids-option 5000-ip-sweep icmp ip-sweep threshold 5000
```

2. Enable the screen in the security zone.

```
[edit]
user@host# set security zones security-zone zone-1 screen 5000-ip-sweep
```

3. If you are done configuring the device, commit the configuration.

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 154](#)
- [Verifying the Security Screen Configuration | 155](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones
Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 5000-ip-sweep
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option 5000-ip-sweep
Screen object status:
```

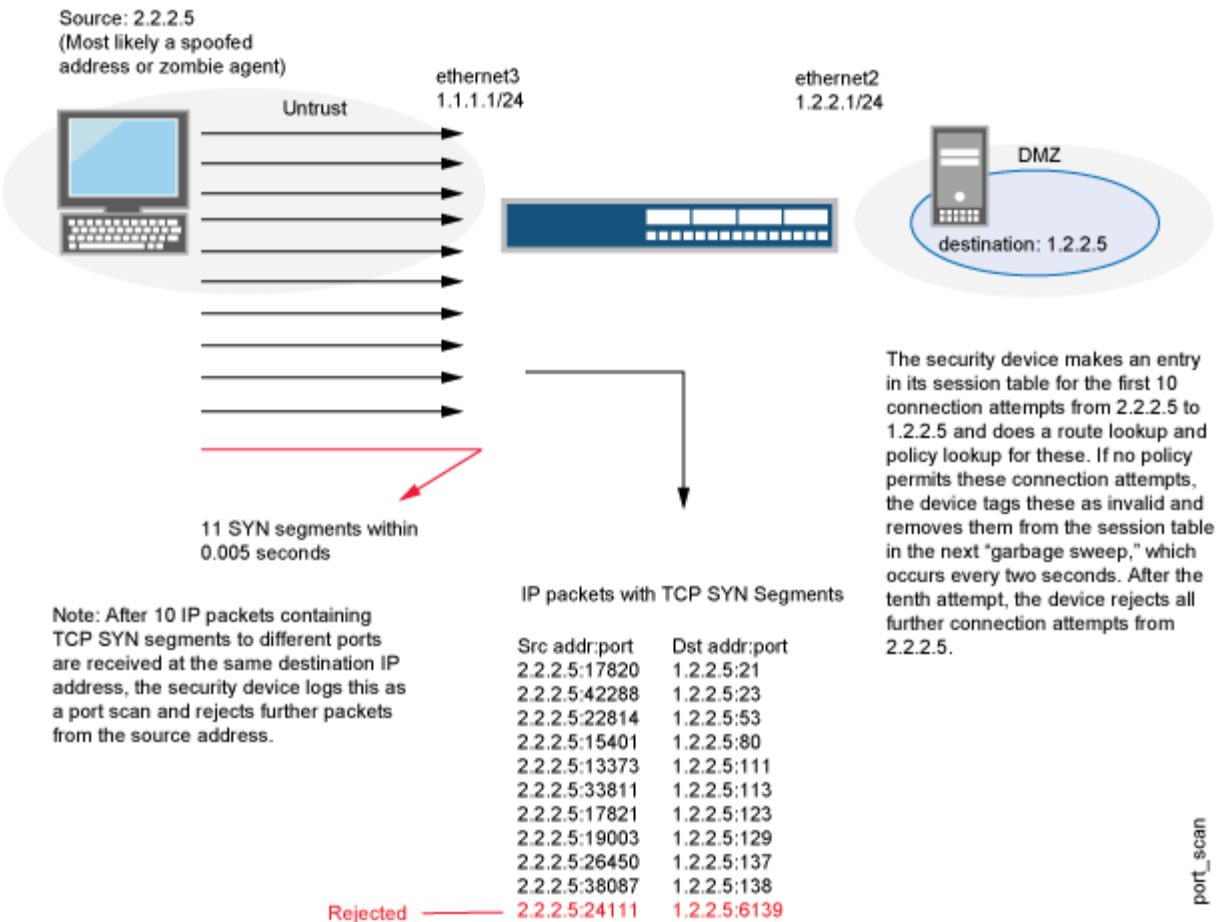
Name	Value
ICMP address sweep threshold	5000

Understanding TCP Port Scanning

A port scan occurs when one source IP address sends IP packets containing TCP SYN segments to 10 different destination ports within a defined interval (5000 microseconds is the default). The purpose of this attack is to scan the available services in the hopes that at least one port will respond, thus identifying a service to target.

Junos OS internally logs the number of different ports scanned from one remote source. Using the default settings, if a remote host scans 10 ports in 0.005 seconds (5000 microseconds), then the device flags this as a port scan attack and rejects all further packets from the remote source, regardless of the destination IP address, for the remainder of the specified timeout period. See [Figure 24 on page 156](#).

Figure 24: Port Scan



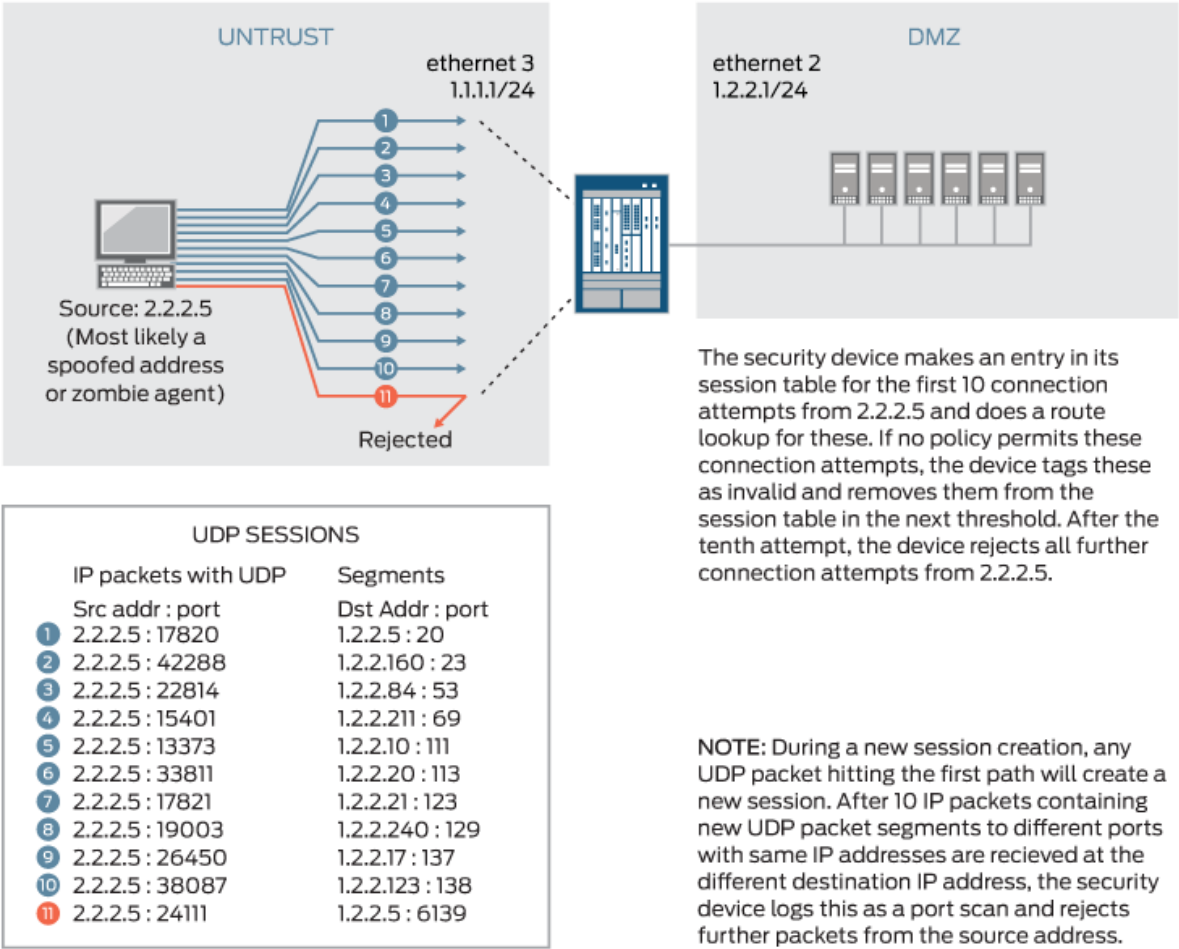
NOTE: Junos OS supports port scanning for both IPv4 and IPv6 traffic.

