



Threat Detection

Threat detection's portscan detector is a mechanism designed to help you detect and prevent portscan activity in all types of traffic to protect networks from eventual attacks. Portscan traffic can be detected efficiently in both allowed and denied traffic.

Portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker determines the types of network protocols or services a host supports and sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

- [Portscan detection and prevention, on page 1](#)
- [Requirements and prerequisites for threat detection, on page 3](#)
- [Guidelines and limitations for threat detection, on page 4](#)
- [Best practices for portscan prevention, on page 4](#)
- [Configure portscan detection and prevention, on page 5](#)
- [Monitoring threat detection, on page 7](#)
- [History for threat detection, on page 9](#)

Portscan detection and prevention

Use threat detection to identify port scan activity. You can use the system to detect port scans and issue events when they are found. Optionally, you can configure the system to prevent port scans by automatically blocking scanners. When preventing port scans, the system sends you events and also blocks the attacker for a duration period that you set.

Pre-defined sensitivity levels for portscan detection

When configuring detection settings, you select from the following pre-defined sensitivity levels. Except for custom, each level has pre-set values for each protocol for the number of ports (TCP/UDP), protocols (IP), or hosts (TCP/UDP/IP/ICMP) that must be scanned within a set time interval (in seconds). Also, all types of scan/sweep are enabled.



Note When counting ports/protocols, threat detection increments the number if the port/protocol in the current packet differs from the previous packet. For example, if you have an application that opens connections in 10 set ports randomly, the total number of ports scanned could mount so quickly that your port number will be exceeded within the interval. The system does not count unique ports only.

Exceeding the number within the interval can indicate a scanning attack. Portscan events are generated only when the port/protocol/host numbers are exceeded for the moving time interval window.

- **Low**—This level uses the shortest time window for portscan detection, coupled with high counts for port/protocol/host. Thus, you should see portscan events for the most aggressive scanners only. Select this sensitivity level to suppress false positives, but remember that some types of portscans (slow scans, filtered scans) might be missed. For more detail on how low sensitivity detection works, see [Detection in the low sensitivity level, on page 3](#).
 - **Interval** (TCP/UDP/IP/ICMP)—60 seconds.
 - **TCP/UDP portscan Number of Ports**—120.
 - **TCP/UDP portsweep Number of Hosts**—180.
 - **IP protocol scan Number of Protocols**—30.
 - **IP protocol sweep Number of Hosts**—25.
 - **ICMP host sweep Number of Hosts**—50.
- **Medium**—This level uses moderate values for both the interval and port/protocol/host counts. However, very active hosts such as network address translators and proxies might generate false positives. Add such hosts to the ignore scanner list. This is the default sensitivity level and a good place to start.
 - **Interval** (TCP/UDP/IP/ICMP)—90 seconds.
 - **TCP/UDP portscan Number of Ports**—90.
 - **TCP/UDP portsweep Number of Hosts**—150.
 - **IP protocol scan Number of Protocols**—15.
 - **IP protocol sweep Number of Hosts**—20.
 - **ICMP host sweep Number of Hosts**—30.
- **High**—This level uses a much longer time window for portscan detection, coupled with lower counts for port/protocol/host. With this level, you are most likely to see events for even the least aggressive port scans/sweeps, so you are more likely to notice all attackers. On the other hand, this level would likely result in the most portscan events issued, and potentially the highest number of false positives.
 - **Interval** (TCP/UDP/IP/ICMP)—600 seconds (10 minutes).
 - **TCP/UDP portscan Number of Ports**—60.
 - **TCP/UDP portsweep Number of Hosts**—100.
 - **IP protocol scan Number of Protocols**—10.
 - **IP protocol sweep Number of Hosts**—10.

- **ICMP host sweep Number of Hosts**—20.
- **Custom**—If you want to configure any setting differently than one of the pre-defined sensitivity levels, or disable a particular type of scan/sweep, the level automatically switches to custom. If you want to adjust the options, first select the level that most closely matches what you want, then edit the values as appropriate.

Detection in the low sensitivity level

If you select the low sensitivity level, the system tracks negative responses for TCP, UDP, and ICMP initial packets. Only if the number of unsuccessful connections is more than the rejection threshold (10% in low sensitivity) and the port/IP protocol count is more than the configured threshold, is an alert triggered. This mitigates false positives.

Rejection threshold applies to low sensitivity (or equivalent custom settings) only; it does not apply to other sensitivity levels or their custom equivalents.

If there is a mix of allowed and blocked traffic, the number of rejected ports or hosts is calculated based on the difference between allowed and blocked traffic. In the case of only blocked traffic, the rejection threshold is not considered.

