



Cisco Secure Firewall Management Center Snort 3 Configuration Guide, Version 7.2

First Published: 2022-06-06

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2022 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1

An Overview of Network Analysis and Intrusion Policies 1

- About Network Analysis and Intrusion Policies 1
- Snort Inspection Engine 2
- Snort 3 2
- Guidelines and Limitations for Network Analysis and Intrusion Policies 4
- How Policies Examine Traffic For Intrusions 5
 - Decoding, Normalizing, and Preprocessing: Network Analysis Policies 6
 - Access Control Rules: Intrusion Policy Selection 7
 - Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets 8
 - Intrusion Event Generation 9
- System-Provided and Custom Network Analysis and Intrusion Policies 10
 - System-Provided Network Analysis and Intrusion Policies 11
 - Benefits of Custom Network Analysis and Intrusion Policies 12
 - Benefits of Custom Network Analysis Policies 12
 - Benefits of Custom Intrusion Policies 13
 - Limitations of Custom Policies 14
- Prerequisites for Network Analysis and Intrusion Policies 16

CHAPTER 2

Migrate from Snort 2 to Snort 3 17

- Snort 3 Inspection Engine 17
- Prerequisites for Network Analysis and Intrusion Policies 18
- How to Migrate from Snort 2 to Snort 3 18
 - Prerequisites for Migrating from Snort 2 to Snort 3 18
 - Enable Snort 3 on an Individual Device 19
 - Enable Snort 3 on Multiple Devices 19
 - Convert Snort 2 Custom IPS Rules to Snort 3 20

Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3	20
Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3	21
View Snort 2 and Snort 3 Base Policy Mapping	22
Synchronize Snort 2 Rules with Snort 3	22
Deploy Configuration Changes	23

PART I

Intrusion Detection and Prevention in Snort 3 27

CHAPTER 3

Get Started with Snort 3 Intrusion Policies 29

Overview of Intrusion Policies	29
Prerequisites for Network Analysis and Intrusion Policies	30
Create a Custom Snort 3 Intrusion Policy	30
Edit Snort 3 Intrusion Policies	31
Rule Action Logging	31
Change the Base Policy of an Intrusion Policy	32
Manage Intrusion Policies	32
Access Control Rule Configuration to Perform Intrusion Prevention	33
Access Control Rule Configuration and Intrusion Policies	34
Configure an Access Control Rule to Perform Intrusion Prevention	34
Deploy Configuration Changes	35

CHAPTER 4

Tune Intrusion Policies Using Rules 37

Overview of Tuning Intrusion Rules	37
Intrusion Rule Types	38
Prerequisites for Network Analysis and Intrusion Policies	39
Custom Rules in Snort 3	39
View Snort 3 Intrusion Rules in an Intrusion Policy	40
Intrusion Rule Action	40
Intrusion Rule Action Options	40
Set Intrusion Rule Action	41
Intrusion Event Notification Filters in an Intrusion Policy	42
Intrusion Event Thresholds	42
Set Intrusion Event Thresholds	42
Set Threshold for an Intrusion Rule in Snort 3	43

View and Delete Intrusion Event Thresholds	44
Intrusion Policy Suppression Configuration	44
Intrusion Policy Suppression Types	45
Set Suppression for an Intrusion Rule in Snort 3	45
View and Delete Suppression Conditions	46
Add Intrusion Rule Comments	46
Snort 2 Custom Rules Conversion to Snort 3	47
Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3	47
Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3	48
Add Custom Rules to Rule Groups	49
Add Rule Groups with Custom Rules to an Intrusion Policy	50
Manage Custom Rules in Snort 3	50
Delete Custom Rules	51
Delete Rule Groups	52

CHAPTER 5

Tailor Intrusion Protection for Your Network Assets	53
Snort 3 Rule Changes in LSP Updates	53
Overview of Secure Firewall Recommended Rules	53
Prerequisites for Network Analysis and Intrusion Policies	54
Generate New Secure Firewall Recommendations in Snort 3	55

PART II

Advanced Network Analysis in Snort 3	59
---	-----------

CHAPTER 6

Get Started with Network Analysis Policies	61
Overview of Network Analysis Policies	61
Manage Network Analysis Policies	62
Snort 3 Definitions and Terminologies for Network Analysis Policy	62
Prerequisites for Network Analysis and Intrusion Policies	65
Custom Network Analysis Policy Creation for Snort 3	65
Network Analysis Policy Mapping	68
View Network Analysis Policy Mapping	68
Create a Network Analysis Policy	69
Modify the Network Analysis Policy	69
Search for an Inspector on the Network Analysis Policy Page	70

Copy the Inspector Configuration	70
Customize the Network Analysis Policy	71
Make Inline Edit for an Inspector to Override Configuration	74
Revert Unsavd Changes during Inline Edits	75
View the List of Inspectors with Overrides	75
Revert Overridden Configuration to Default Configuration	76
Validate Snort 3 Policies	77
Examples of Custom Network Analysis Policy Configuration	79
Network Analysis Policy Settings and Cached Changes	90

PART III Encrypted Visibility Engine for Snort 3 91

CHAPTER 7	Encrypted Visibility Engine 93
	Overview of Encrypted Visibility Engine 93
	How EVE Works 94
	Configure EVE 95
	View EVE Events 95
	View EVE Dashboard 96
	Configure EVE Exception Rules 96
	Add Exception Rule from Unified Events 97

PART IV Elephant Flow Detection for Snort 3 99

CHAPTER 8	Elephant Flow Detection 101
	About Elephant Flow Detection and Remediation 101
	Elephant Flow Upgrade from Intelligent Application Bypass 101
	Configure Elephant Flow 102



CHAPTER 1

An Overview of Network Analysis and Intrusion Policies

The Snort inspection engine is an integral part of the Secure Firewall Threat Defense (formerly Firepower Threat Defense) device. This chapter provides an overview of Snort 3 and the network analysis and intrusion policies. It also provides an insight into system-provided and custom network analysis and intrusion policies.

- [About Network Analysis and Intrusion Policies, on page 1](#)
- [Snort Inspection Engine, on page 2](#)
- [Snort 3, on page 2](#)
- [Guidelines and Limitations for Network Analysis and Intrusion Policies, on page 4](#)
- [How Policies Examine Traffic For Intrusions, on page 5](#)
- [System-Provided and Custom Network Analysis and Intrusion Policies, on page 10](#)
- [Prerequisites for Network Analysis and Intrusion Policies, on page 16](#)

About Network Analysis and Intrusion Policies

Network analysis and intrusion policies work together as part of the intrusion detection and prevention feature.

- The term *intrusion detection* generally refers to the process of passively monitoring and analyzing network traffic for potential intrusions and storing attack data for security analysis. This is sometimes referred to as "IDS."
- The term *intrusion prevention* includes the concept of intrusion detection, but adds the ability to block or alter malicious traffic as it travels across your network. This is sometimes referred to as "IPS."

In an intrusion prevention deployment, when the system examines packets:

- A **network analysis policy** governs how traffic is *decoded* and *preprocessed* so it can be further evaluated, especially for anomalous traffic that might signal an intrusion attempt.
- An **intrusion policy** uses *intrusion and preprocessor rules* (sometimes referred to collectively as *intrusion rules*) to examine the decoded packets for attacks based on patterns. Intrusion policies are paired with *variable sets*, which allow you to use named values to accurately reflect your network environment.

Both network analysis and intrusion policies are invoked by a parent access control policy, but at different times. As the system analyzes traffic, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (additional preprocessing and intrusion rules) phase. Together,

network analysis and intrusion policies provide broad and deep packet inspection. They can help you detect, alert on, and protect against network traffic that could threaten the availability, integrity, and confidentiality of hosts and their data.

The system is delivered with several similarly named network analysis and intrusion policies (for example, Balanced Security and Connectivity) that complement and work with each other. By using system-provided policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and inspector rule states, as well as provides the initial configurations for inspectors and other advanced settings.

You can also create custom network analysis and intrusion policies. You can tune settings in custom policies to inspect traffic in the way that matters most to you so that you can improve both the performance of your managed devices and your ability to respond effectively to the events they generate.

You create, edit, save, and manage network analysis and intrusion policies using similar policy editors in the web interface. When you are editing either type of policy, a navigation panel appears on the left side of the web interface; the right side displays various configuration pages.

Refer to the videos for additional support and information:

- [Snort 3 Condensed Overview](#)
- [Snort 3 Extended Overview](#)

Snort Inspection Engine

The Snort inspection engine is an integral part of the Secure Firewall Threat Defense (formerly Firepower Threat Defense) device. The inspection engine analyzes traffic in real time to provide deep packet inspection. Network analysis and intrusion policies together utilize the Snort inspection engine's capabilities to detect and protect against intrusions.

Snort 3

Snort 3 is the latest version of the Snort inspection engine, which has vast improvements compared to the earlier version of Snort. The older version of Snort is Snort 2. Snort 3 is more efficient, and it provides better performance and scalability.

Snort 3 is architecturally redesigned to inspect more traffic with equivalent resources when compared to Snort 2. Snort 3 provides simplified and flexible insertion of traffic parsers. Snort 3 also provides new rule syntax that makes rule writing easier and shared object rule equivalents visible.

The other significant changes with Snort 3 are:

- Unlike Snort 2, which uses multiple Snort instances, Snort 3 associates multiple threads with a single Snort instance. This uses less memory, improves Snort reload times, and supports more intrusion rules and a larger network map. The number of Snort threads varies by platform and is the same as the number of Snort 2 instances for each platform. Usage is virtually transparent.
- Snort version per threat defense—The Snort inspection engine is threat defense specific and not Secure Firewall Management Center (formerly Firepower Management Center) specific. Management Center can manage several threat defenses, each with either versions of Snort (Snort 2 and Snort 3). Although the management center's intrusion policies are unique, the system applies Snort 2 or Snort 3 version of an intrusion policy for intrusion protection depending on the device's selected inspection engine.

- Decoder rules—Packet decoder rules fire only in the default intrusion policy. The system ignores decoder rules that you enable in other policies.
- Shared object rules—Snort 3 supports some but not all shared object (SO) intrusion rules (rules with a generator ID (GID) of 3). Enabled shared object rules that are not supported do not trigger.
- Multi-layer inspection for Security Intelligence—Snort 3 detects the innermost IP address regardless of the layer.
- Platform support—Snort 3 requires threat defense 7.0 or later. It is not supported with ASA FirePOWER or NGIPSv.
- Managed Devices—An management center with version 7.0 can simultaneously support version 6.4, 6.5, 6.6, 6.7, and 7.0 Snort 2 threat defenses, and version 7.0 Snort 3 threat defenses.
- Traffic interruption when switching Snort versions—Switching Snort versions interrupts traffic inspection and a few packets might drop during deployment.
- Unified policies—Irrespective of the underlying Snort engine version that is enabled in the managed threat defenses, the access control policies, intrusion policies, and network analysis policies configured in the management center work seamlessly in applying the policies. All intrusion policies in management center version 7.0 and above have two versions available, Snort 2 version and Snort 3 version. The intrusion policy is unified, which means that it has a common name, base policy, and inspection mode, although there are two versions of the policy (Snort 2 version and Snort 3 version). The Snort 2 and the Snort 3 versions of the intrusion policy can be different in terms of the rule settings. However, when the intrusion policy is applied on a device, the system automatically identifies the Snort version enabled on the device and applies the rule settings configured for that version.
- Lightweight Security Package (LSP)—Replaces the Snort Rule Updates (SRU) for Snort 3 next-generation intrusion rule and configuration updates. Downloading updates downloads both the Snort 3 LSP and the Snort 2 SRU.

LSP updates provide new and updated intrusion rules and inspector rules, modified states for existing rules, and modified default intrusion policy settings for management center and threat defense versions 7.0 or above. When you upgrade an management center from version 6.7 or lower to 7.0, it supports both LSPs and SRUs. LSP updates may also delete system-provided rules, provide new rule categories and default variables, and modify default variable values. For more information on LSP updates, see the *Update Intrusion Rules* topic in the latest version of the *Firepower Management Center Configuration Guide*.
- Mapping of Snort 2 and Snort 3 rules and presets—Snort 2 and Snort 3 rules are mapped and the mapping is system-provided. However, it is not a one-to-one mapping. The system-provided intrusion base policies are pre-configured for both Snort 2 and Snort 3, and they provide the same intrusion prevention although with different rule sets. The system-provided base policies for Snort 2 and Snort 3 are mapped with each other for the same intrusion prevention settings. For more information, see [View Snort 2 and Snort 3 Base Policy Mapping, on page 22](#).
- Synchronizing Snort 2 and Snort 3 rule override—When a threat defense is upgraded to 7.0, you can upgrade the inspection engine of the threat defense to the Snort 3 version. Management Center maps all the overrides in the existing rules of the Snort 2 version of the intrusion policies to the corresponding Snort 3 rules using the mapping provided by Talos. However, if there are additional overrides performed after the upgrade or if you have installed a new threat defense of version 7.0, they have to be manually synchronized. For more information, see [Synchronize Snort 2 Rules with Snort 3 , on page 22](#).

- Custom intrusion rules—You can create custom intrusion rules in Snort 3. You can also import the custom intrusion rules that exist for Snort 2 to Snort 3. For more information, see [Custom Rules in Snort 3, on page 39](#).
- Switching between Snort 2 and Snort 3 engines—threat defenses that support Snort 3 can also support Snort 2. Switching from Snort 3 to Snort 2 is not recommended from the efficacy perspective.

**Important**

Although you can switch Snort versions freely, intrusion rule changes in one version of Snort will not be updated in the other version automatically. If you change the rule action for a rule in one version of Snort, ensure you replicate the change in the other version before switching the Snort version. System provided synchronization option only synchronizes the changes in the Snort 2 version of the intrusion policy to the Snort 3 version, and not the other way around.

Guidelines and Limitations for Network Analysis and Intrusion Policies

- A high percentage of traffic with small packets causes Snort performance to decrease. This behaviour is observed even when all the preprocessors are disabled.
- When you attempt to deploy a configuration change on a threat defense device with low memory, snort deployment is also triggered. This results in high consumption of RSS memory. Snort memory usage is also impacted if you deploy large configurations on the device, such as multiple IPS policies containing a large number of snort IPS rules, network objects, and access-control lists. You can mitigate such memory issues by optimizing the configuration. For best practices on how to configure access control rules to optimize the configuration, see [Best Practices for Access Control Rules](#).
- If you increase the memory of a Threat Defense Virtual instance, you must redeploy the configuration for Snort 3 to utilize the additional memory.

**Note**

The Snort 3 memory allocation is not automatically adjusted when you increase the memory of the Threat Defense Virtual instance. You must redeploy the configuration to regenerate relevant configuration files, such as `memory_allocation.lua`, which apply the updated resource limits to Snort 3.

Feature Limitations of Snort 3 for Management Center-Managed Threat Defense

The following table lists the features that are supported on Snort 2 but not supported on Snort 3 for management center-managed threat defense devices.

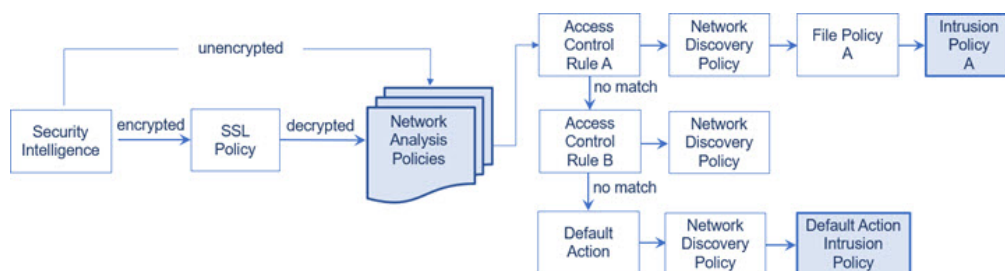
Table 1: Feature Limitations of Snort 3

Policy/Area	Features not supported
Access Control Policy	The following application settings: <ul style="list-style-type: none"> • Safe Search • YouTube EDU
Intrusion Policy	<ul style="list-style-type: none"> • Policy layers • Global rule thresholding • Logging configuration: <ul style="list-style-type: none"> • SNMP • SRU rule updates as Snort 3 supports only LSP rule updates
Other features	Event logging with FQDN names

How Policies Examine Traffic For Intrusions

When the system analyzes traffic as part of your access control deployment, the network analysis (decoding and preprocessing) phase occurs before and separately from the intrusion prevention (intrusion rules and advanced settings) phase.