Understanding UDP Port Scanning

UDP port scan gives statistical information on a session threshold. As the incoming packets traverse the screen, the sessions are established. The number of sessions threshold enforced is based on zone, source IP, and the threshold period and does not allow more than 10 new sessions in the configured

threshold period, for each zone and source IP address. The UDP port scan is disabled by default. When the UDP port scan is enabled, the default threshold period is 5000 microseconds. This value can be manually set to a range of 1000-1,000,000 microseconds. This feature protects some exposed public UDP services against DDoS attacks. See [Figure 25 on page 157](#).

Figure 25: UDP Port Scan



Enhancing Traffic Management by Blocking Port Scans

IN THIS SECTION

- Requirements | 158
- Overview | 158
- Configuration | 159
- Verification | 160

This example shows how to enhance traffic management by configuring a screen to block port scans originating from a particular security zone.

Requirements

Before you begin, understand how port scanning works. See ["Understanding TCP Port Scanning" on page 155](#).

Overview

IN THIS SECTION

- Topology | 158

You can use a port scan to block IP packets containing TCP SYN segments or UDP segments sent to different ports from the same source address within a defined interval. The purpose of this attack is to scan the available services in the hopes that at least one port will respond. Once a port responds, it is identified as a service to target.

In this example, you configure a 5000 port-scan screen to block port scans originating from a particular security zone and then assign the screen to the zone called zone-1.

Topology

Configuration

IN THIS SECTION

- Procedure | 159

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 screen ids-option 5000-port-scan tcp port-scan threshold 5000
set security screen ids-option 10000-port-scan udp port-scan threshold 10000
set security zones security-zone zone-1 screen 5000-port-scan
```

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 screen to block port scans:

1. Configure the screen.

```
[edit security]
user@host# set security screen ids-option 5000-port-scan tcp port-scan threshold 5000
user@host# set security screen ids-option 10000-port-scan udp port-scan threshold 10000
```

2. Enable the screen in the security zone.

```
[edit security]
user@host# set security zones security-zone zone-1 screen 5000-port-scan
```

Results

From configuration mode, confirm your configuration by entering the `show security screen ids-option 5000-port-scan` and `show security zones` 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 screen ids-option 5000-port-scan
tcp {
    port-scan threshold 5000;
}
udp {
    port-scan threshold 10000;
}
```

```
[edit]
user@host# show security zones
security-zone zone-1 {
    screen 5000-port-scan;
}
```

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 160](#)
- [Verifying the Security Screen Configuration | 161](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones
Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: 5000-port-scan
  Interfaces bound: 0
  Interfaces:
```

Meaning

The sample output shows that the screen for zone-1 is enabled for port scan blocking.

Verifying the Security Screen Configuration

Purpose

Verify the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option 5000-port-scan
Screen object status:
Name                                     Value
TCP port scan threshold                 5000
UDP port scan threshold                 10000
```

Meaning

The sample output shows that the port scan blocking is operational with TCP and UDP threshold.

SEE ALSO

| [Attacker Evasion Techniques](#) | 176

Operating System Identification Probes

IN THIS SECTION

- [Understanding Operating System Identification Probes](#) | 162
- [Understanding Domain Name System Resolve](#) | 163
- [Understanding TCP Headers with SYN and FIN Flags Set](#) | 163
- [Example: Blocking Packets with SYN and FIN Flags Set](#) | 164
- [Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set](#) | 168
- [Example: Blocking Packets With FIN Flag Set and Without ACK Flag Set](#) | 169
- [Understanding TCP Header with No Flags Set](#) | 172
- [Example: Blocking Packets with No Flags Set](#) | 173

Prior to launching an exploit, an attacker might probe the targeted host, trying to learn its operating system. Various operating systems react to TCP anomalies in different ways. With that knowledge, an attacker can decide which further attack might inflict more damage to the device, the network, or both. For more information, see the following topics:

Understanding Operating System Identification Probes

Before launching an exploit, attackers might try to probe the targeted host to learn its operating system (OS). With that knowledge, they can better decide which attack to launch and which vulnerabilities to exploit. Junos OS can block reconnaissance probes commonly used to gather information about OS types.

Understanding Domain Name System Resolve

Prior to Junos OS Release 12.1X47, DNS resolution was performed with only UDP as a transport. Messages carried by UDP are restricted to 512 bytes; longer messages are truncated and the traffic class (TC) bit is set in the header. The maximum length of UDP DNS response messages is 512 bytes, but the maximum length of TCP DNS response messages is 65,535 bytes. A DNS resolver knows whether the response is complete if the TC bit is set in the header. Hence, a TCP DNS response can carry more information than a UDP DNS response.

There are three types of DNS resolve behaviors:

- UDP DNS resolve
- TCP DNS resolve
- UDP/TCP DNS resolve



NOTE: A policy uses UDP/TCP DNS resolve to resolve IP addresses. In UDP/TCP DNS resolve, UDP DNS resolve is first used, and when it gets truncated TCP DNS resolve is used.