These criteria are not used for UDP/ICMP connections on interfaces configured in inline sets.

For example, in low sensitivity mode, the port count threshold is 120. Thus, the rejection count threshold is 10% of 120, which is 12. Following are examples of how the system would issue portscan events under this configuration:

- An attacker initiates connections with 131 ports of the target and the target positively acknowledges all the initiations. Port count = 131, which is greater than the threshold, but a portscan alert is not triggered because there are no negative acknowledgements.
- An attacker initiates connections with 131 ports of the target and the target positively acknowledges 121 initiations and negatively acknowledges 10 initiations. Port count = 131, which is greater than the threshold, but reject port count = 10, which is lesser than the rejection threshold. Therefore, a portscan alert is not triggered.
- An attacker initiates connections with 134 ports of the target and the target positively acknowledges 121 initiations and negatively acknowledges 13 initiations. Port count = 134, which is greater than the threshold, reject port count = 13 is also higher than the rejection threshold. Therefore a portscan alert is triggered.

Requirements and prerequisites for threat detection

Model Support

Threat Defense running version 7.2+ and Snort 3.

Supported Domains

Any

User Roles

Admin

Access Admin

Network Admin

Guidelines and limitations for threat detection

- Threat detection requires Snort 3. The NAP portscan configuration is always ignored for a device running Snort 3; you must configure portscan using threat detection. For Snort 2, you can configure port scan through the NAP policy only. If there are Snort 2 devices assigned to the access control policy, the threat detection settings will not be deployed to those unsupported devices.
- Threat detection works on traffic that passes through the device only. It does not work on traffic directed to the device.
- In a high availability setup, port scanning statistics are not synchronized to the standby unit. However, blocked hosts are synchronized and continue to be blocked until the duration period expires in case of a failover.
- (Devices running Threat Defense versions 7.2-7.7). For nodes in a cluster, detection and prevention happen on the individual cluster node. That is, if node B detects and blocks traffic from a host, node A will not be aware of that action because port scan statistics are not synchronized across cluster nodes.
- (Devices running Threat Defense version 10.0+). For nodes in a cluster, detection and prevention happen at the cluster level. Portscans can be detected when they happen across nodes or in an individual node. Shunned hosts are shunned on all devices in the cluster. Shuns are released at the same time on all nodes. Statistics are available at the cluster level.
- For inline sets, or for interfaces that are configured to be part of an equal-cost multipath (ECMP) traffic zone, detection and prevention are done at the zone level. Portscan statistics for a host are accumulated across all interfaces of a zone. Similarly, when a host crosses configured thresholds, it is blocked across all interfaces of the corresponding zone.
- Although the portscan events generated by the threat detection feature are the same as the ones Snort issues for port scan, you do not need to enable port scanning intrusion rules to get the events. Threat detection works regardless of your intrusion policy implementation.

Best practices for portscan prevention

Portscan prevention mode can result in unintended traffic outage. In prevention mode, hosts are blocked from further scanning of networks on all protocols for the configured duration. Review the detection and prevention parameters carefully to ensure legitimate traffic is not blocked.

Before configuring portscan in prevention mode, we strongly recommend the following:

1. Start using portscan in detection mode.
2. Observe the generated portscan events.

3. Tune the sensitivity level, and monitored networks, ignore scanner list, and ignore target list. If a pre-defined sensitivity level does not work well for your situation, configure custom settings as needed.
4. Repeat the process until false positives are eliminated and the event rate represents an accurate picture of port scanning in your network. Ensure that you are comfortable with blocking the remaining identified scanners.

Configure portscan detection and prevention

Portscan is a form of network reconnaissance that is often used by attackers as a prelude to an attack. In a portscan, an attacker determines the types of network protocols or services a host supports and sends specially crafted packets to a targeted host. By examining the packets that the host responds with, the attacker can often determine which ports are open on the host and, either directly or by inference, which application protocols are running on these ports.

You can enable threat detection to watch for port scanning activity and optionally, automatically block scanners for a period of time.

Before you begin

FQDN, wildcard mask, any, any-ipv4, and any-ipv6 network objects are not supported for portscan configuration. These objects are not shown in the **Monitor**, **Ignore Scanner**, **Ignore Target**, and **Exclude** fields.