The following diagram shows, in a simplified fashion, the order of traffic analysis in an inline, intrusion prevention and AMP for Networks deployment. It illustrates how the access control policy invokes other policies to examine traffic, and in which order those policies are invoked. The network analysis and intrusion policy selection phases are highlighted.



In an inline deployment (that is, where relevant configurations are deployed to devices using routed, switched, or transparent interfaces, or inline interface pairs), the system can block traffic without further inspection at almost any step in the illustrated process. Security Intelligence, the SSL policy, network analysis policies, file policies, and intrusion policies can all either drop or modify traffic. Only the network discovery policy, which passively inspects packets, cannot affect the flow of traffic.

Similarly, at each step of the process, a packet could cause the system to generate an event. Intrusion and preprocessor events (sometimes referred to collectively as *intrusion events*) are indications that a packet or its contents may represent a security risk.



Tip The diagram does not reflect that access control rules handle encrypted traffic when your SSL inspection configuration allows it to pass, or if you do not configure SSL inspection. By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured.

Note that for a single connection, although the system selects a network analysis policy before an access control rule as shown in the diagram, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

Decoding, Normalizing, and Preprocessing: Network Analysis Policies

Without decoding and preprocessing, the system could not appropriately evaluate traffic for intrusions because protocol differences would make pattern matching impossible. Network analysis policies govern these traffic-handling tasks:

- **after** traffic is filtered by Security Intelligence
- **after** encrypted traffic is decrypted by an optional SSL policy
- **before** traffic can be inspected by file or intrusion policies

A network analysis policy governs packet processing in phases. First the system decodes packets through the first three TCP/IP layers, then continues with normalizing, preprocessing, and detecting protocol anomalies:

- The packet decoder converts packet headers and payloads into a format that can be easily used by the inspectors and later, intrusion rules. Each layer of the TCP/IP stack is decoded in turn, beginning with the data link layer and continuing through the network and transport layers. The packet decoder also detects various anomalous behaviors in packet headers.
- In inline deployments, the inline normalization preprocessor reformats (normalizes) traffic to minimize the chances of attackers evading detection. It prepares packets for examination by other inspectors and intrusion rules, and helps ensure that the packets the system processes are the same as the packets received by the hosts on your network.
- Various network and transport layers inspectors detect attacks that exploit IP fragmentation, perform checksum validation, and perform TCP and UDP session preprocessing.

Note that some advanced transport and network inspector settings apply globally to all traffic handled by the target devices of an access control policy. You configure these in the access control policy rather than in a network analysis policy.

- Various application-layer protocol decoders normalize specific types of packet data into formats that the intrusion rules engine can analyze. Normalizing application-layer protocol encodings allows the system to effectively apply the same content-related intrusion rules to packets whose data is represented differently, and to obtain meaningful results.
- The Modbus, DNP3, CIP, and s7commplus SCADA inspectors detect traffic anomalies and provide data to intrusion rules. Supervisory Control and Data Acquisition (SCADA) protocols monitor, control, and acquire data from industrial, infrastructure, and facility processes such as manufacturing, production, water treatment, electric power distribution, airport and shipping systems, and so on.

- Several inspectors allow you to detect specific threats, such as Back Orifice, portscans, SYN floods and other rate-based attacks.

Note that you configure the sensitive data inspector, which detects sensitive data such as credit card numbers and Social Security numbers in ASCII text, in intrusion policies.



Note When TLS Server Identity is disabled, Snort 3 does not perform SNI mismatch detection. It only evaluates the SNI in the Client Hello packet and bypasses the validation of the certificate's Common Name (CN) in the Server Hello packet.

In a newly created access control policy, one default network analysis policy governs preprocessing for *all* traffic for *all* intrusion policies invoked by the same parent access control policy. Initially, the system uses the Balanced Security and Connectivity network analysis policy as the default, but you can change it to another system-provided or custom network analysis policy. In a more complex deployment, advanced users can tailor traffic preprocessing options to specific security zones, networks, and VLANs by assigning different custom network analysis policies to preprocess matching traffic.



Note For an access control policy with rule action as **Trust** and a prefilter rule with action as **Fastpath** with logging options disabled, you will observe that the end-of-flow events are still generated in the system. The events are not visible on the management center event pages.

Access Control Rules: Intrusion Policy Selection

After initial preprocessing, access control rules (when present) evaluate traffic. In most cases, the first access control rule that a packet matches is the rule that handles that traffic; you can monitor, trust, block, or allow matching traffic.

When you allow traffic with an access control rule, the system can inspect the traffic for discovery data, malware, prohibited files, and intrusions, in that order. Traffic not matching any access control rule is handled by the access control policy's default action, which can also inspect for discovery data and intrusions.



Note All packets, **regardless** of which network analysis policy preprocesses them, are matched to configured access control rules—and thus are potentially subject to inspection by intrusion policies—in top-down order.

The diagram in [How Policies Examine Traffic For Intrusions, on page 5](#) shows the flow of traffic through a device in an inline, intrusion prevention and AMP for Networks deployment, as follows:

- Access Control Rule A allows matching traffic to proceed. The traffic is then inspected for discovery data by the network discovery policy, for prohibited files and malware by File Policy A, and then for intrusions by Intrusion Policy A.
- Access Control Rule B also allows matching traffic. However, in this scenario, the traffic is not inspected for intrusions (or files or malware), so there are no intrusion or file policies associated with the rule. Note that by default, traffic that you allow to proceed is inspected by the network discovery policy; you do not need to configure this.

- In this scenario, the access control policy's default action allows matching traffic. The traffic is then inspected by the network discovery policy, and then by an intrusion policy. You can (but do not have to) use a different intrusion policy when you associate intrusion policies with access control rules or the default action.

The example in the diagram does not include any blocking or trusting rules because the system does not inspect blocked or trusted traffic.

Intrusion Inspection: Intrusion Policies, Rules, and Variable Sets

You can use intrusion prevention as the system's last line of defense before traffic is allowed to proceed to its destination. Intrusion policies govern how the system inspects traffic for security violations and, in inline deployments, can block or alter malicious traffic. The main function of intrusion policies is to manage which intrusion and preprocessor rules are enabled and how they are configured.

Intrusion and Inspector Rules

An intrusion rule is a specified set of keywords and arguments that detects attempts to exploit vulnerabilities on your network; the system uses an intrusion rule to analyze network traffic to check if it matches the criteria in the rule. The system compares packets against the conditions specified in each rule and, if the packet data matches all the conditions specified in a rule, the rule triggers.

The system includes the following types of rules created by Cisco Talos Intelligence Group (Talos):

- *shared object intrusion rules*, which are compiled and cannot be modified (except for rule header information such as source and destination ports and IP addresses)
- *standard text intrusion rules*, which can be saved and modified as new custom instances of the rule.
- *preprocessor rules*, which are rules associated with inspectors and packet decoder detection options in the network analysis policy. You cannot copy or edit inspector rules. Most inspector rules are disabled by default; you must enable them to use inspectors to generate events and, in an inline deployment, drop offending packets.

When the system processes packets according to an intrusion policy, first a rule optimizer classifies all activated rules in subsets based on criteria such as: transport layer, application protocol, direction to or from the protected network, and so on. Then, the intrusion rules engine selects the appropriate rule subsets to apply to each packet. Finally, a multi-rule search engine performs three different types of searches to determine if the traffic matches the rule:

- The protocol field search looks for matches in particular fields in an application protocol.
- The generic content search looks for ASCII or binary byte matches in the packet payload.
- The packet anomaly search looks for packet headers and payloads that, rather than containing specific content, violate well-established protocols.

In a custom intrusion policy, you can tune detection by enabling and disabling rules, as well as by writing and adding your own standard text rules. You can also use Cisco recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.



Note When there are insufficient packets to process specific traffic against a block rule, the system continues to evaluate the remaining traffic against other rules. If any of the remaining traffic matches a rule which is set to block, then the session is blocked. However, if the system analyses the remaining traffic to be passed, then the traffic status shows pending on the rule which is stuck for want of complete packets.

Variable Sets

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Most variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.

The system provides a single default variable set, which is comprised of predefined default variables. Most system-provided shared object rules and standard text rules use these predefined default variables to define networks and port numbers. For example, the majority of the rules use the variable `$HOME_NET` to specify the protected network and the variable `$EXTERNAL_NET` to specify the unprotected (or outside) network. In addition, specialized rules often use other predefined variables. For example, rules that detect exploits against web servers use the `$HTTP_SERVERS` and `$HTTP_PORTS` variables.



Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends that you modify key default variables in the default set. When you use variables that accurately reflect your network environment, processing is optimized and the system can monitor relevant systems for suspicious activity. Advanced users can create and use custom variable sets for pairing with one or more custom intrusion policies.



Important If you are creating a custom variable set, do not use a number as the first character in a custom variable set name (for example, 3Snort). This will cause Snort 3 validation to fail when you deploy a configuration to threat defense firewall on the management center.

Intrusion Event Generation

When the system identifies a possible intrusion, it generates an *intrusion or preprocessor event* (sometimes collectively called *intrusion events*). Managed devices transmit their events to the management center, where you can view the aggregated data and gain a greater understanding of the attacks against your network assets. In an inline deployment, managed devices can also drop or replace packets that you know to be harmful.

Each intrusion event in the database includes an event header and contains information about the event name and classification; the source and destination IP addresses; ports; the process that generated the event; and the date and time of the event, as well as contextual information about the source of the attack and its target. For packet-based events, the system also logs a copy of the decoded packet header and payload for the packet or packets that triggered the event.

The packet decoder, the preprocessors, and the intrusion rules engine can all cause the system to generate an event. For example:

- If the packet decoder (configured in the network analysis policy) receives an IP packet that is less than 20 bytes, which is the size of an IP datagram without any options or payload, the decoder interprets this

as anomalous traffic. If, later, the accompanying decoder rule in the intrusion policy that examines the packet is enabled, the system generates a inspector event.

- If the IP defragmentation preprocessor encounters a series of overlapping IP fragments, the inspector interprets this as a possible attack and, when the accompanying inspector rule is enabled, the system generates a inspector event.
- Within the intrusion rules engine, most standard text rules and shared object rules are written so that they generate intrusion events when triggered by packets.

As the database accumulates intrusion events, you can begin your analysis of potential attacks. The system provides you with the tools you need to review intrusion events and evaluate whether they are important in the context of your network environment and your security policies.

System-Provided and Custom Network Analysis and Intrusion Policies

Creating a new access control policy is one of the first steps in managing traffic flow using the system. By default, a newly created access control policy invokes system-provided network analysis and intrusion policies to examine traffic.

The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.



Note how:

- A default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided *Balanced Security and Connectivity network analysis policy* is the default.
- The default action of the access control policy allows all non-malicious traffic, as determined by the system-provided *Balanced Security and Connectivity intrusion policy*. Because the default action allows traffic to pass, the discovery feature can examine it for host, application, and user data before the intrusion policy can examine and potentially block malicious traffic.
- The policy uses default Security Intelligence options (global Block and Do Not Block lists only), does not decrypt encrypted traffic with an SSL policy, and does not perform special handling and inspection of network traffic using access control rules.

A simple step you can take to tune your intrusion prevention deployment is to use a different set of system-provided network analysis and intrusion policies as your defaults. Cisco delivers several pairs of these policies with the system.

Or, you can tailor your intrusion prevention deployment by creating and using custom policies. You may find that the inspector options, intrusion rule, and other advanced settings configured in those policies do not address the security needs of your network. By tuning your network analysis and intrusion policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

System-Provided Network Analysis and Intrusion Policies

Cisco delivers several pairs of network analysis and intrusion policies with the system. By using system-provided network analysis and intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos provides intrusion and inspector rule states as well as initial configurations for inspectors and other advanced settings.

No system-provided policy covers every network profile, traffic mix, or defensive posture. Each covers common cases and network setups that provide a starting point for a well-tuned defensive policy. Although you can use system-provided policies as-is, Cisco strongly recommends that you use them as the base for custom policies that you tune to suit your network.



Tip Even if you use system-provided network analysis and intrusion policies, you should configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify key default variables in the default set.

As new vulnerabilities become known, Talos releases intrusion rule updates also known as *Lightweight Security Package* (LSP). These rule updates can modify any system-provided network analysis or intrusion policy, and can provide new and updated intrusion rules and inspector rules, modified states for existing rules, and modified default policy settings. Rule updates may also delete rules from system-provided policies and provide new rule categories, as well as modify the default variable set.

If a rule update affects your deployment, the web interface marks affected intrusion and network analysis policies as out of date, as well as their parent access control policies. You must re-deploy an updated policy for its changes to take effect.

For your convenience, you can configure rule updates to automatically re-deploy affected intrusion policies, either alone or in combination with affected access control policies. This allows you to easily and automatically keep your deployment up-to-date to protect against recently discovered exploits and intrusions.

To ensure up-to-date preprocessing settings, you **must** re-deploy access control policies, which also deploys any associated SSL, network analysis, and file policies that are different from those currently running, and can also can update default values for advanced preprocessing and performance options.

Cisco delivers the following network analysis and intrusion policies with the system:

Balanced Security and Connectivity network analysis and intrusion policies

These policies are built for both speed and detection. Used together, they serve as a good starting point for most organizations and deployment types. The system uses the Balanced Security and Connectivity policies and settings as defaults in most cases.

Connectivity Over Security network analysis and intrusion policies

These policies are built for organizations where connectivity (being able to get to all resources) takes precedence over network infrastructure security. The intrusion policy enables far fewer rules than those enabled in the Security over Connectivity policy. Only the most critical rules that block traffic are enabled.

Security Over Connectivity network analysis and intrusion policies

These policies are built for organizations where network infrastructure security takes precedence over user convenience. The intrusion policy enables numerous network anomaly intrusion rules that could alert on or drop legitimate traffic.

Maximum Detection network analysis and intrusion policies

These policies are built for organizations where network infrastructure security is given even more emphasis than is given by the Security Over Connectivity policies, with the potential for even greater operational impact. For example, the intrusion policy enables rules in a large number of threat categories including malware, exploit kit, old and common vulnerabilities, and known in-the-wild exploits.

No Rules Active intrusion policy

In the No Rules Active intrusion policy, all intrusion rules, and all advanced settings except intrusion rule thresholds, are disabled. This policy provides a starting point if you want to create your own intrusion policy instead of basing it on the enabled rules in one of the other system-provided policies.



Note Depending on the system-provided base policy that is selected, the settings of the policy vary. To view the policy settings, click the **Edit** icon next to the policy and then click the **Base Policy** drop-down box.

Benefits of Custom Network Analysis and Intrusion Policies

You may find that the inspector options, intrusion rules, and other advanced settings configured in the system-provided network analysis and intrusion policies do not fully address the security needs of your organization.

Building custom policies can improve the performance of the system in your environment and can provide a focused view of the malicious traffic and policy violations occurring on your network. By creating and tuning custom policies you can configure, at a very granular level, how the system processes and inspects the traffic on your network for intrusions.

All custom policies have a base policy, also called a base layer, which defines the default settings for all configurations in the policy. A layer is a building block that you can use to efficiently manage multiple network analysis or intrusion policies.

In most cases, you base custom policies on system-provided policies, but you can use another custom policy. However, all custom policies have a system-provided policy as the eventual base in a policy chain. Because rule updates can modify system-provided policies, importing a rule update may affect you even if you are using a custom policy as your base. If a rule update affects your deployment, the web interface marks affected policies as out of date.

Benefits of Custom Network Analysis Policies

By default, one network analysis policy preprocesses all unencrypted traffic handled by the access control policy. That means that all packets are decoded and preprocessed according to the same settings, regardless of the intrusion policy (and therefore intrusion rule set) that later examines them.

Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default. A simple way to tune preprocessing is to create and use a custom network analysis policy as the default.

Tuning options available vary by inspector, but some of the ways you can tune inspectors and decoders include:

- You can disable inspectors that do not apply to the traffic you are monitoring. For example, the HTTP Inspect inspector normalizes HTTP traffic. If you are confident that your network does not include any web servers using Microsoft Internet Information Services (IIS), you can disable the inspector option that looks for IIS-specific traffic and thereby reduce system processing overhead.