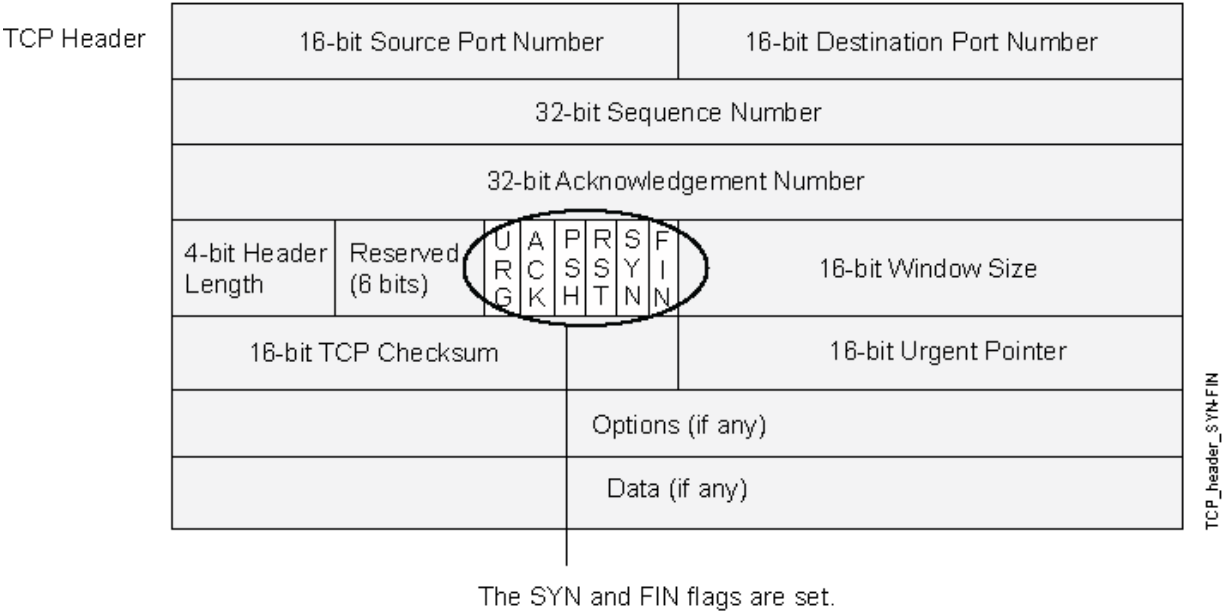


NOTE: A Routing Engine policy supports a maximum of 1024 IPv4 address prefixes and 256 IPv6 address prefixes that can be sent to the PFE. If the maximum number of IPv4 or IPv6 address prefixes exceeds the limits, the addresses over the limitations will not be sent to the PFE and a syslog message is generated. The maximum number of addresses in a TCP DNS response is 4094 for IPv4 addresses and 2340 for IPv6 addresses, but only 1024 IPv4 addresses and 256 IPv6 addresses are loaded to the PFE.

Understanding TCP Headers with SYN and FIN Flags Set

Both the SYN and FIN control flags are not normally set in the same TCP segment header. The SYN flag synchronizes sequence numbers to initiate a TCP connection. The FIN flag indicates the end of data transmission to finish a TCP connection. Their purposes are mutually exclusive. A TCP header with the SYN and FIN flags set is anomalous TCP behavior, causing various responses from the recipient, depending on the OS. See [Figure 26 on page 164](#).

Figure 26: TCP Header with SYN and FIN Flags Set



An attacker can send a segment with both flags set to see what kind of system reply is returned and thereby determine what kind of OS is on the receiving end. The attacker can then use any known system vulnerabilities for further attacks.

When you enable this screen option, Junos OS checks if the SYN and FIN flags are set in TCP headers. If it discovers such a header, it drops the packet.

NOTE: Junos OS supports TCP header with SYN and FIN flags set protection for both IPv4 and IPv6 traffic.

Example: Blocking Packets with SYN and FIN Flags Set

IN THIS SECTION

- Requirements | 165
- Overview | 165
- Configuration | 165
- Verification | 166

This example shows how to create a screen to block packets with the SYN and FIN flags set.

Requirements

Before you begin, understand how TCP headers with SYN and FIN flags work. See "[Understanding TCP Headers with SYN and FIN Flags Set](#)" on page 163.

Overview

IN THIS SECTION

- [Topology](#) | 165

The TCP header with the SYN and FIN flags set cause different responses from a targeted device depending on the OS it is running. The syn-fin screen is enabled for the security zone.

In this example, you create a screen called screen-1 in a security zone to block packets with the SYN and FIN flags set.

Topology

Configuration

IN THIS SECTION

- [Procedure](#) | 165

Procedure

Step-by-Step Procedure

To block packets with both the SYN and FIN flags set:

1. Configure the screen.

```
[edit]  
user@host# set security screen ids-option screen-1 tcp syn-fin
```

2. Enable the screen in the security zone.

```
[edit ]  
user@host# set security zones security-zone zone-1 screen screen-1
```