Procedure

-
- | | |
|---------------|---|
| Step 1 | In the access control policy editor, click Advanced Settings from the More drop-down arrow at the end of the packet flow line. Then, click Edit (✎) next to Threat Detection . |
| Step 2 | <p>In the Threat Detection window, select the Portscan mode:</p> <ul style="list-style-type: none">• Disable—Turn off threat detection. This is the default mode. You can click Revert to Defaults to return to this unconfigured state.• Detection—Perform portscan detection, but alert on problems only. Do not take action against potential attackers. We suggest you use this mode initially until you fine-tune the threat detection settings to avoid excessive false positives.• Prevention—Perform portscan detection and actively block identified scanners, that is, hosts that are performing the port scan. |
| Step 3 | <p>Configure the Traffic Selection options.</p> <p>The traffic selection options determine which networks are monitored, the type of connections monitored, and whether any scanners or target hosts should be exempted from the monitored networks. By default, the system monitors permitted connections on all networks.</p> <ul style="list-style-type: none">• Detection On Traffic—Select the types of connection that will be monitored for portscan activity: Permitted, Denied, or All traffic. The default is Permitted.• Monitor—Select the network objects that define the networks to monitor for portscan or sweep activity. The default is any network, IPv4 or IPv6. Use this option to limit scanning to untrusted networks. |

- **Ignore Scanner**—Select the network objects that define the hosts or networks, from within the range of the monitored networks, that should be ignored. For example, if you have set up your own scanner to test your network, you can exempt the address of your scanner to avoid unnecessary reporting on the address. Do not include addresses that are outside the monitored networks, as these addresses are already ignored.
- **Ignore Target**—Select the network objects that define the hosts or networks that should be ignored as targets, that is, victims of a portscan or sweep.

Step 4 Click the **Configuration** tab and select the scanning sensitivity level.

The pre-defined sensitivity levels, **Low**, **Medium**, and **High**, set the port scanning options to values that are increasingly aggressive. For example, if you select low, you would expect to see fewer port scanning events, and you could potentially miss attackers more easily than if you selected medium or high. On the other hand, if you select high, you might see more events and also potentially more false positives. The default level is medium. For more information on these levels, see [Pre-defined sensitivity levels for portscan detection, on page 1](#).

As you select the levels, you can see the related values within the protocol sections: **TCP**, **UDP**, **IP**, and **ICMP**. If you change any of the preset values, or disable a type of scan, the sensitivity mode automatically changes to **Custom**.

Within each protocol section, the options are:

- **Interval**—The time range, in seconds, within which the configured values for portscan or portsweep are exceeded. For example, if you select 90 seconds, and 60 as the number of TCP portscan ports, a scanner would need to try 60 ports on a host within 90 seconds for it to be considered a portscan. The system generates events only if the number of ports, protocols, or hosts (for a portsweep) are exceeded within the specified interval.

You can specify a range between 30-600 seconds. The longer the period, the more likely a host might be identified as a scanner.
- **Portscan (TCP/UDP)**—Select whether to monitor for port scanning against single hosts, and specify the number of ports that must be scanned within the interval to count as a portscan attack. The allowed range is 1-256.
- **Portsweep (TCP/UDP)**—Select whether to monitor for port sweeping against multiple hosts, and specify the number of hosts that must be scanned for a given port within the interval to count as a portsweep attack. The allowed range is 1-256.
- **Protocol Scan (IP)**—Select whether to monitor for protocol scanning against single hosts, and specify the number of protocols that must be scanned within the interval to count as a protocol scan attack. The allowed range is 1-255.
- **Protocol Sweep (IP)**—Select whether to monitor for protocol sweeping against multiple hosts, and specify the number of hosts that must be scanned for a given protocol within the interval to count as a protocol sweep attack. The allowed range is 1-256.
- **Hostsweep (ICMP)**—Select whether to monitor for ICMP host sweeping against multiple hosts, and specify the number of hosts that must be scanned within the interval to count as a hostsweep attack. The allowed range is 1-256.

Step 5 If you selected prevention mode, click the **Prevention** tab and configure the options.

In prevention mode, hosts are automatically blocked from further scanning of networks on all protocols for the configured duration. Review the detection and prevention parameters carefully to ensure legitimate traffic is not blocked.

- **Exclude**—Select the network objects that define the hosts or networks, from within the range of the monitored networks, that should be excluded from automatic blocking. Even if these hosts violate your scanning detection parameters, the system will not block them.
- **Duration**—How long, in seconds, automatically blocked scanner hosts should be prevented from sending traffic of any kind through the device. After the duration period ends, the hosts are automatically cleared and can again send traffic through the device. The allowed range is 600-2592000 seconds. The default is 3600 seconds (1 hour).

If you need to manually unblock a host, SSH to the firewall that is blocking the host and use the **clear threat-detection portscan attacker** command.

Step 6 Click **OK** to save the threat detection settings.