Note If you disable a inspector in a custom network analysis policy, but the system needs to use that inspector to later evaluate packets against an enabled intrusion or inspector rule, the system automatically enables and uses the inspector although the inspector remains disabled in the network analysis policy web interface.

- Specify ports, where appropriate, to focus the activity of certain inspectors. For example, you can identify additional ports to monitor for DNS server responses or encrypted SSL sessions, or ports on which you decode telnet, HTTP, and RPC traffic.

For advanced users with complex deployments, you can create multiple network analysis policies, each tailored to preprocess traffic differently. Then, you can configure the system to use those policies to govern the preprocessing of traffic using different security zones, networks, or VLANs. (Note that ASA FirePOWER modules cannot restrict preprocessing by VLAN.)



Note Tailoring preprocessing using custom network analysis policies—especially multiple network analysis policies—is an advanced task. Because preprocessing and intrusion inspection are so closely related, you **must** be careful to allow the network analysis and intrusion policies examining a single packet to complement each other.

Benefits of Custom Intrusion Policies

In a newly created access control policy initially configured to perform intrusion prevention, the default action allows all traffic, but first inspects it with the system-provided Balanced Security and Connectivity intrusion policy. Unless you add access control rules or change the default action, all traffic is inspected by that intrusion policy.

To customize your intrusion prevention deployment, you can create multiple intrusion policies, each tailored to inspect traffic differently. Then, configure an access control policy with rules that specify which policy inspects which traffic. Access control rules can be simple or complex, matching and inspecting traffic using multiple criteria including security zone, network or geographical location, VLAN, port, application, requested URL, or user.

The main function of intrusion policies is to manage which intrusion and inspector rules are enabled and how they are configured, as follows:

- Within each intrusion policy, you should verify that all rules applicable to your environment are enabled, and improve performance by disabling rules that are not applicable to your environment. You can specify which rules should drop or modify malicious packets.
- Cisco recommendations allow you to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.
- You can modify existing rules and write new standard text rules as needed to catch new exploits or to enforce your security policies.

Other customizations you might make to an intrusion policy include:

- The sensitive data preprocessor detects sensitive data such as credit card numbers and Social Security numbers in ASCII text. Note that other inspectors that detect specific threats (back orifice attacks, several

portscan types, and rate-based attacks that attempt to overwhelm your network with excessive traffic) are configured in network analysis policies.

- Global thresholds cause the system to generate events based on how many times traffic matching an intrusion rule originates from or is targeted to a specific address or address range within a specified time period. This helps prevent the system from being overwhelmed with a large number of events.
- Suppressing intrusion event notifications and setting thresholds for individual rules or entire intrusion policies can also prevent the system from being overwhelmed with a large number of events.
- In addition to the various views of intrusion events within the web interface, you can enable logging to syslog facilities or send event data to an SNMP trap server. Per policy, you can specify intrusion event notification limits, set up intrusion event notification to external logging facilities, and configure external responses to intrusion events. Note that in addition to these per-policy alerting configurations, you can globally enable or disable email alerting on intrusion events for each rule or rule group. Your email alert settings are used regardless of which intrusion policy processes a packet.

Limitations of Custom Policies

Because preprocessing and intrusion inspection are so closely related, you **must** be careful that your configuration allows the network analysis and intrusion policies processing and examining a single packet to complement each other.

By default, the system uses one network analysis policy to preprocess all traffic handled by managed devices using a single access control policy. The following diagram shows how a newly created access control policy in an inline, intrusion-prevention deployment initially handles traffic. The preprocessing and intrusion prevention phases are highlighted.



Notice how a default network analysis policy governs the preprocessing of *all* traffic handled by the access control policy. Initially, the system-provided Balanced Security and Connectivity network analysis policy is the default.

A simple way to tune preprocessing is to create and use a custom network analysis policy as the default. However, if you disable an inspector in a custom network analysis policy but the system needs to evaluate preprocessed packets against an enabled intrusion or inspector rule, the system automatically enables and uses the inspector although it remains disabled in the network analysis policy web interface.



Note In order to get the performance benefits of disabling an inspector, you **must** make sure that none of your intrusion policies have enabled rules that require that inspector.


An additional challenge arises if you use multiple custom network analysis policies. For advanced users with complex deployments, you can tailor preprocessing to specific security zones, networks, and VLANs by assigning custom network analysis policies to preprocess matching traffic. (Note that ASA FirePOWER cannot restrict preprocessing by VLAN.) To accomplish this, you add custom *network analysis rules* to your access control policy. Each rule has an associated network analysis policy that governs the preprocessing of traffic that matches the rule.



Tip You configure network analysis rules as an advanced setting in an access control policy. Unlike other types of rules, network analysis rules invoke—rather than being contained by—network analysis policies.

The system matches packets to any configured network analysis rules in top-down order by rule number. Traffic that does not match any network analysis rule is preprocessed by the default network analysis policy. While this allows you a great deal of flexibility in preprocessing traffic, keep in mind that all packets, **regardless** of which network analysis policy preprocessed them, are subsequently matched to access control rules—and thus to potential inspection by intrusion policies—in their own process. In other words, preprocessing a packet with a particular network analysis policy does **not** guarantee that the packet will be examined with any particular intrusion policy. You **must** carefully configure your access control policy so it invokes the correct network analysis and intrusion policies to evaluate a particular packet.

The following diagram shows in focused detail how the network analysis policy (preprocessing) selection phase occurs before and separately from the intrusion prevention (rules) phase. For simplicity, the diagram eliminates the discovery and file/malware inspection phases. It also highlights the default network analysis and default-action intrusion policies.



In this scenario, an access control policy is configured with two network analysis rules and a default network analysis policy:

- Network Analysis Rule A preprocesses matching traffic with Network Analysis Policy A. Later, you want this traffic to be inspected by Intrusion Policy A.
- Network Analysis Rule B preprocesses matching traffic with Network Analysis Policy B. Later, you want this traffic to be inspected by Intrusion Policy B.
- All remaining traffic is preprocessed with the default network analysis policy. Later, you want this traffic to be inspected by the intrusion policy associated with the access control policy's default action.

After the system preprocesses traffic, it can examine the traffic for intrusions. The diagram shows an access control policy with two access control rules and a default action:

- Access Control Rule A allows matching traffic. The traffic is then inspected by Intrusion Policy A.
- Access Control Rule B allows matching traffic. The traffic is then inspected by Intrusion Policy B.
- The access control policy's default action allows matching traffic. The traffic is then inspected by the default action's intrusion policy.

Each packet's handling is governed by a network analysis policy and intrusion policy pair, but the system does **not** coordinate the pair for you. Consider a scenario where you misconfigure your access control policy so that Network Analysis Rule A and Access Control Rule A do not process the same traffic. For example, you could intend the paired policies to govern the handling of traffic on a particular security zone, but you mistakenly use different zones in the two rules' conditions. This could cause traffic to be incorrectly preprocessed. For this reason, tailoring preprocessing using network analysis rules and custom policies is an **advanced** task.

Note that for a single connection, although the system selects a network analysis policy before an access control rule, some preprocessing (notably application layer preprocessing) occurs after access control rule selection. This does **not** affect how you configure preprocessing in custom network analysis policies.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.



CHAPTER 2

Migrate from Snort 2 to Snort 3

Starting with Version 7.0, Snort 3 is the default inspection engine for new threat defense deployments with management center. If you are still using the Snort 2 inspection engine, switch to Snort 3 now for improved detection and performance.

Upgrading threat defense to Version 7.2 through 7.6 also upgrades eligible Snort 2 devices to Snort 3. For devices that are ineligible because they use custom intrusion or network analysis policies, manually upgrade to Snort 3 as described here.

Although you can switch individual devices back, you should not. Snort 2 will be deprecated in a future release and will eventually prevent threat defense upgrade.

- [Snort 3 Inspection Engine, on page 17](#)
- [Prerequisites for Network Analysis and Intrusion Policies, on page 18](#)
- [How to Migrate from Snort 2 to Snort 3, on page 18](#)
- [View Snort 2 and Snort 3 Base Policy Mapping, on page 22](#)
- [Synchronize Snort 2 Rules with Snort 3 , on page 22](#)
- [Deploy Configuration Changes, on page 23](#)

Snort 3 Inspection Engine

Snort 3 is the default inspection engine for newly registered threat defense devices of version 7.0 and later. However, for threat defense devices of lower versions, Snort 2 is the default inspection engine. When you upgrade a managed threat defense device to version 7.0 or later, the inspection engine remains on Snort 2. To use Snort 3 in upgraded threat defenses of version 7.0 and later, you must explicitly enable it. When Snort 3 is enabled as the inspection engine of the device, the Snort 3 version of the intrusion policy that is applied on the device (through the access control policies) is activated and applied to all the traffic passing through the device.

You can switch Snort versions when required. Snort 2 and Snort 3 intrusion rules are mapped and the mapping is system-provided. However, you may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. If you change the rule action for one rule in Snort 2, that change is not retained if you switch to Snort 3 without synchronizing Snort 2 with Snort 3. For more information on synchronization, see [Synchronize Snort 2 Rules with Snort 3 , on page 22](#).

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

How to Migrate from Snort 2 to Snort 3

Migrating from Snort 2 to Snort 3 requires you to switch the inspection engine of the threat defense device from Snort 2 to Snort 3.

Depending on your requirements, the tasks to complete the migration of your device from Snort 2 to Snort 3 is listed in the following table:

Step	Task	Links to Procedures
1	Enable Snort 3	<ul style="list-style-type: none"> • Enable Snort 3 on an Individual Device, on page 19 • Enable Snort 3 on Multiple Devices, on page 19
2	Convert Snort 2 custom rules to Snort 3	<ul style="list-style-type: none"> • Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3, on page 20 • Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3, on page 21
3	Synchronize Snort 2 rules with Snort 3	Synchronize Snort 2 Rules with Snort 3 , on page 22

Prerequisites for Migrating from Snort 2 to Snort 3

The following are the recommended prerequisites that you must consider before migrating your device from Snort 2 to Snort 3.

- Have a working knowledge of Snort. To learn about the Snort 3 architecture, see [Snort 3 Adoption](#).
- Back up your management center. See [Backup the Management Center](#).
- Back up your intrusion policy. See [Exporting Configurations](#).
- Clone your intrusion policy. To do this, you can use an existing policy as the base policy to create a copy of your intrusion policy. In the **Intrusion Policies** page, click **Create Policy** and choose an existing intrusion policy from the **Base Policy** dropdown list.

Enable Snort 3 on an Individual Device

**Important**

During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

Procedure

Step 1 Choose **Devices > Device Management**.

Step 2 Click the device to go to the device home page.

Note

The device is marked as Snort 2 or Snort 3, showing the current version on the device.

Step 3 Click the **Device** tab.

Step 4 In the Inspection Engine section, click **Upgrade**.

Note

In case you want to disable Snort 3, click **Revert to Snort 2** in the Inspection Engine section.

Step 5 Click **Yes**.

What to do next

Deploy the changes on the device. See, [Deploy Configuration Changes, on page 23](#).

The system converts your policy configurations during the deployment process to make them compatible with the selected Snort version.

Enable Snort 3 on Multiple Devices

To enable Snort 3 on multiple devices, ensure all the required threat defense devices are on version 7.0 or later.

**Important**

During the deployment process, there could be a momentary traffic loss because the current inspection engine needs to be shut down.

Procedure

Step 1 Choose **Devices > Device Management**.

Step 2 Select all the devices on which you want to enable or disable Snort 3.

Note

The devices are marked as Snort 2 or Snort 3, showing the current version on the device.

Step 3 Click **Select Bulk Action** drop-down list and choose **Upgrade to Snort 3**.

Note

To disable Snort 3, click **Downgrade to Snort 2**.

Step 4 Click **Yes**.

What to do next

Deploy the changes on the device. See, [Deploy Configuration Changes, on page 23](#).

The system converts your policy configurations during the deployment process to make them compatible with the selected Snort version.

Convert Snort 2 Custom IPS Rules to Snort 3

If you are using a rule set from a third-party vendor, contact that vendor to confirm that their rules successfully convert to Snort 3 or to obtain a replacement rule set written natively for Snort 3. If you have custom rules that you have written yourself, familiarize with writing Snort 3 rules prior to conversion, so you can update your rules to optimize Snort 3 detection after conversion. See the links below to learn more about writing rules in Snort 3.

- <https://blog.snort.org/2020/08/how-rules-are-improving-in-snort-3.html>
- <https://blog.snort.org/2020/10/talos-transition-to-snort-3.html>

You can refer to other blogs at <https://blog.snort.org/> to learn more about Snort 3 rules.

See the following procedures to convert Snort 2 rules to Snort 3 rules using the system-provided tool.

- [Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3, on page 20](#)
- [Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3, on page 21](#)



Important Snort 2 network analysis policy (NAP) settings *cannot* be copied to Snort 3 automatically. NAP settings have to be manually replicated in Snort 3.

Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3

Procedure

- Step 1** Choose **Objects > Intrusion Rules**.
- Step 2** Click **Snort 3 All Rules** tab.
- Step 3** Ensure **All Rules** is selected in the left pane.
- Step 4** Click the **Tasks** drop-down list and choose:

- **Convert Snort 2 rules and import**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and import them into management center as Snort 3 custom rules.
- **Convert Snort 2 rules and download**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and download them into your local system.

Step 5 Click **OK**.

Note

- If you selected **Convert and import** in the previous step, then all the converted rules are saved under a newly created rule group **All Snort 2 Converted Global** under **Local Rules**.
- If you selected **Convert and download** in the previous step, then save the rules file locally. You can review the converted rules in the downloaded file and later upload them by following the steps in [Add Custom Rules to Rule Groups](#), on page 49.

Refer to the video [Converting Snort 2 Rules to Snort 3](#) for additional support and information.

What to do next


Deploy configuration changes; see [Deploy Configuration Changes](#), on page 23.

Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3


Procedure

Step 1 Choose **Policies > Intrusion**.

Step 2 In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

Step 3 Click the **Sync** icon **Snort out-of-Sync** () of the intrusion policy.

Note

If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in green **Snort in-Sync** (). It indicates that there are no custom rules to be converted.

Step 4 Read through the summary and click the **Custom Rules** tab.

Step 5 Choose:

- **Import converted rules to this policy**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and import them into management center as Snort 3 custom rules.
- **Download converted rules**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and download them into your local system. You can review the converted rules in the downloaded file and later upload the file by clicking the upload icon.

Step 6 Click **Re-Sync**.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

View Snort 2 and Snort 3 Base Policy Mapping

Procedure

-
- Step 1** Choose **Policies > Intrusion**.
- Step 2** Ensure the **Intrusion Policies** tab is selected.
- Step 3** Click **IPS Mapping**.
- Step 4** In the **IPS Policy Mapping** dialog box, click **View Mappings** to view the Snort 3 to Snort 2 intrusion policy mapping.
- Step 5** Click **OK**.
-

Synchronize Snort 2 Rules with Snort 3

To ensure that the Snort 2 version settings and custom rules are retained and carried over to Snort 3, the management center provides the synchronization functionality. Synchronization helps Snort 2 rule override settings and custom rules, which you may have altered and added over the last few months or years, to be replicated on the Snort 3 version. This utility helps to synchronize Snort 2 version policy configuration with Snort 3 version to start with similar coverage.

If the management center is upgraded from 6.7 or earlier to 7.0 or later version, the system synchronizes the configuration. If the management center is a fresh 7.0 or later version, you can upgrade to a higher version, and the system will not synchronize any content during upgrade.

Before upgrading a device to Snort 3, if changes are made in Snort 2 version, you can use this utility to have the latest synchronization from Snort 2 version to Snort 3 version so that you start with a similar coverage.




Note On moving to Snort 3, it is recommended that you manage the Snort 3 version of the policy independently and do not use this utility as a regular operation.




-
- Important**
- Only the Snort 2 rule overrides and custom rules are copied to Snort 3 and not the other way around. You may not find a one-to-one mapping of all the intrusion rules in Snort 2 and Snort 3. Your changes to rule actions for rules that exist in both versions are synchronized when you perform the following procedure.
 - Synchronization *does not* migrate the threshold and suppression settings of any custom or system-provided rules from Snort 2 to Snort 3.
-

Procedure

- Step 1** Choose **Policies > Intrusion**.
- Step 2** Ensure the **Intrusion Policies** tab is selected.
- Step 3** Click **Show Snort 3 Sync status**.
- Step 4** Identify the intrusion policy that is out-of-sync.
- Step 5** Click the **Sync** icon **Snort out-of-Sync** ().