3. If you are done configuring the device, commit the configuration.

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 166](#)
- [Verifying the Security Screen Configuration | 167](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                Value
  TCP SYN FIN          enabled
```

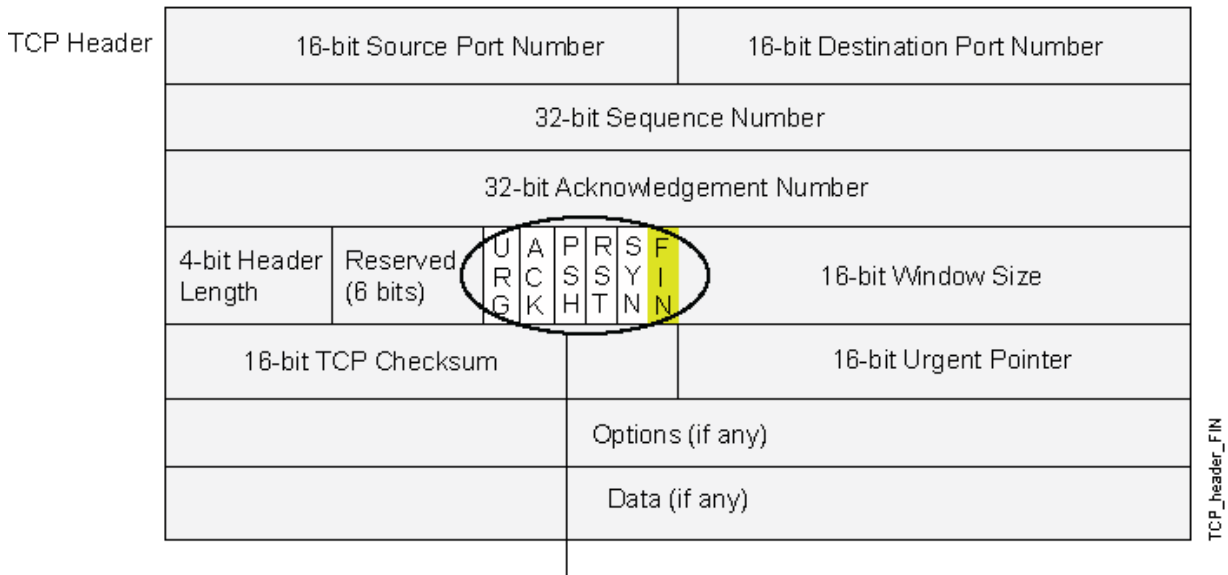
Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set

Figure 27 on page 168 shows TCP segments with the FIN control flag set (to signal the conclusion of a session and terminate the connection). Normally, TCP segments with the FIN flag set also have the ACK flag set (to acknowledge the previous packet received). Because a TCP header with the FIN flag set but not the ACK flag is anomalous TCP behavior, there is no uniform response to this. The OS might respond by sending a TCP segment with the RST flag set. Another might completely ignore it. The victim's response can provide the attacker with a clue as to its OS. (Other purposes for sending a TCP segment with the FIN flag set are to evade detection while performing address and port scans and to evade defenses on guard for a SYN flood by performing a FIN flood instead.)



NOTE: Vendors have interpreted RFC 793, *Transmission Control Protocol*, variously when designing their TCP/IP implementations. When a TCP segment arrives with the FIN flag set but not the ACK flag, some implementations send RST segments, while others drop the packet without sending an RST.

Figure 27: TCP Header with FIN Flag Set



Only the FIN flag is set.

When you enable this screen option, Junos OS checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.



NOTE: Junos OS supports TCP header with SYN and FIN flags set protection for both IPv4 and Ipv6 traffic.

Example: Blocking Packets With FIN Flag Set and Without ACK Flag Set

IN THIS SECTION

- [Requirements | 169](#)
- [Overview | 169](#)
- [Configuration | 170](#)
- [Verification | 170](#)

This example shows how to create a screen to block packets with the FIN flag set but the ACK flag not set.

Requirements

Before you begin, understand how TCP headers work. See "[Understanding TCP Headers With FIN Flag Set and Without ACK Flag Set](#)" on page 168.

Overview

The TCP segments with the FIN flag set also have the ACK flag set to acknowledge the previous packet received. Because a TCP header with the FIN flag set but the ACK flag not set is anomalous TCP behavior, there is no uniform response to this. When you enable the `fin-no-ack` screen option, Junos OS checks if the FIN flag is set but not the ACK flag in TCP headers. If it discovers a packet with such a header, it drops the packet.

In this example, you create a screen called `screen-1` to block packets with the FIN flag set but the ACK flag not set.

Configuration

IN THIS SECTION

- [Procedure | 170](#)

Procedure

Step-by-Step Procedure

To block packets with the FIN flag set but the ACK flag not set:

1. Configure the screen.

```
[edit ]
user@host# set security screen ids-option screen-1 tcp fin-no-ack
```

2. If you are done configuring the device, commit the configuration.

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 171](#)
- [Verifying the Security Screen Configuration | 171](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

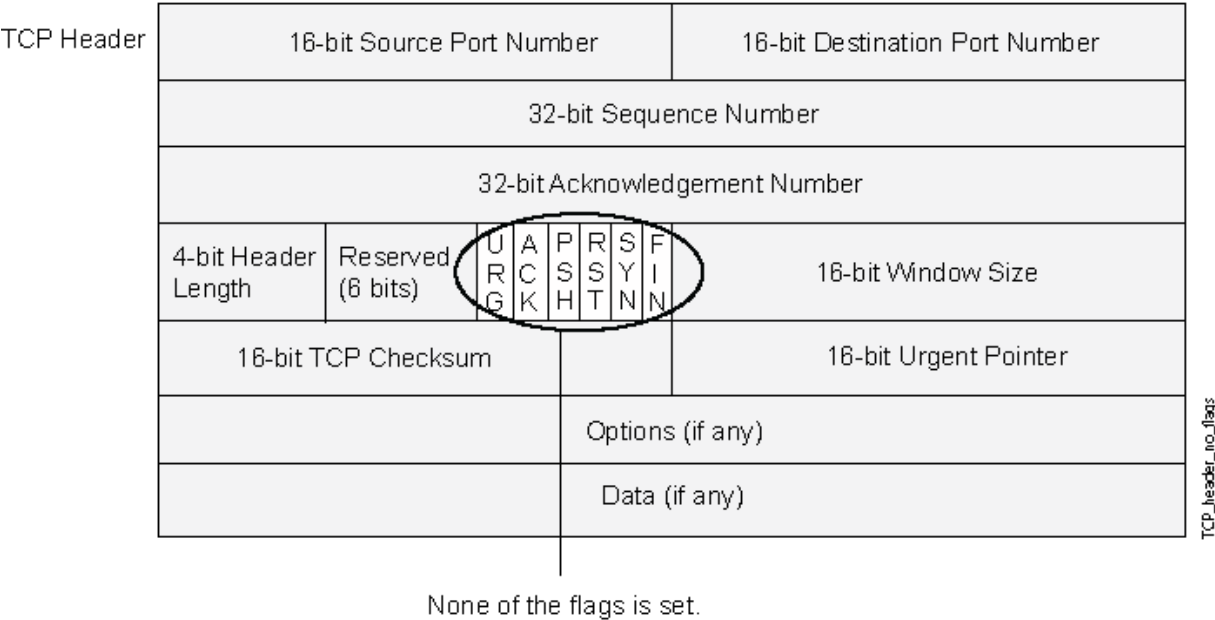
```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

Name                               Value
TCP FIN no ACK                     enabled
```

Understanding TCP Header with No Flags Set

A normal TCP segment header has at least one flag control set. A TCP segment with no control flags set is an anomalous event. Because different operating systems respond differently to such anomalies, the response (or lack of response) from the targeted device can provide a clue as to the type of OS it is running. See [Figure 28 on page 172](#).

Figure 28: TCP Header with No Flags Set



When you enable the device to detect TCP segment headers with no flags set, the device drops all TCP packets with a missing or malformed flags field.



NOTE: Junos OS supports TCP header with no flags set protection for both IPv4 and IPv6 traffic.

Example: Blocking Packets with No Flags Set

IN THIS SECTION

- Requirements | 173
- Overview | 173
- Configuration | 173
- Verification | 174

This example shows how to create a screen to block packets with no flags set.

Requirements

Before you begin, understand how a TCP header with no flags set works. See ["Understanding TCP Header with No Flags Set" on page 172](#).

Overview

A normal TCP segment header has at least one flag control set. A TCP segment with no control flags set is an anomalous event. Because different operating systems respond differently to such anomalies, the response (or lack of response) from the targeted device can provide a clue as to the type of OS it is running.

When you enable the device to detect TCP segment headers with no flags set, the device drops all TCP packets with a missing or malformed flags field.

In this example, you create a screen called screen-1 to block packets with no flags set.

Configuration

IN THIS SECTION

- Procedure | 174

Procedure

Step-by-Step Procedure

To block packets with no flags set:

1. Configure the screen.