Step 7 Click **Save** to save the access control policy.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#).

Monitoring threat detection

The following topics explain how to monitor portscan activity

Viewing portscan alerts

Portscan activity is alerted through the existing portscan-specific intrusion events. Intrusion events with generator ID (GID) 122 and Snort ID from SIDs 1 through 27 are generated. For these events, the (*port_scan*) string is prepended in the event messages. The events include packet information along with packet data containing the statistics that triggered the alert.

To see portscan events, go to **Events & Logs** > + **Show more** > **Intrusion** > **Events**.

Portscan issues these events regardless of your intrusion policy or NAP configuration. Events are issued only when scanners exceed the number of configured ports/protocols/hosts for the various types of scan or sweep within the configured time interval for the associated protocol. A port scan from one host generates one event per set interval as soon as the threshold is met. If the same host initiates a new port scan during the same interval, no event is reported.

The following table shows the possible events.

Table 1: Portscan events

Portscan type	Intrusion event
TCP Regular, Decoy, Distributed Scan	122:1 (port_scan) TCP portscan
TCP Portsweep	122:3 (port_scan) TCP portsweep

Portscan type	Intrusion event
TCP Distributed Scan	122:4 (port_scan) TCP distributed portscan
IP Regular, Decoy, Distributed Protocol Scan n	122:9 (port_scan) IP protocol scan
IP Protocol Sweep	122:11 (port_scan) IP protocol sweep
IP Distributed Scan	122:12 (port_scan) IP distributed protocol scan
UDP Regular, Decoy, Distributed Scan	122:17 (port_scan) UDP portscan
UDP Portsweep	122:19 (port_scan) UDP portsweep
UDP Distributed Scan	122:20 (port_scan) UDP distributed portscan
ICMP Sweep	122:25 (port_scan) ICMP sweep
Portscan Block	122:100 (port_scan) host blocked due to portscan activity

Monitoring portscan on the firewall

To monitor portscan, log into the device CLI and use the following commands.

- **show threat-detection portscan** [attacker | target | shun]

Shows the IP addresses of scanners, those that have been shunned (blocked), and hosts that have been targeted by scans or sweeps.

- **show threat-detection portscan statistics** [host [ipv4_address | ipv6_address]] [protocol {tcp | udp | ip | icmp}]

Shows statistics related to the portscan system. You can specify host, protocol, or host and protocol to filter the output to the desired information.

- **clear threat-detection portscan** [attacker | target | shun] [ipv4_address mask | ipv6_address/prefix]

Manually unblocks scanners (attackers) or identified targets. Enter the command without parameters to clear all attackers, targets, or shunned hosts.

- **clear threat-detection portscan statistics** [host [ipv4_address | ipv6_address]] [protocol {tcp | udp | ip | icmp}]

Erases statistics related to portscan, so that you can more clearly see the current state of scanning through this device. Enter the command without parameters to clear all statistics. Alternatively, specify a host, protocol, or host and protocol, to limit the reset to the specified items.

Unblocking a host

If you configure threat detection in prevention mode, and the system blocks a host that you know is not an attacker, you can manually unblock the host before host is automatically unblocked when the duration period expires.

To manually unblock the host, log into the device CLI where the host is blocked and enter the **clear threat-detection portscan attacker** command. For example:

```
> clear threat-detection portscan attacker 10.2.0.100 255.255.255.255
1 tracker object deleted and 1 shun entry removed
```

Consider adding the host IP to the exclude list in the prevention configuration.

History for threat detection

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Description
Threat detection support for clustering.	10.0	10.0	<p>You can now configure threat detection at the cluster level. For nodes in a cluster, detection and prevention happen at the cluster level. Portscans can be detected when they happen across nodes or in an individual node. Shunned hosts are shunned on all devices in the cluster. Shuns are released at the same time on all nodes. Statistics are available at the cluster level.</p> <p>There are no changes in how threat detection is configured.</p>
Improved detection in low sensitivity.	7.4	7.4	<p>The method for identifying port scans and sweeps has been improved when operating in low sensitivity level. Changes are automatic; there are no new configuration settings.</p>
Improved portscan detection.	7.2	7.2 running Snort 3	<p>With an improved portscan detector, you can easily configure the system to detect or prevent portscans. You can refine the networks you want to protect, set the sensitivity, and so on. For devices running Snort 2 and for devices running Version 7.1 and earlier, continue to use the network analysis policy for portscan detection.</p> <p>New/modified screens: We added Threat Detection to the access control policy's Advanced tab.</p> <p>New/Modified commands: clear threat-detection portscan, show threat-detection portscan.</p>