Note

If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in green **Snort in-Sync** ().

- Step 6** Read through the summary and download a copy of the summary if required.
- Step 7** Click **Re-Sync**.

Note

- The synchronized settings will be applicable on the Snort 3 intrusion engine only if it is applied on a device, and after a successful deployment.
- Snort 2 custom rules can be converted to Snort 3 using the system-provided tool. If you have any Snort 2 custom rules click the Custom Rules tab and follow the on-screen instructions to convert the rules. For more information, see [Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3, on page 21](#).

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Deploy Configuration Changes

After you change configurations, deploy them to the affected devices.



Note

This topic covers the basic steps involved in deploying configuration changes. We *strongly* recommend that you refer the *Deploy Configuration Changes* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide* to understand the prerequisites and implications of deploying the changes before proceeding with the steps.



Caution

When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic.

Procedure

Step 1 On the Secure Firewall Management Center menu bar, click **Deploy** and choose **Deployment**.

The GUI page lists the devices with out-of-date configurations having **Pending** status.

- The **Modified By** column lists the users who have modified the policies or objects. Expand the device listing to view the users who have modified the policies for each policy listing.

Note


Username is not provided for deleted policies and objects.

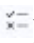
- The **Inspect Interruption** column indicates if traffic inspection interruption might occur in the device during deployment.

If this column is blank for a device, it indicates that there will be no traffic inspection interruptions on that device during deployment.


- The **Last Modified Time** column specifies the last time you made configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment.

Step 2 Identify and choose the devices on which you want to deploy configuration changes.

- Search—Search for the device name, type, domain, group, or status in the search box.
- Expand—Click **Expand Arrow** () to view device-specific configuration changes to be deployed.

When you check a check box adjacent to a device, all the changes made to the device and listed under the device, are pushed for deployment. However, you can use **Policy selection** () to select individual policies or specific configurations to deploy while withholding the remaining changes without deploying them.

Note

- When the status in the **Inspect Interruption** column indicates (**Yes**) that deploying will interrupt inspection, and perhaps traffic, on a threat defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ().
- When there are changes to interface groups, security zones, or objects, the impacted devices are shown as out-of-date on the management center. To ensure that these changes take effect, the policies with these interface groups, security zones, or objects, also need to be deployed along with these changes. The impacted policies are shown as out-of-date on the **Preview** page on the management center.

Step 3 Click **Deploy**.

Step 4 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- Deploy—Continue deploying without resolving warning conditions. You cannot proceed if the system identifies errors.

- Close—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

During deployment, if there is a deployment failure, there is a possibility that the failure may impact traffic. However, it depends on certain conditions. If there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For details, see the Deploy Configuration Changes topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.



PART I

Intrusion Detection and Prevention in Snort 3

- [Get Started with Snort 3 Intrusion Policies, on page 29](#)
- [Tune Intrusion Policies Using Rules, on page 37](#)
- [Tailor Intrusion Protection for Your Network Assets, on page 53](#)



CHAPTER 3

Get Started with Snort 3 Intrusion Policies

This chapter provides information on managing Snort 3 intrusion policies and access control rule configurations for intrusion detection and prevention.

- [Overview of Intrusion Policies, on page 29](#)
- [Prerequisites for Network Analysis and Intrusion Policies, on page 30](#)
- [Create a Custom Snort 3 Intrusion Policy , on page 30](#)
- [Edit Snort 3 Intrusion Policies, on page 31](#)
- [Change the Base Policy of an Intrusion Policy, on page 32](#)
- [Manage Intrusion Policies, on page 32](#)
- [Access Control Rule Configuration to Perform Intrusion Prevention, on page 33](#)
- [Deploy Configuration Changes, on page 35](#)

Overview of Intrusion Policies

Intrusion policies are defined sets of intrusion detection and prevention configurations that inspect traffic for security violations and, in inline deployments, can block or alter malicious traffic. Intrusion policies are invoked by your access control policy and are the system's last line of defense before traffic is allowed to its destination.

At the heart of each intrusion policy are the intrusion rules. An enabled rule causes the system to generate intrusion events for (and optionally block) traffic matching the rule. Disabling a rule stops processing of the rule.

The system delivers several base intrusion policies, which enable you to take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and inspector rule states (enabled or disabled), as well as provides the initial configurations for other advanced settings.



Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules.

If you create a custom intrusion policy, you can:

- Tune detection by enabling and disabling rules, as well as by writing and adding your own rules.

- Use Secure Firewall recommendations to associate the operating systems, servers, and client application protocols detected on your network with rules specifically written to protect those assets.

An intrusion policy can drop matching packets and generate intrusion events. To configure an intrusion or preprocessor drop rule, set its state to Block.

When tailoring your intrusion policy, especially when enabling and adding rules, keep in mind that some intrusion rules require that traffic first be decoded or preprocessed in a certain way. Before an intrusion policy examines a packet, the packet is preprocessed according to configurations in a network analysis policy. If you disable a required inspector, the system automatically uses it with its current settings, although the inspector remains disabled in the network analysis policy web interface.



Caution

Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

After you configure a custom intrusion policy, you can use it as part of your access control configuration by associating the intrusion policy with one or more access control rules or an access control policy's default action. This forces the system to use the intrusion policy to examine certain allowed traffic before the traffic passes to its final destination. A variable set that you pair with the intrusion policy allows you to accurately reflect your home and external networks and, as appropriate, the servers on your network.

Note that by default, the system disables intrusion inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion inspection configured.

Refer to the video for additional support and information - [Snort 3 Intrusion Policy Overview](#).

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Create a Custom Snort 3 Intrusion Policy

Procedure

- Step 1** Choose **Policies > Intrusion**.
- Step 2** Click **Create Policy**.
- Step 3** Enter a unique **Name** and, optionally, a **Description**.
- Step 4** Choose the **Inspection Mode**.

The selected action determines whether intrusion rules block and alert (**Prevention** mode) or only alert (**Detection** mode).

Note

Before selecting the prevention mode, you might want block rules to alert only so you can identify rules that cause a lot of false positives.

Step 5 Choose the **Base Policy**.

You can use either a system-provided policy or an existing policy as your base policy.

Step 6 Click **Save**.

The new policy has the same settings as its base policy.

What to do next

To customize the policy, see [Edit Snort 3 Intrusion Policies, on page 31](#).

Edit Snort 3 Intrusion Policies

While editing a Snort 3 policy, all the changes are saved instantaneously. No additional action is required to save the changes.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Rule Action Logging

From Management Center 7.2.0 onwards, in the **Intrusion Events** page, the event in the **Inline Result** column displays the same name as the IPS action applied to the rule, so that you can see the action that was applied on the traffic matching the rule.

For the IPS actions, the following table shows the events that are displayed in the **Inline Result** column of the **Intrusion Events** page and **Action** column for **Intrusion Event Type** in the **Unified Events** page.

IPS Action for Snort 3	Inline Result - Management Center 7.1.0 and earlier	Inline Result -Management Center 7.2.0 onwards
Alert	Pass	Alert
Block	Dropped/Would Have Dropped/Partially Dropped	Block/Would Block/Partial Block
Drop	Dropped/Would have dropped	Drop/Would drop
Reject	Dropped/Would have dropped	Reject/Would reject
Rewrite	Allow	Rewrite

**Important**

- In case of a rule without the “Replace” option, the **Rewrite** action is displayed as **Would Rewrite**.
- The **Rewrite** action would also be displayed as **Would Rewrite** if the "Replace" option is specified, but the IPS policy is in Detection mode or the device is in Inline-TAP/Passive mode.

**Note**

In case of backward compatibility (Management Center 7.2.0 managing a Threat Defense 7.1.0 device), the events mentioned are applicable only to the Alert IPS action where **Pass** is displayed as **Alert** for events. For all the other actions, the events for Management Center 7.1.0 are applicable.

Change the Base Policy of an Intrusion Policy

You can choose a different system-provided or custom policy as your base policy.

You can chain up to five custom policies, with four of the five using one of the other four previously created policies as its base policy; the fifth must use a system-provided policy as its base.

Procedure

-
- Step 1** Choose **Policies > Intrusion**.
 - Step 2** Click **Edit** (✎) next to the intrusion policy you want to configure.
 - Step 3** Choose a policy from the **Base Policy** drop-down list.
 - Step 4** Click **Save**.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Manage Intrusion Policies

On the Intrusion Policy page (**Policies > Intrusion**) you can view your current custom intrusion policies, along with the following information:


- Number of access control policies and devices are using the intrusion policy to inspect traffic
- In a multidomain deployment, the domain where the policy was created


In a multidomain deployment, the system displays policies created in the current domain, which you can edit. It also displays policies created in ancestor domains, which you cannot edit. To view and edit policies created in a lower domain, switch to that domain.

Procedure

Step 1 Choose **Policies > Intrusion**.

Step 2 Manage your intrusion policy:

- Create — Click **Create Policy**; see [Create a Custom Snort 3 Intrusion Policy](#), on page 30.
- Delete — Click **Delete** () next to the policy you want to delete. The system prompts you to confirm and informs you if another user has unsaved changes in the policy. Click **OK** to confirm.

If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Edit intrusion policy details — Click **Edit** () next to the policy you want to edit. You can edit the **Name**, **Inspection Mode**, and the **Base Policy** of the intrusion policy.
- Edit intrusion policy settings — Click **Snort 3 Version**; see [Edit Snort 3 Intrusion Policies](#), on page 31.
- Export — If you want to export an intrusion policy to import on another management center, click Export; see the *Exporting Configurations* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.
- Deploy — Choose **Deploy > Deployment**; see [Deploy Configuration Changes](#), on page 23.
- Report — Click **Report**; see the *Generating Current Policy Reports* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*. Generates two reports, one for each policy version.

Access Control Rule Configuration to Perform Intrusion Prevention

An access control policy can have multiple access control rules associated with intrusion policies. You can configure intrusion inspection for any Allow or Interactive Block access control rule, which permits you to match different intrusion inspection profiles against different types of traffic on your network before it reaches its final destination.

Whenever the system uses an intrusion policy to evaluate traffic, it uses an associated *variable set*. Variables in a set represent values commonly used in intrusion rules to identify source and destination IP addresses and ports. You can also use variables in intrusion policies to represent IP addresses in rule suppressions and dynamic rule states.



Tip Even if you use system-provided intrusion policies, Cisco **strongly** recommends you configure the system's intrusion variables to accurately reflect your network environment. At a minimum, modify default variables in the default set.

Understanding System-Provided and Custom Intrusion Policies

Cisco delivers several intrusion policies with the system. By using system-provided intrusion policies, you can take advantage of the experience of the Cisco Talos Intelligence Group (Talos). For these policies, Talos sets intrusion and preprocessor rule states, as well as provides the initial configurations for advanced settings. You can use system-provided policies as-is, or you can use them as the base for custom policies. Building custom policies can improve the performance of the system in your environment and provide a focused view of the malicious traffic and policy violations occurring on your network.

Connection and Intrusion Event Logging

When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, it saves that event to the Management Center. The system also automatically logs the end of the connection where the intrusion occurred to the Management Center database, regardless of the logging configuration of the access control rule.

Access Control Rule Configuration and Intrusion Policies

The number of unique intrusion policies you can use in a single access control policy depends on the model of the target devices; more powerful devices can handle more. Every unique **pair** of intrusion policy and variable set counts as one policy. Although you can associate a different intrusion policy-variable set pair with each Allow and Interactive Block rule (as well as with the default action), you cannot deploy an access control policy if the target devices have insufficient resources to perform inspection as configured.

Configure an Access Control Rule to Perform Intrusion Prevention

You must be an Admin, Access Admin, or Network Admin to perform this task.

Procedure

-
- | | |
|---------------|--|
| Step 1 | In the access control policy editor, create a new rule or edit an existing rule; see the <i>Access Control Rule Components</i> topic in the latest version of the <i>Cisco Secure Firewall Management Center Configuration Guide</i> . |
| Step 2 | Ensure the rule action is set to Allow , Interactive Block , or Interactive Block with reset . |
| Step 3 | Click Inspection . |
| Step 4 | Choose a system-provided or a custom intrusion policy, or choose None to disable intrusion inspection for traffic that matches the access control rule. |
| Step 5 | If you want to change the variable set associated with the intrusion policy, choose a value from the Variable Set drop-down list. |
| Step 6 | Click Save to save the rule. |
| Step 7 | Click Save to save the policy. |
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#), on page 23.

Deploy Configuration Changes

After you change configurations, deploy them to the affected devices.



Note This topic covers the basic steps involved in deploying configuration changes. We *strongly* recommend that you refer the *Deploy Configuration Changes* topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide* to understand the prerequisites and implications of deploying the changes before proceeding with the steps.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic.

Procedure

Step 1 On the Secure Firewall Management Center menu bar, click **Deploy** and choose **Deployment**.

The GUI page lists the devices with out-of-date configurations having **Pending** status.

- The **Modified By** column lists the users who have modified the policies or objects. Expand the device listing to view the users who have modified the policies for each policy listing.

Note


Username is not provided for deleted policies and objects.


- The **Inspect Interruption** column indicates if traffic inspection interruption might occur in the device during deployment.

If this column is blank for a device, it indicates that there will be no traffic inspection interruptions on that device during deployment.


- The **Last Modified Time** column specifies the last time you made configuration changes.
- The **Preview** column allows you to preview the changes for the next deployment.
- The **Status** column provides the status for each deployment.

Step 2 Identify and choose the devices on which you want to deploy configuration changes.

- Search—Search for the device name, type, domain, group, or status in the search box.
- Expand—Click **Expand Arrow** () to view device-specific configuration changes to be deployed.

When you check a check box adjacent to a device, all the changes made to the device and listed under the device, are pushed for deployment. However, you can use **Policy selection** () to select individual policies or specific configurations to deploy while withholding the remaining changes without deploying them.

Note

- When the status in the **Inspect Interruption** column indicates (**Yes**) that deploying will interrupt inspection, and perhaps traffic, on a threat defense device, the expanded list indicates the specific configurations causing the interruption with the **Inspect Interruption** ().
- When there are changes to interface groups, security zones, or objects, the impacted devices are shown as out-of-date on the management center. To ensure that these changes take effect, the policies with these interface groups, security zones, or objects, also need to be deployed along with these changes. The impacted policies are shown as out-of-date on the **Preview** page on the management center.

Step 3 Click **Deploy**.

Step 4 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

You have the following choices:

- **Deploy**—Continue deploying without resolving warning conditions. You cannot proceed if the system identifies errors.
- **Close**—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

During deployment, if there is a deployment failure, there is a possibility that the failure may impact traffic. However, it depends on certain conditions. If there are specific configuration changes in the deployment, the deployment failure may lead to traffic being interrupted. For details, see the Deploy Configuration Changes topic in the latest version of the *Cisco Secure Firewall Management Center Configuration Guide*.



CHAPTER 4

Tune Intrusion Policies Using Rules

This chapter provides information on custom rules in Snort 3, intrusion rule action, intrusion event notification filters in an intrusion policy, converting Snort 2 custom rules to Snort 3, and adding rule groups with custom rules to an intrusion policy.

- [Overview of Tuning Intrusion Rules, on page 37](#)
- [Intrusion Rule Types, on page 38](#)
- [Prerequisites for Network Analysis and Intrusion Policies, on page 39](#)
- [Custom Rules in Snort 3, on page 39](#)
- [View Snort 3 Intrusion Rules in an Intrusion Policy, on page 40](#)
- [Intrusion Rule Action, on page 40](#)
- [Intrusion Event Notification Filters in an Intrusion Policy, on page 42](#)
- [Add Intrusion Rule Comments, on page 46](#)
- [Snort 2 Custom Rules Conversion to Snort 3, on page 47](#)
- [Add Custom Rules to Rule Groups, on page 49](#)
- [Add Rule Groups with Custom Rules to an Intrusion Policy, on page 50](#)
- [Manage Custom Rules in Snort 3, on page 50](#)
- [Delete Custom Rules, on page 51](#)
- [Delete Rule Groups, on page 52](#)

Overview of Tuning Intrusion Rules

You can configure rule states and other settings for shared object rules, standard text rules, and inspector rules.

You enable a rule by setting its rule state to Alert or to Block. Enabling a rule causes the system to generate events on traffic matching the rule. Disabling a rule stops processing of the rule. You can also set your intrusion policy so that a rule set to Block generates events on, and drops, matching traffic.