```
[edit ]
user@host# set security screen ids-option screen-1 tcp tcp-no-flag
```

2. Enable the screen in the security zone.

```
[edit ]
user@host# set security zones security-zone zone-1 screen screen-1
```

3. If you are done configuring the device, commit the configuration.

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 174](#)
- [Verifying the Security Screen Configuration | 175](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                Value
  TCP no flag         enabled
```

RELATED DOCUMENTATION

[Reconnaissance Deterrence Overview](#) | 142

Attacker Evasion Techniques

IN THIS SECTION

- [Understanding Attacker Evasion Techniques | 176](#)
- [Understanding FIN Scans | 177](#)
- [Thwarting a FIN Scan | 177](#)
- [Understanding TCP SYN Checking | 177](#)
- [Setting TCP SYN Checking | 180](#)
- [Setting TCP Strict SYN Checking | 180](#)
- [Understanding IP Spoofing | 180](#)
- [Example: Blocking IP Spoofing | 181](#)
- [Understanding IP Spoofing in Layer 2 Transparent Mode on Security Devices | 184](#)
- [Configuring IP Spoofing in Layer 2 Transparent Mode on Security Devices | 185](#)
- [Understanding IP Source Route Options | 186](#)
- [Example: Blocking Packets with Either a Loose or a Strict Source Route Option Set | 189](#)
- [Example: Detecting Packets with Either a Loose or a Strict Source Route Option Set | 192](#)

An attacker might use the SYN and FIN flags to launch the attack. The inset also illustrates the configuration of Screen options designed to block these probes. For more information, see the following topics:

Understanding Attacker Evasion Techniques

Whether gathering information or launching an attack, it is generally expected that the attacker avoids detection. Although some IP address and port scans are blatant and easily detectable, more wily attackers use a variety of means to conceal their activity. Techniques such as using FIN scans instead of SYN scans—which attackers know most firewalls and intrusion detection programs detect—indicate an evolution of reconnaissance and exploit techniques for evading detection and successfully accomplishing their tasks.

Understanding FIN Scans

A FIN scan sends TCP segments with the FIN flag set in an attempt to provoke a response (a TCP segment with the RST flag set) and thereby discover an active host or an active port on a host. Attackers might use this approach rather than perform an address sweep with ICMP echo requests or an address scan with SYN segments, because they know that many firewalls typically guard against the latter two approaches but not necessarily against FIN segments. The use of TCP segments with the FIN flag set might evade detection and thereby help the attackers succeed in their reconnaissance efforts.

Thwarting a FIN Scan

To thwart FIN scans, take either or both of the following actions:

- Enable the screen option that specifically blocks TCP segments with the FIN flag set but not the ACK flag, which is anomalous for a TCP segment:

```
user@host#set security screen fin-no-ack tcp fin-no-ack
user@host#set security zones security-zone name screen fin-no-ack
```

where *name* is the name of the zone to which you want to apply this screen option .

- Change the packet processing behavior to reject all non-SYN packets that do not belong to an existing session. The SYN check flag is set as the default.




NOTE: Changing the packet flow to check that the SYN flag is set for packets that do not belong to existing sessions also thwarts other types of non-SYN scans, such as a null scan (when no TCP flags are set).

Understanding TCP SYN Checking

By default, Junos OS checks for SYN flags in the first packet of a session and rejects any TCP segments with non-SYN flags attempting to initiate a session. You can leave this packet flow as is or change it so that Junos OS does not enforce SYN flag checking before creating a session. [Figure 29 on page 178](#) illustrates packet flow sequences both when SYN flag checking is enabled and when it is disabled.

Figure 29: SYN Flag Checking



When Junos OS with SYN flag checking enabled receives a non-SYN TCP segment that does not belong to an existing session, it drops the packet. By default, Junos OS does not send a TCP RST to the source host on receiving the non-SYN segment. You can configure the device to send TCP RST to the source host by using the `set security zones security-zone trust tcp-rst` command. If the code bit of the initial non-SYN TCP packet is RST, the device does not send a TCP-RST.

Not checking for the SYN flag in the first packets offers the following advantages:

- **NSRP with Asymmetric Routing**—In an active/active NSRP configuration in a dynamic routing environment, a host might send the initial TCP segment with the SYN flag set to one device (Device-A), but the SYN/ACK might be routed to the other device in the cluster (Device-B). If this asymmetric routing occurs after Device-A has synchronized its session with Device-B, all is well. On the other hand, if the SYN/ACK response reaches Device-B before Device-A synchronizes the session and SYN checking is enabled, Device-B rejects the SYN/ACK, and the session cannot be established. With SYN checking disabled, Device-B accepts the SYN/ACK response—even though there is no existing session to which it belongs—and creates a new session table entry for it.
- **Uninterrupted Sessions**—If you reset the device or even change a component in the core section of a policy and SYN checking is enabled, all existing sessions or those sessions to which the policy change applies are disrupted and must be restarted. Disabling SYN checking avoids such disruptions to network traffic flows.



NOTE: A solution to this scenario is to install the device with SYN checking disabled initially. Then, after a few hours—when established sessions are running through the device—enable SYN checking. The core section in a policy contains the following main components: source and destination zones, source and destination addresses, one or more services, and an action.

However, the previous advantages exact the following security sacrifices:

- **Reconnaissance Holes**—When an initial TCP segment with a non-SYN flag—such as ACK, URG, RST, FIN—arrives at a closed port, many operating systems (Windows, for example) respond with a TCP segment that has the RST flag set. If the port is open, then the recipient does not generate any response.

By analyzing these responses or lack thereof, an intelligence gatherer can perform reconnaissance on the protected network and also on the Junos OS policy set. If a TCP segment is sent with a non-SYN flag set and the policy permits it through, the destination host receiving such a segment might drop it and respond with a TCP segment that has the RST flag set. Such a response informs the perpetrator of the presence of an active host at a specific address and that the targeted port number is closed. The intelligence gatherer also learns that the firewall policy permits access to that port number on that host.