You can filter rules to display a subset of rules, enabling you to select the exact set of rules where you want to change rule states or rule settings.

When an intrusion rule or rule argument requires a disabled inspector, the system automatically uses it with its current configuration even though it remains disabled in the network analysis policy's web interface.



Note

We recommend that you do not modify shared object rules and you only enable or disable these rules for your threat defense device. To create custom Snort rules, contact Cisco support.

Intrusion Rule Types

An intrusion rule is a specified set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities in your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule, and triggers the rule if the data packet meets all the conditions specified in the rule.

An intrusion policy contains:

- *intrusion rules*, which are subdivided into *shared object rules* and *standard text rules*
- *inspector rules*, which are associated with a detection option of the packet decoder or with one of the inspectors included with the system

The following table summarizes attributes of these rule types:

Table 2: Intrusion Rule Types

Type	Generator ID (GID)	Snort ID (SID)	Source	Can Copy?	Can Edit?
shared object rule	3	lower than 1000000	Cisco Talos Intelligence Group (Talos)	yes	limited
standard text rule	1 (Global domain or legacy GID)	lower than 1000000	Talos	yes	limited
	1000 - 2000 (descendant domain)	1000000 or higher	Created or imported by user	yes	yes
preprocessor rule	decoder- or preprocessor-specific	lower than 1000000	Talos	no	no
		1000000 or higher	Generated by the system during option configuration	no	no

You cannot save changes to any rule created by Talos, but you can save a copy of a modified rule as a custom rule. You can modify either variables used in the rule or rule header information (such as source and destination ports and IP addresses). In a multidomain deployment, rules created by Talos belong to the Global domain. Administrators in descendant domains can save local copies of the rules, which they can then edit.

For the rules it creates, Talos assigns default rule states in each default intrusion policy. Most preprocessor rules are disabled by default and must be enabled if you want the system to generate events for preprocessor rules and, in an inline deployment, drop offending packets.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Custom Rules in Snort 3

You can create a custom intrusion rule by importing a local rule file. The rule file can either have a `.txt` or `.rules` extension. The system saves the custom rule in the local rule category, regardless of the method you used to create it. A custom rule must belong to a rule group. However, a custom rule can be a part of two or more groups as well.

When you create a custom intrusion rule, the system assigns it a unique rule number, which has the format `GID:SID:Rev`. The elements of this number are:

- **GID**—Generator ID. For custom rules, it is not necessary to specify the GID. The system automatically generates the GID based on whether you are in the Global domain or a sub-domain while uploading the rules. For all standard text rules, this value is 2000 for a Global domain.
- **SID**—Snort ID. Indicates whether the rule is a local rule or a system rule. When you create a new rule, assign a unique SID to the rule.
SID numbers for local rules start at 1000000, and the SID for each new local rule is incremented by one.
- **Rev**—The revision number. For a new rule, the revision number is one. Each time you modify a custom rule the revision number should be incremented by one.

In a custom standard text rule, you set the rule header settings and the rule keywords and arguments. You can use the rule header settings to focus the rule to only match traffic using a specific protocol and traveling to or from specific IP addresses or ports.

To check whether a SID is enabled or disabled, verify the entries in the `snort.lua` file located in the `./file-contents/ngfw/var/sf/detection_engines/<id>/ips/<id>` directory.

- If the SID is disabled by default, no entry will be present in the file.
- If the SID is manually enabled, you will see an entry with **enable:yes**.
- If the SID is disabled after being manually enabled, the entry remains in the file and will display **enable:no**.



Note

- Snort 3 custom rules cannot be edited. Ensure that custom rules have a valid classification message for `classtype` within the rule text. If you import a rule without a classification or wrong classification, then delete and recreate the rule.
- You can create custom intrusion rules using Snort 3. However, support for tuning and troubleshooting these rules is not available currently.

View Snort 3 Intrusion Rules in an Intrusion Policy

You can adjust how rules are displayed in the intrusion policy. You can also display the details for a specific rule to see rule settings, rule documentation, and other rule specifics.

Procedure

-
- Step 1** Choose **Policies > Intrusion**.
- Step 2** Click **Snort 3 Version** next to the policy.
- Step 3** While viewing the rules, you can:
- Filter the rules.
 - Choose a rule group to see rules related to that group.
 - View an intrusion rule's details.
 - View rule comments.
 - View rule documentation.

See [Edit Snort 3 Intrusion Policies, on page 31](#) for details on performing these tasks.

Intrusion Rule Action

Intrusion rule action allows you to enable or disable the rule within an individual intrusion policy, as well as specify which action the system takes if monitored conditions trigger the rule.

The Cisco Talos Intelligence Group (Talos) sets the default action of each intrusion and inspector rule in each default policy. For example, a rule may be enabled in the Security over Connectivity default policy and disabled in the Connectivity over Security default policy. Talos sometimes uses a rule update to change the default action of one or more rules in a default policy. If you allow rule updates to update your base policy, you also allow the rule update to change the default action of a rule in your policy when the default action changes in the default policy you used to create your policy (or in the default policy it is based on). Note, however, that if you have changed the rule action, the rule update does not override your change.

When you create an intrusion rule, it inherits the default actions of the rules in the default policy you use to create your policy.

Intrusion Rule Action Options

In an intrusion policy, you can set a rule's action to the following values:

Alert

You want the system to detect a specific intrusion attempt and generate an intrusion event when it finds matching traffic. When a malicious packet crosses your network and triggers the rule, the packet is sent to its destination and the system generates an intrusion event. The malicious packet reaches its target, but you are notified through the event logging.

Block

You want the system to detect a specific intrusion attempt, drop the packet containing the attack, and generate an intrusion event when it finds matching traffic. The malicious packet never reaches its target, and you are notified through the event logging.

Disable

You do not want the system to evaluate matching traffic.



Note Choosing either the **Alert** or **Block** options enables the rule. Choosing **Disable** disables the rule.

We **strongly** recommend that you **do not** enable all the intrusion rules in an intrusion policy. The performance of your managed device is likely to degrade if all rules are enabled. Instead, tune your rule set to match your network environment as closely as possible.

Set Intrusion Rule Action

Intrusion rule actions are policy-specific.

Procedure

Step 1 Choose **Policies > Intrusion**.

Step 2 Click **Snort 3 Version** next to the policy you want to edit.

Tip

This page shows the total number of:

- disabled rules
- enabled rules set to Alert
- enabled rules set to Block
- overridden rules

Step 3 Choose the rule or rules where you want to set the rule action.

Step 4 Choose one of the rule actions from the **Rule Action** drop-down list. See [Edit Snort 3 Intrusion Policies, on page 31](#) for more information about the different rule actions.

Step 5 Click **Save**.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Intrusion Event Notification Filters in an Intrusion Policy

The importance of an intrusion event can be based on frequency of occurrence, or on source or destination IP address. In some cases you may not care about an event until it has occurred a certain number of times. For example, you may not be concerned if someone attempts to log into a server until they fail a certain number of times. In other cases, you may only need to see a few occurrences to know there is a widespread problem. For example, if a DoS attack is launched against your web server, you may only need to see a few occurrences of an intrusion event to know that you need to address the situation. Seeing hundreds of the same event only overwhelms your system.

Intrusion Event Thresholds

You can set thresholds for individual rules to limit the number of times the system logs and displays an intrusion event based on how many times the event is generated within a specified time period. This can prevent you from being overwhelmed with a large number of identical events. You can set thresholds per shared object rule, standard text rule, or inspector rule.

Set Intrusion Event Thresholds

To set a threshold, first specify the thresholding type.

Table 3: Thresholding Options

Option	Description
Limit	Logs and displays events for the specified number of packets (specified by the Count argument) that trigger the rule during the specified time period. For example, if you set the type to Limit , the Count to 10, and the Seconds to 60, and 14 packets trigger the rule, the system stops logging events for the rule after displaying the first 10 that occur within the same minute.
Threshold	Logs and displays a single event when the specified number of packets (specified by the Count argument) trigger the rule during the specified time period. Note that the counter for the time restarts after you hit the threshold count of events and the system logs that event. For example, you set the type to Threshold , Count to 10, and Seconds to 60, and the rule triggers 10 times by second 33. The system generates one event, then resets the Seconds and Count counters to 0. The rule then triggers another 10 times in the next 25 seconds. Because the counters reset to 0 at second 33, the system logs another event.
Both	Logs and displays an event once per specified time period, after the specified number (count) of packets trigger the rule. For example, if you set the type to Both , Count to two, and Seconds to 10, the following event counts result: <ul style="list-style-type: none">• If the rule is triggered once in 10 seconds, the system does not generate any events (the threshold is not met)• If the rule is triggered twice in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time)• If the rule is triggered four times in 10 seconds, the system generates one event (the threshold is met when the rule triggers the second time, and following events are ignored)

Secondly, specify tracking, which determines whether the event threshold is calculated per source or destination IP address.

Table 4: Thresholding IP Options

Option	Description
Source	Calculates event instance count per source IP address.
Destination	Calculates event instance count per destination IP address.

Finally, specify the number of instances and time period that define the threshold.

Table 5: Thresholding Instance/Time Options

Option	Description
Count	The number of event instances per specified time period per tracking IP address required to meet the threshold.
Seconds	The number of seconds that elapse before the count resets. If you set the threshold type to limit , the tracking to Source IP , the count to 10, and the seconds to 10, the system logs and displays the first 10 events that occur in 10 seconds from a given source port. If only 7 events occur in the first 10 seconds, the system logs and displays those; if 40 events occur in the first 10 seconds, the system logs and displays 10, then begins counting again when the 10-second time period elapses.

Note that you can use intrusion event thresholding alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event suppression.



Tip You can also add thresholds from within the packet view of an intrusion event.

Set Threshold for an Intrusion Rule in Snort 3

You can set a single threshold for a rule from the Rule Detail page. Adding a threshold overwrites any existing threshold for the rule. The threshold you set for an intrusion rule is applied to each packet thread. However, the configuration is fully applied only within the context of a unique flow. There may be more alerts on different network flows, but there will not be fewer alerts than the configured number.

Procedure

- Step 1** Choose **Objects > Intrusion Rules**.
- Step 2** Click **Snort 3 All Rules** tab.
- Step 3** From an intrusion rule's Alert Configuration column, click the **None** link.
- Step 4** Click **Edit** (✎).
- Step 5** In the Alert Configuration window, click the **Threshold** tab.
- Step 6** From the **Type** drop-down list, choose the type of threshold you want to set:

- Choose **Limit** to limit notification to the specified number of event instances per time period.
- Choose **Threshold** to provide notification for each specified number of event instances per time period.
- Choose **Both** to provide notification once per time period after a specified number of event instances.

Step 7 Choose **Source** or **Destination** in the **Track By** field to indicate whether you want the event instances tracked by source or destination IP address.

Step 8 Enter the number of event instances you want to use as your threshold in the **Count** field.

Step 9 Enter a number that specifies the time period, in seconds, for which event instances are tracked in the **Seconds** field.

Step 10 Click **Save**.

Refer to the video [Snort 3 Suppression and Threshold](#) for additional support and information.

What to do next

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

View and Delete Intrusion Event Thresholds

To view or delete an existing threshold setting for a rule, use the Rules Details view to display the configured settings for a threshold and see if they are appropriate for your system. If they are not, you can add a new threshold to overwrite the existing values.

Procedure

Step 1 Choose **Objects > Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Choose the rule with a configured threshold as shown in the **Alert Configuration** column (the **Alert Configuration** column displays **Threshold** as a link for the rule).

Step 4 To remove the threshold for the rule, click **Threshold** link in the **Alert Configuration** column.

Step 5 Click **Edit** (✎).

Step 6 Click **Threshold** tab.

Step 7 Click **Reset**.

Step 8 Click **Save**.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Intrusion Policy Suppression Configuration

You can suppress intrusion event notification when a specific IP address or range of IP addresses triggers a specific rule or inspector. This is useful for eliminating false positives. For example, if you have a mail server

that transmits packets that look like a specific exploit, you might suppress event notification for that event when it is triggered by your mail server. The rule triggers for all packets, but you only see events for legitimate attacks.

Intrusion Policy Suppression Types

Note that you can use intrusion event suppression alone or in any combination with rate-based attack prevention, the `detection_filter` keyword, and intrusion event thresholding.



Tip You can add suppressions from within the packet view of an intrusion event. You can also access suppression settings by using the **Alert Configuration** column on the intrusion rules editor page (**Objects > Intrusion Rules > Snort 3 All Rules**).

Set Suppression for an Intrusion Rule in Snort 3

You can set one or more suppressions for a rule in your intrusion policy.

Before you begin

Ensure you create the required network objects to be added for source or destination suppression.

Procedure

-
- | | |
|---------------|---|
| Step 1 | Choose Objects > Intrusion Rules . |
| Step 2 | Click Snort 3 All Rules tab. |
| Step 3 | Click the None link in the intrusion rule's Alert Configuration column. |
| Step 4 | Click Edit (✎). |
| Step 5 | From the Suppressions tab, click the add icon Add (+) next to any of the following options: <ul style="list-style-type: none">• Choose Source Networks to suppress events generated by packets originating from a specified source IP address.• Choose Destination Networks to suppress events generated by packets going to a specified destination IP address. |
| Step 6 | Select any of the preset networks in the Network drop-down list. |
| Step 7 | Click Save . |
| Step 8 | (Optional) Repeat the last three steps if required. |
| Step 9 | Click Save in the Alert Configuration window. |
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

View and Delete Suppression Conditions

You may want to view or delete an existing suppression condition. For example, you can suppress event notification for packets originating from a mail server IP address because the mail server normally transmits packets that look like exploits. If you then decommission that mail server and reassign the IP address to another host, you should delete the suppression conditions for that source IP address.

Procedure

-
- Step 1** Choose **Objects > Intrusion Rules**.
 - Step 2** Click **Snort 3 All Rules** tab.
 - Step 3** Choose the rule for which you want to view or delete suppressions.
 - Step 4** Click **Suppression** in the **Alert Configuration** column.
 - Step 5** Click **Edit** (✎).
 - Step 6** Click **Suppressions** tab.
 - Step 7** Remove the suppression by clicking **Clear** (✕) next to the suppression.
 - Step 8** Click **Save**.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Add Intrusion Rule Comments

You can add comments to rules in your intrusion policy. Comments added this way are policy-specific; that is, comments you add to a rule in one intrusion policy are not visible in other intrusion policies.

Procedure

-
- Step 1** Choose **Policies > Intrusion**.
 - Step 2** Click **Snort 3 Version** next to the policy you want to edit.
 - Step 3** In the right side of the page where all the rules are listed, choose the rule where you want to add a comment.
 - Step 4** Click **Comment** (🗨) under the **Comments** column.
 - Step 5** In the **Comments** field, enter the rule comment.
 - Step 6** Click **Add Comment**.
 - Step 7** Click **Save**.

Tip

The system displays a **Comment** (🗨️) next to the rule in the Comments column.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#), on page 23.

Snort 2 Custom Rules Conversion to Snort 3

If you are using custom rules, make sure you are prepared to manage that rule set for Snort 3 prior to conversion from Snort 2 to Snort 3. If you are using a rule set from a third-party vendor, contact that vendor to confirm that their rules will successfully convert to Snort 3 or to obtain a replacement rule set written natively for Snort 3. If you have custom rules that you have written yourself, familiarize with writing Snort 3 rules prior to conversion, so you can update your rules to optimize Snort 3 detection after conversion. See the links below to learn more about writing rules in Snort 3.

- <https://blog.snort.org/2020/08/how-rules-are-improving-in-snort-3.html>
- <https://blog.snort.org/2020/10/talos-transition-to-snort-3.html>

You can refer to other blogs at <https://blog.snort.org/> to learn more about Snort 3 rules.



Important

Snort 2 network analysis policy (NAP) settings *cannot* be copied to Snort 3 automatically. NAP settings have to be manually replicated in Snort 3.

Convert all Snort 2 Custom Rules across all Intrusion Policies to Snort 3

Procedure

Step 1 Choose **Objects > Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Ensure **All Rules** is selected in the left pane.

Step 4 Click the **Tasks** drop-down list and choose:

- **Convert Snort 2 rules and import**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and import them into management center as Snort 3 custom rules.
- **Convert Snort 2 rules and download**—To automatically convert all the Snort 2 custom rules across all the intrusion policies to Snort 3 and download them into your local system.

Step 5 Click **OK**.

Note

- If you selected **Convert and import** in the previous step, then all the converted rules are saved under a newly created rule group **All Snort 2 Converted Global** under **Local Rules**.
- If you selected **Convert and download** in the previous step, then save the rules file locally. You can review the converted rules in the downloaded file and later upload them by following the steps in [Add Custom Rules to Rule Groups](#), on page 49.

Refer to the video [Converting Snort 2 Rules to Snort 3](#) for additional support and information.

What to do next


Deploy configuration changes; see [Deploy Configuration Changes](#), on page 23.

Convert Snort 2 Custom Rules of a Single Intrusion Policy to Snort 3


Procedure

Step 1 Choose **Policies > Intrusion**.

Step 2 In the **Intrusion Policies** tab, click **Show Snort 3 Sync status**.

Step 3 Click the **Sync** icon **Snort out-of-Sync** () of the intrusion policy.

Note

If the Snort 2 and the Snort 3 versions of the intrusion policy are synchronized, then the **Sync** icon is in green **Snort in-Sync** (). It indicates that there are no custom rules to be converted.

Step 4 Read through the summary and click the **Custom Rules** tab.

Step 5 Choose:

- **Import converted rules to this policy**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and import them into management center as Snort 3 custom rules.
- **Download converted rules**—To convert the Snort 2 custom rules in the intrusion policy to Snort 3 and download them into your local system. You can review the converted rules in the downloaded file and later upload the file by clicking the upload icon.

Step 6 Click **Re-Sync**.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#), on page 23.

Add Custom Rules to Rule Groups

Uploading custom rules in the management center adds the custom rules that you have created locally to the list of all the Snort 3 rules.

Procedure

Step 1 Choose **Objects > Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.

Step 3 Click the **Tasks** drop-down list.

Step 4 Click **Upload Snort 3 Rules**.

Step 5 Drag and drop the `.txt` or `.rules` file that contains the Snort 3 custom rules that you have created.

Step 6 Click **OK**.

Note

If there are any errors in the selected file, then you cannot proceed further. You can download the error file and **Replace File** link to upload version 2 of the file, after fixing the errors.

Step 7 Associate rules to a rule group to add the new rules to that group.

You can also create a new custom rule group (by clicking the **Create New Custom Rule Group** link) and then add the rules to the new group.

Note

If there are no existing local rule groups, then proceed by clicking **Create New Custom Rule Group to proceed**. Enter a **Name** for the new rule group and click **Save**.

Step 8 Choose either of the following:

- **Merge Rules** to merge the new rules that you are adding with the existing rules in the rule group.
- **Replace all rules in the group with file contents** to replace all the exiting rules with the new rules that you are adding.

Note

If you chose more than one rule group in the previous step, then only the **Merge Rules** option is available.

Step 9 Click **Next**.

Review the summary to know the new rule IDs that are being added and optionally download it.

Step 10 Click **Finish**.



Important

The rule action of all the uploaded rules is in the disabled state. You have to change them to the required state to ensure the rules are active.

What to do next

- Uploading custom rules in the management center adds the custom rules that you have created to the list of all the Snort 3 rules. To enforce these custom rules on the traffic, add and enable these rules in the required intrusion policies. For information on adding rule groups with custom rules to an intrusion policy, see [Add Rule Groups with Custom Rules to an Intrusion Policy, on page 50](#). For information on enabling custom rules, see [Manage Custom Rules in Snort 3, on page 50](#).
- Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Add Rule Groups with Custom Rules to an Intrusion Policy

Custom rules that are uploaded in the system have to be enabled in an intrusion policy to enforce those rules on the traffic. After uploading custom rules on management center, add the rule group with the new custom rules in the intrusion policy.

Procedure

-
- Step 1** Choose **Policies > Intrusion**.
- Step 2** In the **Intrusion Policies** tab, click the **Snort 3 Version** of the intrusion policy.
- Step 3** Click **Add** (+) next to the Rule Groups search bar.
- Step 4** In the **Add Rule Groups** window, click the **Expand Arrow** (▸) icon next to a rule group to expand the local rule group.
- Step 5** Check the check box next to the uploaded custom rules group.
- Step 6** Click **Save**.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Manage Custom Rules in Snort 3

Custom rules that are uploaded in the system have to be added to an intrusion policy and enabled to enforce those rules on the traffic. You can enable the uploaded custom rules across all policies or selectively on individual policies.

Follow the steps to enable custom rules in one or many intrusion policies:

Procedure

-
- Step 1** Choose **Objects > Intrusion Rules**.
- Step 2** Click **Snort 3 All Rules** tab.
- Step 3** Expand **Local Rules**.

- Step 4** Select the required rule group.
- Step 5** Select the rules by checking the check boxes next to them.
- Step 6** Select **Per Intrusion Policy** from the **Rule Actions** drop-down list.
- Step 7** Choose:
- **All Policies**—to have the same rule actions for all the rules to be added.
 - **Per Intrusion Policy**—to have different rule actions for each intrusion policy.
- Step 8** Set the rule actions:
- If you selected All Policies in the previous step, then select the required rule action from the **Select Override state** drop-down list.
 - If you selected Per Intrusion Policy in the previous step, then select the **Rule Action** against the policy name. To add more policies, click **Add Another**.
- Step 9** Optionally, add a comment in the **Comments** text box.
- Step 10** Click **Save**.

What to do next

Deploy the changes on the device. See, [Deploy Configuration Changes, on page 23](#).

Delete Custom Rules

Procedure

-
- Step 1** Choose **Objects > Intrusion Rules**.
- Step 2** Click **Snort 3 All Rules** tab.
- Step 3** Expand **Local Rules** in the left pane.
- Step 4** Check the check boxes of the rules you want to delete.
- Step 5** Ensure that the rule action for all the rules that you select is **Disable**.
- If required, follow the steps below to disable the rule action for multiple selected rules:
- From the **Rule Actions** drop-down box, select **Per Intrusion Policy**.
 - Select **All Policies** radio button.
 - Select **Disable** from the **Select Override state** drop-down list.
 - Click **Save**.
 - Check the check boxes of the rules you want to delete.
- Step 6** From the **Rule Actions** drop-down list, select **Delete**.
- Step 7** Click **Delete** in the Delete Rules pop-up window.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).

Delete Rule Groups

Before you begin

Exclude the rule group you want to delete from all intrusion policies where you have included it. For steps on excluding a rule group from an intrusion policy, see [Edit Snort 3 Intrusion Policies, on page 31](#).

Procedure

Step 1 Choose **Objects > Intrusion Rules**.

Step 2 Click **Snort 3 All Rules** tab.


Step 3 Expand **Local Rules** in the left pane.

Step 4 Select the rule group to be deleted.

Step 5 Ensure the rule action for all the rules in the group is set to **Disable** before proceeding.

If the rule action for any of the rules is anything other than **Disable**, then you cannot delete the rule group. If required, follow the steps below to disable the rule action for all the rules:

- a) Check the check box below the **Rule Actions** drop-down list to select all the rules in the group.
- b) From the **Rule Actions** drop-down box, select **Per Intrusion Policy**.
- c) Select **All Policies** radio button.
- d) Select **Disable** from the **Select Override state** drop-down list.
- e) Click **Save**.

Step 6 Click the **Delete** () next to the rule group.

Step 7 Click **OK** in the Delete Rule Group pop-up window.

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).



CHAPTER 5

Tailor Intrusion Protection for Your Network Assets

This chapter provides an insight into Secure Firewall recommended rules and generating and applying Secure Firewall recommended rules.

- [Snort 3 Rule Changes in LSP Updates](#) , on page 53
- [Overview of Secure Firewall Recommended Rules](#), on page 53
- [Prerequisites for Network Analysis and Intrusion Policies](#), on page 54
- [Generate New Secure Firewall Recommendations in Snort 3](#), on page 55

Snort 3 Rule Changes in LSP Updates

During regular Snort 3 Lightweight Security Package (LSP) updates, an existing system-defined intrusion rule may be replaced with a new intrusion rule. There could be possibilities of a single rule being replaced with multiple rules, or multiple rules being replaced with a single rule. This occurs when better detection is possible for which rules are combined or expanded. For better management, some existing system-defined rules may also be removed as a part of the LSP update.

To get notifications for changes to any *overridden* system-defined rules during LSP updates, ensure that the **Retain user overrides for deleted Snort 3 rules** check box is checked.

To navigate to the **Retain user overrides for deleted Snort 3 rules** check box, click **System** (⚙️), and then choose **Configuration > Intrusion Policy Preferences**.

By default this check box is checked. When this check box is checked, the system retains the rule overrides in the new replacement rules that are added as a part of the LSP update. The notifications are shown in the **Tasks** tab under the Notifications icon that is located next to **System** (⚙️).

Overview of Secure Firewall Recommended Rules

You can use intrusion rule recommendations to target vulnerabilities associated with host assets detected in the network. For example, operating systems, servers, and client application protocols. This allows you to tailor your intrusion policy to the specific needs of your monitored network.

The system makes an individual set of recommendations for each intrusion policy. It typically recommends rule state changes for standard text rules and shared object rules. However, it can also recommend changes for inspector and decoder rules.

When you generate rule state recommendations, you can use the default settings or configure advanced settings. Advanced settings allow you to:

- Redefine which hosts on your network the system monitors for vulnerabilities
- Influence which rules the system recommends based on rule overhead
- Specify whether to generate recommendations to disable rules

You can also choose to use the recommendations immediately or review the recommendations (and affected rules) before accepting them.

Choosing to use recommended rule states adds a read-only Secure Firewall Recommendations layer to your intrusion policy, and subsequently choosing not to use recommended rule states removes the layer.

You can schedule a task to generate recommendations automatically based on the most recently saved configuration settings in your intrusion policy.

The system does not change rule states that you set manually such as:

- Manually setting the states of specified rules *before* you generate recommendations prevents the system from modifying the states of those rules in the future.
- Manually setting the states of specified rules *after* you generate recommendations overrides the recommended states of those rules.



Tip The intrusion policy report can include a list of rules with rule states that differ from the recommended state.

While displaying the recommendation-filtered Rules page, or after accessing the Rules page directly from the navigation panel or the Policy Information page, you can manually set rule states, sort rules, and take any of the other actions available on the Rules page, such as suppressing rules, setting rule thresholds, and so on.



Note The Cisco Talos Intelligence Group (Talos) determines the appropriate state of each rule in the system-provided policies. If you use a system-provided policy as your base policy, and you allow the system to set your rules to the Secure Firewall recommended rule state, the rules in your intrusion policy match the settings recommended for your network assets.

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Generate New Secure Firewall Recommendations in Snort 3

Generate the Secure Firewall recommendations for the intrusion policy and then follow the steps that are listed here to create new recommended rule settings in Snort 3. Rule overheads are interpreted as **security levels** based on the threshold policies selected by you in Snort 3. Recommended action is based on the selected security level and if it is higher than the base policy, then the recommendation is not just limited to generating the events.

Prior to setting the Secure Firewall recommendations you should ask which of the three points listed below closely matches the goal:

- **Increased Protection**—Enable additional rules based on vulnerabilities found in the host database and do not automatically disable any rules. This will likely result in a larger rule set.
- **Focused Protection**—Enable additional rules and disable existing rules based on vulnerabilities found in the host database. This can increase or decrease the number of rules depending on vulnerabilities discovered.
- **Higher Efficiency**—Use the currently enabled rule set and disable any rules for vulnerabilities not found in the host database. This will likely result in a smaller enabled rule set.

Based on the response, the recommendation actions are as follows:

- Set recommendations to the next highest security level, and uncheck the disable rules.
- Set recommendations to the next highest security level, and check the disable rules.
- Set recommendations to the current security level, and check the disable rules.

Before you begin

Secure Firewall recommendations have the following requirements:

- Ensure that hosts are present in the system to generate recommendations.
- Protected networks configured for recommendations should map to the hosts present in the system

Procedure

Step 1 Choose **Policies > Intrusion**.

Step 2 Click the **Snort 3 Version** button of the intrusion policy.

Step 3 Click the **Recommendations (Not in Use)** layer to configure the rule recommendations.

In the Secure Firewall Rule Recommendations window you can set the following:

- **Security Level:** Click to select the security level. Optionally, you can check the **Accept Recommendation to Disable Rules** checkbox to disable rules that are not enabled at the input security level and in protected networks. Only enable this option if you need to trim your rule set due to a high number of alerts or to improve inspection performance. The security levels are:

- Security level 1: Connectivity Over Security

No Impact—No new rules are enabled and no existing rules are disabled. To increase the protection, select a higher security level.

Lower Security (checkbox is checked)—All rules are disabled except for the rules in the Connectivity Over Security ruleset that match potential vulnerabilities on discovered hosts. It is recommended instead to adjust the Base Policy.

- Security level 2: Balanced Security Over Connectivity

No Impact—No new rules are enabled and no existing rules are disabled. To increase the protection, select a higher security level.

Higher Efficiency (checkbox is checked)—Keeps existing rules that match potential vulnerabilities on discovered hosts and disables rules for vulnerabilities not found on the network.

- Security level 3: Security Over Connectivity

Increased Security—Enables additional rules that match potential vulnerabilities on discovered hosts based on the Maximum Detection ruleset.

Focused Security (checkbox is checked)—Enables additional rules that match vulnerabilities on discovered hosts based on the Security Over Connectivity ruleset, while disabling existing rules that do not match potential vulnerabilities on discovered hosts.

- Security level 4 : Maximum Detection

Increased Security—Enables additional rules that match potential vulnerabilities on discovered hosts based on the Security Over Connectivity ruleset.

Focused Security (checkbox is checked)—Enables additional rules that match vulnerabilities on discovered hosts based on the Maximum Detection ruleset, while disabling existing rules that do not match potential vulnerabilities on discovered hosts.

Note

Maximum Detection enables a very high number of rules and may impact performance. We recommend you to review and test this setting before deploying into a production environment.

- **Protected Networks:** Specifies the monitored networks or individual hosts to examine for recommendations. You can select one or more system or custom defined network objects from the drop-down list. By default, any IPv4 or IPv6 networks are selected, if no selection is done.

Important

The Secure Firewall Rule Recommendations depend on network discovery. Protected Networks apply to any hosts discovered within the ranges configured in your Network Discovery policy. For more information, see the chapter [Network Discovery Policies](#) in the *Cisco Secure Firewall Management Center Device Configuration guide*.

Click the **Add** + button to create a new network object of type Host or Network and click **Save**.

Step 4 Generate and apply recommendations:

- **Generate:** Generates the recommendations for an intrusion policy. This action lists the rules under Recommended Rules (Not in use).
- **Generate and Apply:** Generates and applies the recommendations for an intrusion policy. This action lists the rules under Recommended Rules (In use).

Recommendations are generated successfully. A new recommendation tab appears with all the recommended rules with their corresponding recommended actions. Rule action preset filters are also available for this tab, in addition with new recommendations.

Step 5

You can verify the recommendations and then choose to apply them accordingly:

- **Accept**—Applies the previously generated recommendations for an Intrusion policy.
- **Refresh**—Regenerates and updates the rule recommendations for an Intrusion policy.
- **Edit**—It opens the Recommendations dialog box, you can provide the recommendation input values and then generate the recommendations.
- **Remove All**—Revert or remove the applied recommended rules from the policy and also removes the recommendation tab.

Under **All Rules**, there is a Recommended Rules section which displays the recommended rules.

Note

Final action for an Intrusion rule is applied based on the rule action priority order and following is the rule action priority order:

Rule Override > Generated Recommendations > Group Override > Base Policy Default Action

For enabled recommendations, management center considers the current state: group overrides, base policy, and recommendation configurations and priority order of actions is:

pass > block > reject > drop > rewrite > alert

What to do next

Deploy configuration changes; see [Deploy Configuration Changes, on page 23](#).



PART II

Advanced Network Analysis in Snort 3

- [Get Started with Network Analysis Policies, on page 61](#)



CHAPTER 6

Get Started with Network Analysis Policies

This chapter provides an insight into network analysis policy basics, prerequisites, and how to manage network analysis policies. It also provides information on custom network analysis policy creation and network analysis policy settings.