By enabling SYN flag checking, Junos OS drops TCP segments without a SYN flag if they do not belong to an existing session. It does not return a TCP RST segment. Consequently, the scanner gets no replies regardless of the policy set or whether the port is open or closed on the targeted host.

- **Session Table Floods**—If SYN checking is disabled, an attacker can bypass the Junos OS SYN flood protection feature by flooding a protected network with a barrage of TCP segments that have non-SYN flags set. Although the targeted hosts drop the packets—and possibly send TCP RST segments in reply—such a flood can fill up the session table of the Juniper Networks device. When the session table is full, the device cannot process new sessions for legitimate traffic.

By enabling SYN checking and SYN flood protection, you can thwart this kind of attack. Checking that the SYN flag is set on the initial packet in a session forces all new sessions to begin with a TCP segment that has the SYN flag set. SYN flood protection then limits the number of TCP SYN segments per second so that the session table does not become overwhelmed.

If you do not need SYN checking disabled, Juniper Networks strongly recommends that it be enabled (its default state for an initial installation of Junos OS). You can enable it with the `set flow tcp-syn-check` command. With SYN checking enabled, the device rejects TCP segments with non-SYN flags set unless they belong to an established session.

Setting TCP SYN Checking

With SYN checking enabled, the device rejects TCP segments with non-SYN flags set unless they belong to an established session. Enabling SYN checking can help prevent attacker reconnaissance and session table floods. TCP SYN checking is enabled by default.

To disable SYN checking:

```
user@host#set security flow tcp-session no-syn-check
```

Setting TCP Strict SYN Checking

With strict SYN checking enabled, the device enables the strict three-way handshake check for the TCP session. It enhances security by dropping data packets before the three-way handshake is done. TCP strict SYN checking is disabled by default.



NOTE: The strict-syn-check option cannot be enabled if no-syn-check or no-syn-check-in-tunnel is enabled.



NOTE: When you enable strict-syn-check the SYN packets carrying data are dropped.

To enable strict SYN checking:

```
user@host#set security flow tcp-session strict-syn-check
```

Understanding IP Spoofing

One method of attempting to gain access to a restricted area of the network is to insert a false source address in the packet header to make the packet appear to come from a trusted source. This technique is called IP spoofing. The mechanism to detect IP spoofing relies on route table entries. For example, if a packet with source IP address 10.1.1.6 arrives at ge-0/0/1, but Junos OS has a route to 10.1.1.0/24 through ge-0/0/0, a check for IP spoofing discovers that this address arrived at an invalid interface as defined in the route table. A valid packet from 10.1.1.6 can only arrive via ge-0/0/0, not ge-0/0/1. Therefore, Junos OS concludes that the packet has a spoofed source IP address and discards it.



NOTE: Junos OS detects and drops both IPv4 and IPv6 spoofed packets.

Example: Blocking IP Spoofing

IN THIS SECTION

- [Requirements | 181](#)
- [Overview | 181](#)
- [Configuration | 181](#)
- [Verification | 182](#)

This example shows how to configure a screen to block IP spoof attacks.

Requirements

Before you begin, understand how IP Spoofing works. See ["Understanding IP Spoofing" on page 180](#).

Overview

One method of attempting to gain access to a restricted area of a network is to insert a bogus source address in the packet header to make the packet appear to come from a trusted source. This technique is called IP spoofing.

In this example, you configure a screen called screen-1 to block IP spoof attacks and enable the screen in the zone-1 security zone.

Configuration

IN THIS SECTION

- [Procedure | 182](#)

Procedure

Step-by-Step Procedure

To block IP spoofing:

1. Configure the screen.

```
[edit ]
user@host# set security screen ids-option screen-1 ip spoofing
```

2. Enable the screen in the security zone.

```
[edit]
user@host# set security zone security-zone zone-1 screen screen-1
```

3. If you are done configuring the device, commit the configuration.

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 182](#)
- [Verifying the Security Screen Configuration | 183](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

  Name                Value
  IP spoofing         enabled
```

Understanding IP Spoofing in Layer 2 Transparent Mode on Security Devices

In an IP spoofing attack, the attacker gains access to a restricted area of the network and inserts a false source address in the packet header to make the packet appear to come from a trusted source. IP spoofing is most frequently used in denial-of-service (DoS) attacks. When SRX Series Firewalls are operating in transparent mode, the IP spoof-checking mechanism makes use of address book entries. Address books only exist on the Routing Engine. IP spoofing in Layer 2 transparent mode is performed on the Packet Forwarding Engine. Address book information cannot be obtained from the Routing Engine each time a packet is received by the Packet Forwarding Engine. Therefore, address books attached to the Layer 2 zones must be pushed to the Packet Forwarding Engine.



NOTE: IP spoofing in Layer 2 transparent mode does not support DNS and wildcard addresses.

When a packet is received by the Packet Forwarding Engine, the packet's source IP address is checked to determine if it is in the incoming zone's address-book. If the packet's source IP address is in the incoming zone's address book, then this IP address is allowed on the interface, and traffic is passed.

If the source IP address is not present in the incoming zone's address-book, but exists in other zones', then the IP address is considered a spoofed IP. Accordingly, actions such as drop and logging can be taken depending on the screen configuration (alarm-without-drop).



NOTE: If the alarm-without-drop option is configured, the Layer 2 and Layer 3 spoofing packets only trigger an alarm message, but the packets are not dropped.

If a packet's source IP address is not present in the incoming zone's address book or other zones', then you cannot determine if the IP is spoofed or not. In such instances, the packet is passed.

Junos OS takes into account the following match conditions while it searches for source IP addresses in the address book:

- **Host-match**—The IP address match found in the address-book is an address without a prefix.
- **Prefix-match**—The IP address match found in the address-book is an address with a prefix.
- **Any-match**—The IP address match found in the address-book is "any", "any-IPv4", or "any-IPv6".
- **No-match**—No IP address match is found.

Configuring IP Spoofing in Layer 2 Transparent Mode on Security Devices

You can configure the IP spoof-checking mechanism to determine whether or not an IP is being spoofed.

To configure IP spoofing in Layer 2 transparent mode:

1. Set the interface in Layer 2 transparent mode.

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 family ethernet-switching
```

2. (Optional) Set the zone in Layer 2 transparent mode.

```
[edit]
user@host# set security zones security-zone untrust interfaces ge-0/0/1.0
```

3. Configure the address book.

```
[edit]
user@host# set security address-book my-book address myadd1 10.1.1.0/24
user@host# set security address-book my-book address myadd2 10.1.2.0/24
```

4. Apply the address book to the zone.

```
[edit]
user@host# set security address-book my-book attach zone untrust
```

5. Configure screen IP spoofing.

```
[edit]
user@host# set security screen ids-option my-screen ip spoofing
```

6. Apply the screen to the zone.

```
[edit]
user@host# set security zones security-zone untrust screen my-screen
```

7. (Optional) Configure the alarm-without-drop option.