- [Overview of Network Analysis Policies, on page 61](#)
- [Manage Network Analysis Policies, on page 62](#)
- [Snort 3 Definitions and Terminologies for Network Analysis Policy, on page 62](#)
- [Prerequisites for Network Analysis and Intrusion Policies, on page 65](#)
- [Custom Network Analysis Policy Creation for Snort 3, on page 65](#)
- [Network Analysis Policy Settings and Cached Changes, on page 90](#)

Overview of Network Analysis Policies

Network analysis policies govern many traffic preprocessing options, and are invoked by advanced settings in your access control policy. Network analysis-related preprocessing occurs after Security Intelligence matching and SSL decryption, but before intrusion or file inspection begins.

By default, the system uses the *Balanced Security and Connectivity* network analysis policy to preprocess all traffic handled by an access control policy. However, you can choose a different default network analysis policy to perform this preprocessing. For your convenience, the system provides a choice of several non-modifiable network analysis policies, which are tuned for a specific balance of security and connectivity by the Cisco Talos Intelligence Group (Talos). You can also create a custom network analysis policy with custom preprocessing settings.



Tip System-provided intrusion and network analysis policies are similarly named but contain different configurations. For example, the Balanced Security and Connectivity network analysis policy and the Balanced Security and Connectivity intrusion policy work together and can both be updated in intrusion rule updates. However, the network analysis policy governs mostly preprocessing options, whereas the intrusion policy governs mostly intrusion rules. Network analysis and intrusion policies work together to examine your traffic.

You can also tailor traffic preprocessing options to specific security zones, networks, and VLANs by creating multiple custom network analysis policies, then assigning them to preprocess different traffic. (Note that ASA FirePOWER cannot restrict preprocessing by VLAN.)

Manage Network Analysis Policies

Under your user name in the toolbar, the system displays a tree of available domains. To switch domains, choose the domain you want to access.

Procedure

Step 1 Choose one of the following paths to access the network analysis policy.

- **Policies > Access Control heading > Access Control**, then click **Network Analysis Policy**
- **Policies > Access Control heading > Intrusion**, then click **Network Analysis Policies**
- **Policies > Intrusion > Network Analysis Policies**

Note

If your custom user role limits access to the first path listed here, use the second path to access the policy.

Step 2 Manage your network analysis policy:

- **Create**—If you want to create a new network analysis policy, click **Create Policy**.

Two versions of the network analysis policy are created, a **Snort 2 Version** and a **Snort 3 Version**.

- For the Snort 2 version, see *Custom Network Analysis Policy Creation for Snort 2* in the *Cisco Secure Firewall Management Center Configuration Guide*.
- For the Snort 3 version, see [Custom Network Analysis Policy Creation for Snort 3, on page 65](#).
- **Delete**—If you want to delete a network analysis policy, click the **Delete** icon, then confirm that you want to delete the policy. You cannot delete a network analysis policy if an access control policy references it.
If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- **Edit**—If you want to edit an existing network analysis policy, click the **Edit** icon.
If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- **Report**—Click the **Report** icon; see *Generating Current Policy Reports* in the *Cisco Secure Firewall Management Center Configuration Guide*.

Snort 3 Definitions and Terminologies for Network Analysis Policy

The following table lists the Snort 3 concepts and terms used in the Network Analysis Policy.

Table 6: Snort 3 Definitions and Terminologies for Network Analysis Policy

Term	Description
Inspectors	Inspectors are plugins that process packets (similar to the Snort 2 preprocessor).
Binder inspector	<p>Binder inspector defines the flow when a particular inspector has to be accessed and taken into consideration.</p> <p>When the traffic matches the conditions defined in the binder inspector, only then do the values/configurations for that inspector come into effect.</p> <p>For more information, see <i>Binder Inspector</i> in Custom Network Analysis Policy Creation for Snort 3, on page 65.</p>
Singleton inspectors	<p>Singleton inspectors contain one instance. These inspectors do not support adding more instances like multiton inspectors. Settings of singleton inspector are applied to the entire traffic matching that inspector and not to a specific traffic segment.</p> <p>For more information, see <i>Singleton Inspectors</i> in Custom Network Analysis Policy Creation for Snort 3, on page 65.</p>
Multiton inspectors	<p>Multiton inspectors contain multiple instances which you can configure as needed. These inspectors support configuring settings based on specific conditions, such as network, port, and VLAN. One set of supported settings is called an instance.</p> <p>For more information, see <i>Multiton Inspectors</i> in Custom Network Analysis Policy Creation for Snort 3, on page 65.</p>
Schema	<p>The schema file is based on the OpenAPI JSON specification, and it validates the content that you upload or download. You can download the schema file and open it using any third-party JSON editor, such as Swagger editor. The schema file helps you to identify what parameters can be configured for inspectors with their corresponding allowed values, range, and accepted patterns to be used.</p> <p>For more information, see Customize the Network Analysis Policy, on page 71.</p>

Term	Description
Sample file	<p>It is a pre-existing template that contains example configurations to help you with configuring the inspectors.</p> <p>You can refer to the example configurations included in the sample file and make any changes that you may require.</p> <p>For more information, see Customize the Network Analysis Policy, on page 71.</p>
Full configuration	<p>You can download the entire inspector configurations in a single file.</p> <p>All information regarding the inspector configuration is available in this file.</p> <p>The full configuration is a merged configuration of the default configuration (rolled out as a part of the LSP updates by Cisco Talos) and the custom NAP inspector configurations.</p> <p>For more information, see Customize the Network Analysis Policy, on page 71.</p>
Overridden configuration	<p>In the Snort 3 Version of the network analysis policy page:</p> <ul style="list-style-type: none"> • Under Actions > Upload, you can click Overridden Configuration to upload the JSON file that contains the overridden configuration. • Under Actions > Download, you can click Overridden Configuration to download the inspector configuration that has been overridden. <p>If you have not overridden any inspector configuration, then this option is disabled. When you override the inspector configuration, then this option is enabled automatically to allow you to download.</p> <p>For more information, see Customize the Network Analysis Policy, on page 71.</p>

Related Topics

[Custom Network Analysis Policy Creation for Snort 3, on page 65](#)

[Customize the Network Analysis Policy, on page 71](#)

[Network Analysis Policy Mapping, on page 68](#)

Prerequisites for Network Analysis and Intrusion Policies

To allow the Snort inspection engine to process traffic for intrusion and malware analysis, you must have the IPS license enabled for the Threat Defense device.

You must be an Admin user to manage network analysis, intrusion policies, and perform migration tasks.

Custom Network Analysis Policy Creation for Snort 3

The default network analysis policy is tuned for typical network requirements and optimal performance. Usually, the default network analysis policy suffices most network requirements and you might not need to customize the policy. However, when you have a specific network requirement or when you are facing performance issues, the default network analysis policy can be customized. Note that customizing the network analysis policy is an advanced configuration that should be done only by advanced users or Cisco support.

Network analysis policy configuration for Snort 3 is a data-driven model, which is based on JSON and JSON Schema. Schema is based on the OpenAPI specification, and it helps you get a view of the supported inspectors, settings, settings type, and valid values. The Snort 3 inspectors are plugins that process packets (similar to the Snort 2 preprocessor). Network analysis policy configuration is available to download in the JSON format.

In Snort 3, the list of inspectors and settings are not in a one-to-one mapping with the Snort 2 list of preprocessors and settings. Also, the number of inspectors and settings available in management center is a subset of the inspectors and settings that Snort 3 supports. See <https://snort.org/snort3> for more information on Snort 3. See <https://www.cisco.com/go/snort3-inspectors> for more information on the inspectors available in management center.



Note

- While upgrading the management center to the 7.0 release, the changes that were done in the Snort 2 version of the network analysis policy are not migrated to Snort 3 after the upgrade.
 - Unlike the intrusion policy, there is no option to synchronize Snort 2 network analysis policy settings to Snort 3.
-

Default Inspector Updates

Lightweight Security Package (LSP) updates may contain new inspectors or modifications to integer ranges for existing inspector configurations. Following the installation of an LSP, new inspectors and/or updated ranges will be available under **Inspectors** in the **Snort 3 Version** of your network analysis policy.

Binder Inspector

Binder inspector defines the flow when a particular inspector has to be accessed and taken into consideration. When the traffic matches the conditions defined in the binder inspector, only then the values/configurations for that inspector come into effect. For example:

For the *imap* inspector, the binder defines the following condition when it has to be accessed. That is when:

- Service is equal to imap.
- Role is equal to any.

If these conditions are met, then use the type imap.

▼ binder	
185	{
186	"when": {
187	"service": "imap",
188	"role": "any"
189	},
190	"use": {
191	"type": "imap"
192	}
193	},

Singleton Inspectors

Singleton inspectors contain a single instance. These inspectors do not support adding more instances like multiton inspectors. Settings of singleton inspector are applied to the entire traffic and not to a specific traffic segment.

For example:

```
{
  "normalizer":{
    "enabled":true,
    "type":"singleton",
    "data":{
      "ip4":{
        "df":true
      }
    }
  }
}
```


Multiton Inspectors

Multiton inspectors contain multiple instances which you can configure as needed. These inspectors support configuring settings based on specific conditions, such as network, port, and VLAN. One set of supported settings is called an instance. There is a default instance, and you can also add additional instances based on specific conditions. If the traffic matches that condition, the settings from that instance are applied. Otherwise, the settings from the default instance are applied. Also, the name of the default instance is the same as the inspector's name.

For a multiton inspector, when you upload the overridden inspector configuration, you also need to include/define a matching binder condition (conditions under when the inspector has to be accessed or used) for each instance in the JSON file, otherwise, the upload will result in an error. You can also create new instances, but make sure that you include a binder condition for every new instance that you create to avoid errors.

For example:

- Multiton inspector where the default instance is modified.

```
{
  "http_inspect":{
    "enabled":true,
    "type":"multiton",
    "instances":[
      {
        "name":"http_inspect",
        "data":{
          "response_depth":5000
        }
      }
    ]
  }
}
```

- Multiton inspector where the default instance and default binder is modified.

```
{
  "http_inspect":{
    "enabled":true,
    "type":"multiton",
    "instances":[
      {
        "name":"http_inspect",
        "data":{
          "response_depth":5000
        }
      }
    ]
  },
  "binder":{
    "type":"binder",
    "enabled":true,
    "rules":[
      {
        "use":{
          "type":"http_inspect"
        },
        "when":{
          "role":"any",
          "ports":"8080",
          "proto":"tcp",
          "service":"http"
        }
      }
    ]
  }
}
```

```

    }
  ]
}

```

- Multiton inspector where a custom instance and a custom binder is added.

```

{
  "http_inspect":{
    "enabled":true,
    "type":"multiton",
    "instances":[
      {
        "name":"http_inspect1",
        "data":{
          "response_depth":5000
        }
      }
    ]
  },
  "binder":{
    "type":"binder",
    "enabled":true,
    "rules":[
      {
        "use":{
          "type":"http_inspect",
          "name":"http_inspect1"
        },
        "when":{
          "role":"any",
          "ports":"8080",
          "proto":"tcp",
          "service":"http"
        }
      }
    ]
  }
}

```

Network Analysis Policy Mapping

For network analysis policies, Cisco Talos provides mapping information, which is used to find the corresponding Snort 2 version of the policies for the Snort 3 version.

This mapping ensures that the Snort 3 version of policies has its equivalent Snort 2 version.

View Network Analysis Policy Mapping

Procedure

Step 1 Go to **Policies > Intrusion > Network Analysis Policies**.

Step 2 Click **NAP Mapping**.

Step 3 Expand the arrow for **View Mappings**.

The Snort 3 network analysis policies that are automatically mapped to a Snort 2 equivalent policy are displayed.

Step 4 Click **OK**.

Create a Network Analysis Policy

All the existing network analysis policies are available in management center with their corresponding Snort 2 and Snort 3 versions. When you create a new network analysis policy, it is created with both the Snort 2 version and the Snort 3 version.

Procedure

Step 1 Go to **Policies > Intrusion > Network Analysis Policies**.

Step 2 Click **Create Policy**.

Step 3 Enter the **Name** and **Description**.

Step 4 Choose the **Inspection Mode** from the available choices.

- **Detection**
- **Prevention**

Step 5 Select a **Base Policy** and click **Save**.

Note

Configure Network Analysis Policy (NAP) in **Prevention** mode if you are using Snort 3 and SSL Decryption or TLS Server Identity.

The new network analysis policy is created with its corresponding **Snort 2 Version** and **Snort 3 Version**.

Modify the Network Analysis Policy

You can modify the network analysis policy to change its name, description, or the base policy.

Procedure

Step 1 Go to **Policies > Intrusion > Network Analysis Policies**.

Step 2 Click **Edit** to change the name, description, inspection mode, or the base policy.

Note

If you edit the network analysis policy name, description, base policy, and inspection mode, the edits are applied to both the Snort 2 and Snort 3 versions. If you want to change the inspection mode for a specific version, then you can do that from within the network analysis policy page for that respective version.

Step 3 Click **Save**.

Search for an Inspector on the Network Analysis Policy Page

On the Snort 3 version of the network analysis policy page, you may need to search for an inspector by entering any relevant text in the search bar.

Procedure

-
- Step 1** Go to **Policies > Intrusion > Network Analysis Policies**.
- Step 2** Go to the **Snort 3 Version** of the network analysis policy.
- Step 3** Enter an inspector's name or any relevant text to search for in the **Search** bar.

All the inspectors matching the text you search for are displayed.

For example, if you enter **pop**, then the pop inspector and the binder inspector are shown as matching results on the screen.

Related Topics

- [Examples of Custom Network Analysis Policy Configuration](#), on page 79
- [View the List of Inspectors with Overrides](#), on page 75
- [Snort 3 Definitions and Terminologies for Network Analysis Policy](#), on page 62
- [Customize the Network Analysis Policy](#), on page 71
- [Make Inline Edit for an Inspector to Override Configuration](#), on page 74

Copy the Inspector Configuration

You can copy the inspector configuration for the Snort 3 version of the network analysis policy according to your requirements.

Procedure

-
- Step 1** Go to **Policies > Intrusion > Network Analysis Policies**.
- Step 2** Go to the **Snort 3 Version** of the network analysis policy.
- Step 3** Under **Inspectors**, expand the required inspector for which you want to copy the configuration.
- The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.
- Step 4** Click the **Copy to clipboard** icon to copy the inspector configuration to the clipboard for one or both of the following.
- **Default Configuration** in the left column
 - **Overridden Configuration** in the right column
- Step 5** Paste the copied inspector configuration to a JSON editor to make any edits you may require.
-

Related Topics

[Customize the Network Analysis Policy](#), on page 71

Customize the Network Analysis Policy

You can customize the Snort 3 version of the network analysis policy according to your requirements.

Procedure

Step 1 Go to **Policies > Intrusion > Network Analysis Policies**.

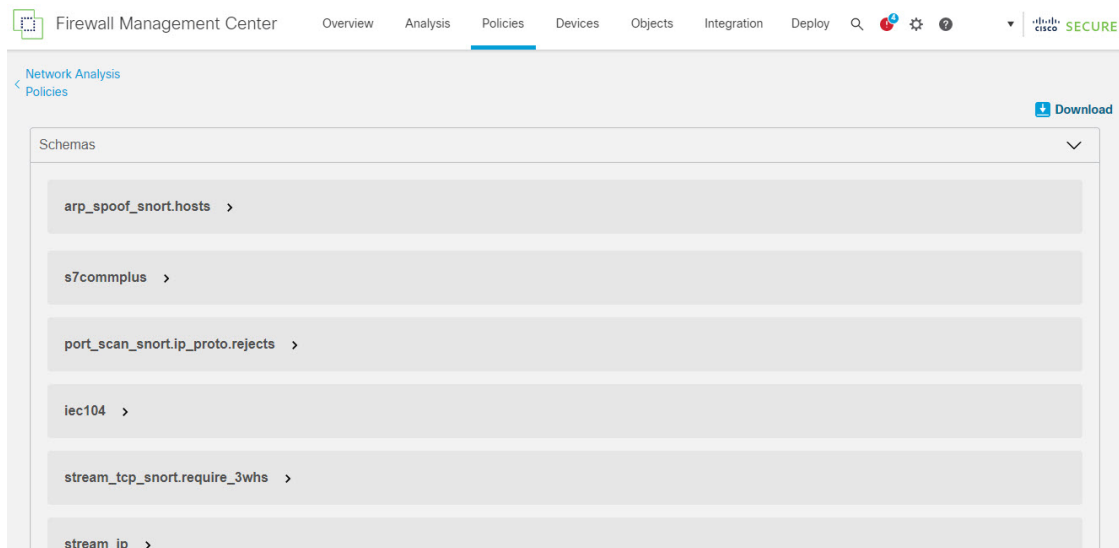
Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Click the **Actions** drop-down menu.