[edit]

```
user@host# set security screen ids-option my-screen alarm-without-drop
```

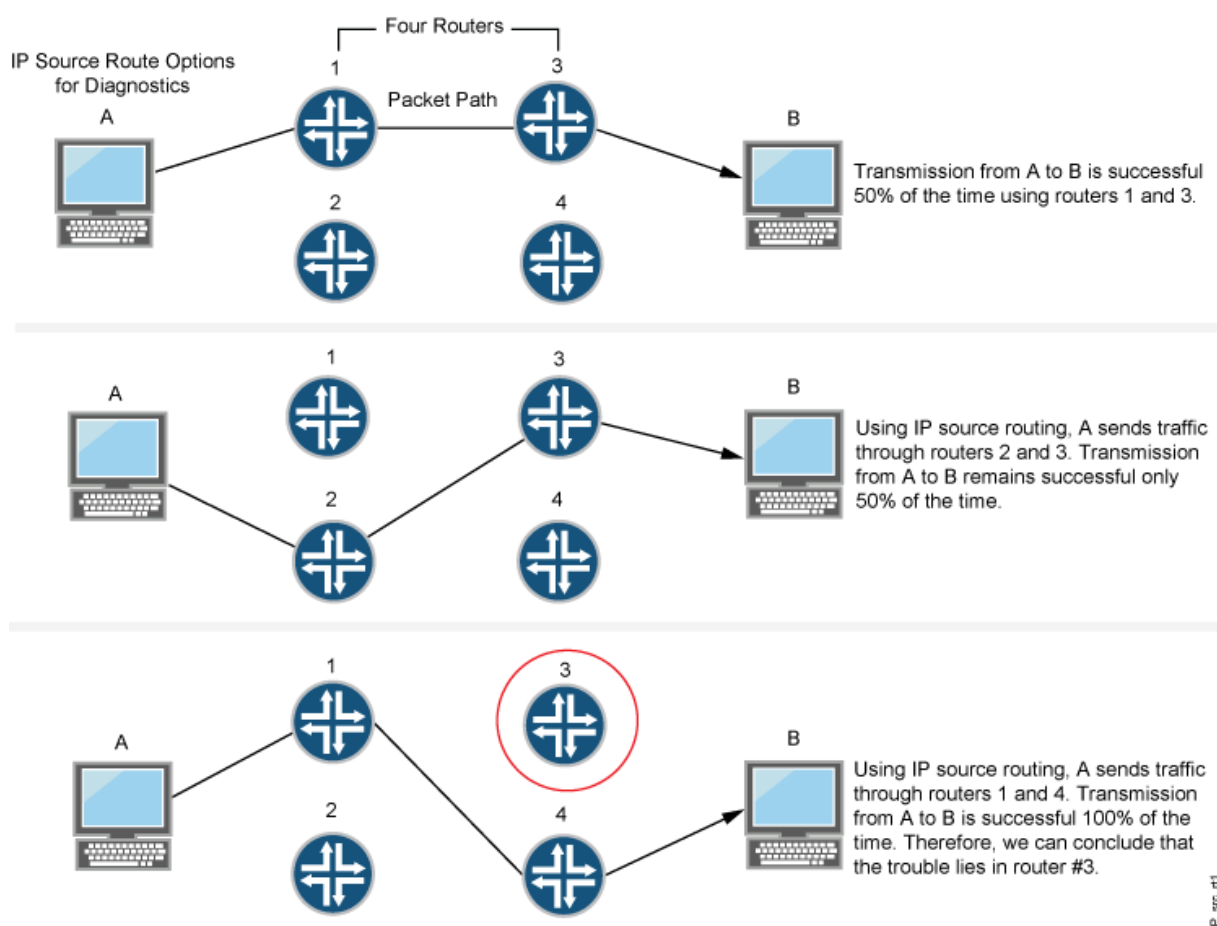


NOTE: If the alarm-without-drop option is configured, the Layer 2 spoofing packet only triggers an alarm message, but the packet is not dropped.

Understanding IP Source Route Options

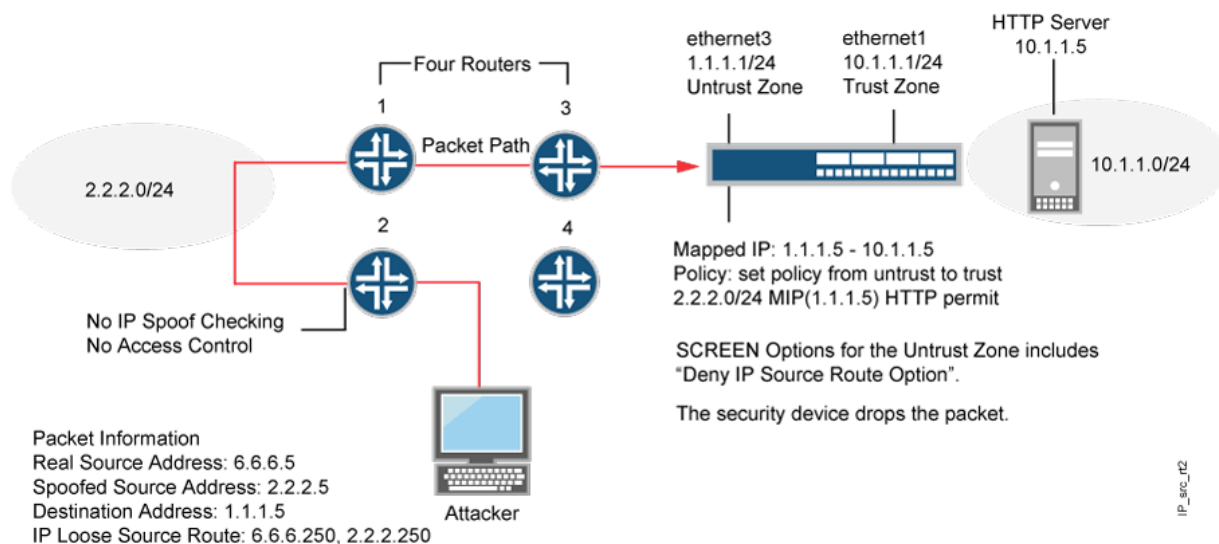
Source routing was designed to allow users at the source of an IP packet transmission to specify the IP addresses of the devices (also referred to as “hops”) along the path that they want an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis. If, for example, the transmission of a packet to a particular destination meets with irregular success, you might first use either the record route or the timestamp IP option to discover the addresses of devices along the path or paths that the packet takes. You can then use either the loose or the strict source route option to direct traffic along a specific path, using the addresses you learned from the results that the record route or timestamp options produced. By changing device addresses to alter the path and sending several packets along different paths, you can note changes that either improve or lessen the success rate. Through analysis and the process of elimination, you might be able to deduce where the trouble lies. See [Figure 30 on page 187](#).

Figure 30: IP Source Routing



Although the uses of IP source route options were originally benign, attackers have learned to put them to more devious uses. They can use IP source route options to hide their true address and access restricted areas of a network by specifying a different path. For an example showing how an attacker can put both deceptions to use, consider the following scenario as illustrated in [Figure 31 on page 188](#).

Figure 31: Loose IP Source Route Option for Deception



Junos OS only allows traffic 2.2.2.0/24 if it comes through ethernet1, an interface bound to zone_external. Devices 3 and 4 enforce access controls but devices 1 and 2 do not. Furthermore, device 2 does not check for IP spoofing. The attacker spoofs the source address and, by using the loose source route option, directs the packet through device 2 to the 2.2.2.0/24 network and from there out device 1. Device 1 forwards it to device 3, which forwards it to the Juniper Networks device. Because the packet came from the 2.2.2.0/24 subnet and has a source address from that subnet, it seems to be valid. However, one remnant of the earlier chicanery remains: the loose source route option. In this example, you have enabled the deny IP source route screen option for zone_external. When the packet arrives at ethernet3, the device rejects it.

You can enable the device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface. The screen options are as follows:

- **Deny IP Source Route Option**—Enable this option to block all IP traffic that employs the loose or strict source route option. Source route options can allow an attacker to enter a network with a false IP address.
- **Detect IP Loose Source Route Option**—The device detects packets where the IP option is 3 (Loose Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies a partial route list for a packet to take on its journey from source to destination. The packet must proceed in the order of addresses specified, but it is allowed to pass through other devices in between those specified.
- **Detect IP Strict Source Route Option**—The device detects packets where the IP option is 9 (Strict Source Routing) and records the event in the screen counters list for the ingress interface. This option specifies the complete route list for a packet to take on its journey from source to destination.

The last address in the list replaces the address in the destination field. Currently, this screen option is applicable to IPv4 only.

Example: Blocking Packets with Either a Loose or a Strict Source Route Option Set

IN THIS SECTION

- Requirements | 189
- Overview | 189
- Configuration | 190
- Verification | 190

This example shows how to block packets with either a loose or a strict source route option set.

Requirements

Before you begin, understand how IP source route options work. See ["Understanding IP Source Route Options" on page 186](#).

Overview

Source routing allows users at the source of an IP packet transmission to specify the IP addresses of the devices (also referred to as “hops”) along the path that they want an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis.

You can enable the device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface.

In this example you create the screen called screen-1 to block packets with either a loose or a strict source route option set and enable the screen in the zone-1 security zone.

Configuration

IN THIS SECTION

- [Procedure | 190](#)

Procedure

Step-by-Step Procedure

To block packets with either the loose or the strict source route option set:

1. Configure the screen.

```
[edit ]
user@host# set security screen ids-option screen-1 ip source-route-option
```

2. Enable the screen in the security zone.

```
[edit ]
user@host# set security zones security-zone zone-1 screen screen-1
```

3. If you are done configuring the device, commit the configuration.

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

Verification

IN THIS SECTION

- [Verifying the Screens in the Security Zone | 191](#)
- [Verifying the Security Screen Configuration | 191](#)

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones
Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:
```

Name	Value
IP source route option	enabled

Example: Detecting Packets with Either a Loose or a Strict Source Route Option Set

IN THIS SECTION

- Requirements | 192
- Overview | 192
- Configuration | 192
- Verification | 193

This example shows how to detect packets with either a loose or a strict source route option set.

Requirements

Before you begin, understand how IP source route options work. See ["Understanding IP Source Route Options" on page 186](#).

Overview

Source routing allows users at the source of an IP packet transmission to specify the IP addresses of the devices (also referred to as "hops") along the path that they want an IP packet to take on its way to its destination. The original intent of the IP source route options was to provide routing control tools to aid diagnostic analysis.

You can enable the device to either block any packets with loose or strict source route options set or detect such packets and then record the event in the counters list for the ingress interface.

In this example, you create two screens called screen-1 and screen-2 to detect and record, but not block, packets with a loose or strict source route option set and enable the screens in zones zone-1 and zone-2.

Configuration

IN THIS SECTION

- Procedure | 193

Procedure

Step-by-Step Procedure

To detect and record, but not block, packets with a loose or strict source route option set:

1. Configure the loose source screen.

```
[edit]
user@host# set security screen ids-option screen-1 ip loose-source-route-option
```

2. Configure the strict source route screen.

```
[edit]
user@host# set security screen ids-option screen-2 ip strict-source-route-option
```



NOTE: Currently, this screen option supports IPv4 only.

3. Enable the screens in the security zones.

```
[edit]
user@host# set security zones security-zone zone-1 screen screen-1
user@host# set security zones security-zone zone-2 screen screen-2
```

4. If you are done configuring the device, commit the configuration.

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

Verification

IN THIS SECTION

● [Verifying the Screens in the Security Zone](#) | 194

● Verifying the Security Screen Configuration | 194

Confirm that the configuration is working properly.

Verifying the Screens in the Security Zone

Purpose

Verify that the screen is enabled in the security zone.

Action

From operational mode, enter the `show security zones` command.

```
[edit]
user@host> show security zones

Security zone: zone-1
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-1
  Interfaces bound: 1
  Interfaces:
    ge-1/0/0.0
Security zone: zone-2
  Send reset for non-SYN session TCP packets: Off
  Policy configurable: Yes
  Screen: screen-2
  Interfaces bound: 1
  Interfaces:
    ge-2/0/0.0
```

Verifying the Security Screen Configuration

Purpose

Display the configuration information about the security screen.

Action

From operational mode, enter the `show security screen ids-option screen-name` command.

```
[edit]
user@host> show security screen ids-option screen-1
Screen object status:

Screen object status:

Name                                     Value
IP loose source route option            enabled
```

```
[edit]
user@host> show security screen ids-option screen-2
Screen object status:

Screen object status:

Name                                     Value
IP strict source route option           enabled
```

RELATED DOCUMENTATION

[Reconnaissance Deterrence Overview | 142](#)

[IP Address Sweep and Port Scan | 142](#)

5

CHAPTER

Configuration Statements and Operational Commands

[Junos CLI Reference Overview](#) | 197

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](#)
- [Operational Commands](#)