The following options are displayed:

- View Schema
- Download Schema / Download Sample File / Template
- Download Full Configuration
- Download Overridden Configuration
- Upload Overridden Configuration

Step 4 Click **View Schema** to open the schema file directly in a browser.



Step 5 You can download the schema file, sample file / template, full configuration, or overridden configuration as needed.

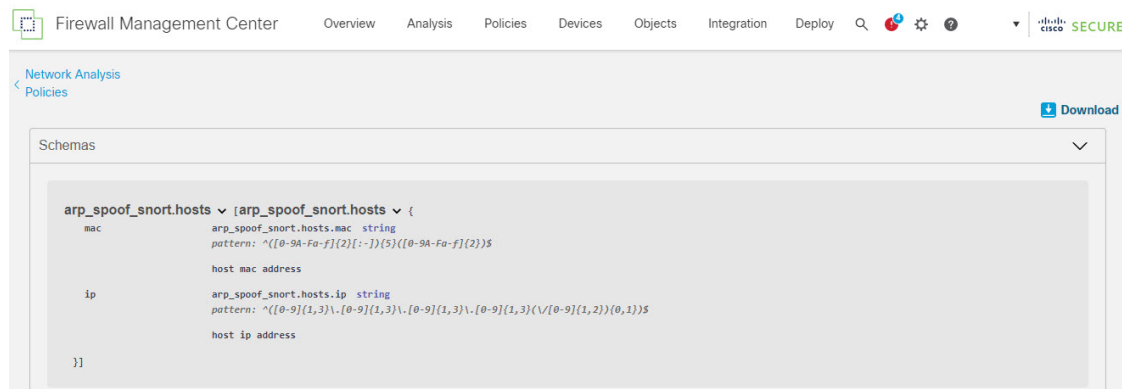
These options provide you an insight about the allowed values, range, and patterns, existing and default inspector configurations, and overridden inspector configurations.

a) Click **Download Schema** to download the schema file.

The schema file validates the content that you upload or download. You can download the schema file and open it using any third-party JSON editor. The schema file helps you to identify what parameters can be configured for inspectors with their corresponding allowed values, range, and accepted patterns to be used.

For example, for the `arp_spoof_snort` inspector, you can configure the hosts. The hosts include the `mac` and `ip` address values. The schema file shows the following accepted pattern for these values.

- **mac-pattern:** `^([0-9A-Fa-f]{2}[:-]){5}([0-9A-Fa-f]{2})$`
- **ip-pattern:** `^([0-9]{1,3}\. [0-9]{1,3}\. [0-9]{1,3}\. [0-9]{1,3}){0,1}$`



You must provide the values, range, patterns according to the accepted ones in the schema file to be able to successfully override the inspector configuration, otherwise, you get an error message.

- Click **Download Sample File / Template** to use a pre-existing template that contains example configurations to help you with configuring the inspectors.

You can refer to the example configurations included in the sample file and make any changes that you may require.

- Click **Download Full Configuration** to download the entire inspector configurations in a single JSON file.

Instead of expanding the inspectors separately, you can download the full configuration to look out for the information you need. All information regarding the inspector configuration is available in this file.

- Click **Download Overridden Configuration** to download the inspector configuration that has been overridden.

Step 6 To override the existing configuration, follow the steps.

You can choose to override an inspector configuration using the following ways.

- Make inline edits for an inspector directly on the management center. See the topic **Make Inline Edit for an Inspector to Override Configuration** in the **Getting Started with Network Analysis Policies** chapter of the *Cisco Secure Firewall Management Center Snort 3 Configuration Guide*.
- Continue to follow the current procedure of using the **Actions** drop-down menu to upload the overridden configuration file.

If you chose to make inline edits directly in the management center, then you don't need to follow the current procedure further. Otherwise, you must follow this procedure completely.

- Under **Inspectors**, expand the required inspector for which you want to override the default configuration.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

You may need to search for an inspector by entering any relevant text in the search bar.

- b) Click the **Copy to clipboard** icon to copy the default inspector configuration to the clipboard.
- c) Create a JSON file and paste the default configuration in it.
- d) Keep the inspector configuration that you want to override, and remove all the other configuration and instances from the JSON file.

You can also use the **Sample File / Template** to understand how to override the default configuration. This is a sample file that includes JSON snippets explaining how you can customize the network analysis policy for Snort 3.

- e) Make changes to the inspector configuration as needed.

Validate the changes and make sure they conform to the schema file. For multiton inspectors, make sure that the binder conditions for all instances are included in the JSON file. See *Multiton Inspectors* in the topic **Custom Network Analysis Policy Creation for Snort 3** in the *Cisco Secure Firewall Management Center Snort 3 Configuration Guide* for more information.

- f) If you are copying any further default inspector configurations, append that inspector configuration to the existing file that contains the overridden configuration.

Note

The copied inspector configuration must comply with the JSON standards.

- g) Save the overridden configuration file to your system.

Step 7

From the **Actions** drop-down menu, choose Upload Overridden Configuration to upload the JSON file that contains the overridden configuration.

Caution

Upload only the changes that you require. You should not upload the entire configuration as it makes the overrides sticky in nature and therefore, any subsequent changes to the default configuration as part of the LSP updates would not be applied.

You can drag and drop a file or click to browse to the JSON file saved in your system that contains the overridden inspector configuration.

- **Merge inspector overrides** – Content in the uploaded file is merged with the existing configuration if there is no common inspector. If there are common inspectors, then the content in the uploaded file (for common inspectors) takes precedence over the previous content, and it replaces the previous configuration for those inspectors.
- **Replace inspector overrides** – Removes all previous overrides and replaces them with the new content in the uploaded file.

Attention

Choosing this option deletes all the previous overrides. Make an informed decision before you override the configuration using this option.

If any error occurs while uploading the overridden inspectors, you see the error in the **Upload Overridden Configuration File** pop-up window. You can also download the file with the error, fix the error, and reupload the file.

Step 8

In the **Upload Overridden Configuration File** pop-up window, click **Import** to upload the overridden inspector configuration.

After you upload the overridden inspector configuration, you will see an orange icon next to the inspector that signifies that it is an overridden inspector.

Also, the **Overridden Configuration** column under the inspector shows the overridden value.

You can also view all the overridden inspectors using the **Show Overrides Only** checkbox adjacent to the Search bar.

Note

Make sure that you always download the overridden configuration, open the JSON file, and append any new changes/overrides to the inspector configurations to this file. This action is needed so that you do not lose the old overridden configurations.

Step 9 (Optional) Take a backup of the overridden configuration file on your system before making any new inspector configuration changes.

Tip

We recommend that you take the backup from time to time as you override the inspector configuration.

Related Topics

[Revert Overridden Configuration to Default Configuration](#), on page 76

[View the List of Inspectors with Overrides](#), on page 75

[Search for an Inspector on the Network Analysis Policy Page](#), on page 70

[Copy the Inspector Configuration](#), on page 70

Make Inline Edit for an Inspector to Override Configuration

For the Snort 3 version of the network analysis policy, you can make an inline edit for the inspector configuration to override the configuration according to your requirements.

Alternatively, you can also use the **Actions** drop-down menu to upload the overridden configuration file. See [Customize the Network Analysis Policy, on page 71](#) for more information.

Procedure

Step 1 Go to **Policies > Intrusion > Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to override the default setting.

The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.

Step 4 Under the **Overridden Configuration** in the right column, click **Edit Inspector** (Pencil) icon to make changes to the inspector configuration.

The Override Configuration pop-up appears where you can make the required edits.

Note

- Make sure that you keep only those settings that you want to override. If you leave a setting with the same value, that field becomes sticky. This means if that setting is changed in the future by Talos, the current value will be retained.
- If you are adding or deleting any custom instance, make sure that you add or delete a binder rule for that instance in the binder inspector as well.

Step 5 Click **OK**.

If there are any errors according to the JSON standards, it shows you an error message.

Step 6 Click **Save** to save the changes.

If the changes conform to the OpenAPI schema specification, the management center allows you to save the configuration, otherwise, the **Error saving overridden configuration** pop-up appears that shows the errors. You can also download the file with the errors.

Related Topics

[Customize the Network Analysis Policy](#), on page 71

[Revert Unsaved Changes during Inline Edits](#), on page 75

[Revert Overridden Configuration to Default Configuration](#), on page 76

[Examples of Custom Network Analysis Policy Configuration](#), on page 79

Revert Unsaved Changes during Inline Edits

While making inline edits to override the configuration for an inspector, you can revert any unsaved changes. Note that this action reverts all unsaved changes to the most recently saved value, but does not revert the configuration to the default configuration for an inspector.

Procedure

Step 1 Go to **Policies > Intrusion > Network Analysis Policies**.

Step 2 Go to the **Snort 3 Version** of the network analysis policy.

Step 3 Under **Inspectors**, expand the required inspector for which you want to revert the unsaved changes.

The default configuration is displayed in the left column and the overridden configuration is displayed in the right column under the inspector.

Step 4 Under the **Overridden Configuration** on the right column, click the **Cross X** icon to revert any unsaved changes for the inspector.

Alternatively, you can click **Cancel** to cancel the changes.

If you do not have any unsaved changes to the inspector configuration, then this option is not visible.

Related Topics

[Revert Overridden Configuration to Default Configuration](#), on page 76

[Make Inline Edit for an Inspector to Override Configuration](#), on page 74

View the List of Inspectors with Overrides

You can view a list of all the overridden inspectors.

Procedure

-
- Step 1** Go to **Policies > Intrusion > Network Analysis Policies**.
- Step 2** Go to the **Snort 3 Version** of the network analysis policy.
- Step 3** Check the **Show Overrides Only** checkbox adjacent to the Search bar to view the list of overridden inspectors.
- All the overridden inspectors are shown with an orange icon next to their names to help you identify them.
-

Related Topics

- [Search for an Inspector on the Network Analysis Policy Page](#), on page 70
- [Make Inline Edit for an Inspector to Override Configuration](#), on page 74
- [Customize the Network Analysis Policy](#), on page 71

Revert Overridden Configuration to Default Configuration

You can revert any changes that you made to override the default configuration for an inspector. This action reverts the overridden configuration to the default configuration for an inspector.

Procedure

-
- Step 1** Go to **Policies > Intrusion > Network Analysis Policies**.
- Step 2** Go to the **Snort 3 Version** of the network analysis policy.
- Step 3** Under **Inspectors**, expand the required inspector for which you want to revert the overridden configuration.
- The overridden inspectors are shown with the orange icon next to their name.
- The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector. Under the **Overridden Configuration** on the right column, click **Revert to default configuration** (back arrow) icon to revert the overridden configuration for the inspector to the default configuration.
- If you did not make any changes to the default configuration for the inspector, then this option is disabled.
- Step 4** Click **Revert** to confirm the decision.
- Step 5** Click **Save** to save the changes.
- If you do not want to save the changes, you can click **Cancel** or the **Cross X** icon.
-

Related Topics

- [Revert Unsaved Changes during Inline Edits](#), on page 75
- [Customize the Network Analysis Policy](#), on page 71
- [Make Inline Edit for an Inspector to Override Configuration](#), on page 74
- [Examples of Custom Network Analysis Policy Configuration](#), on page 79

Validate Snort 3 Policies

To validate the Snort 3 policies, here is a list of basic information that user can make note of:

- Current version of the management center can manage multiple threat defense versions.
- Current version of management center supports NAP configurations which are not applicable to previous version of threat defense devices.
- Current NAP Policy and validations will work based on the current version support.
- Changes may include content which is not valid for previous versions of threat defenses.
- Policy configuration changes are accepted if they are valid configuration for the current version and which is performed using current Snort 3 binary and NAP schema.
- For previous version threat defenses, validation is performed during deployment using NAP schema and Snort 3 binary for that specific version. If there is any configuration which is not applicable for the given version, user is provided information or warning that we will not deploy the configuration which is not supported on the given version and remaining configuration will get deployed.

In this procedure, when we associate the NAP policy to an Access Control Policy and deploy it on a device, for example any inspector like rate filter configuration is applied to validate the Snort 3 policies.

Procedure

-
- Step 1** **Steps to Override NAP Policy Configuration:** Under **Inspectors** in the **Snort 3 Version** of the network analysis policy, expand the required inspector for which you want to override the default setting.
- The default configuration is displayed on the left column and the overridden configuration is displayed on the right column under the inspector.
- Step 2** Under the **Overridden Configuration** on the right column, click **Edit Inspector** (Pencil) icon to make changes to any inspector like `rate_filter`.
- The Override Configuration pop-up appears where you can make the required edits to the `rate_filter` inspector.
- Step 3** Click **OK**.
- Step 4** Click **Save** to save the changes.
- Alternatively, you can also use the **Actions** drop-down menu to upload the overridden configuration file.
- Step 5** Click the **Actions** drop-down menu in the **Snort 3 Version** of the network analysis policy.
- Step 6** Under **Upload** you can click **Overridden Configuration** to upload the JSON file that contains the overridden configuration.

Caution

Upload only the changes that you require. You should not upload the entire configuration as it makes the overrides sticky in nature and therefore, any subsequent changes to the default configuration as part of the LSP updates will not be applied.

You can drag and drop a file or click to browse to the JSON file saved in your system that contains the overridden inspector configuration.

- **Merge inspector overrides** – Content in the uploaded file is merged with the existing configuration if there is no common inspector. If there are common inspectors, then the content in the uploaded file (for common inspectors) takes precedence over the previous content, and it replaces the previous configuration for those inspectors.
- **Replace inspector overrides** – Removes all previous overrides and replaces them with the new content in the uploaded file.

Attention

As choosing this option deletes all the previous overrides, make an informed decision before you override the configuration using this option.

If any error occurs while uploading the overridden inspectors, you see the error on the **Upload Overridden Configuration File** pop-up window. You can also download the file with the error, then fix the error and reupload the file.

- Step 7** **Steps to Associate NAP Policy to Access Control Policy:** In the access control policy editor, click **Advanced**, then click **Edit** next to the Network Analysis and Intrusion Policies section.
- Step 8** From the **Default Network Analysis Policy** drop-down list, select a default network analysis policy.
- If you choose a user-created policy, you can click **Edit** to edit the policy in a new window. You cannot edit system-provided policies.
- Step 9** Click **OK**.
- Step 10** Click **Save** to save the policy.
- Step 11** Alternatively, in the access control policy editor, click **Advanced**, then click **Edit** next to the Network Analysis and Intrusion Policies section.
- Step 12** Click **Add Rule**.
- Step 13** Configure the rule's conditions by clicking the conditions you want to add.
- Step 14** Click **Network Analysis** and choose the **Network Analysis Policy** you want to use to preprocess the traffic matching this rule.
- Step 15** Click **Add**.
- Step 16** **Deployment:** On the management center menu bar, click **Deploy** and then select **Deployment**.
- Step 17** Identify and choose the devices on which you want to deploy configuration changes.
- **Search**—Search for the device name, type, domain, group, or status in the search box.
 - **Expand**—Click **Expand Arrow** to view device-specific configuration changes to be deployed.
- By selecting the device check box, all the changes for the device, which are listed under the device, are pushed for deployment. However, you can use the **Policy Selection** to select individual policies or configurations to deploy while withholding the remaining changes without deploying them.
- Optionally, use **Show or Hide Policy** to selectively view or hide the associated unmodified policies.
- Step 18** Click **Deploy**.
- Step 19** If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Validation Messages** window. To view complete details, click the arrow icon before the warnings or errors.

Note

It shows a warning that Snort 3 Network analysis policy contains inspectors or attributes that are not valid for this threat defense version, following the invalid settings will be skipped in deployment: Invalid inspectors are : ["rate_filter"] only for devices lower than 7.1 version.

Examples of Custom Network Analysis Policy Configuration

This is a sample file that includes JSON snippets explaining how you can customize the network analysis policy for Snort 3. You can choose to override an inspector configuration using the following ways:

- Make inline edits for an inspector directly on the management center. See [Make Inline Edit for an Inspector to Override Configuration, on page 74](#).
- Use the **Actions** drop-down menu to upload the overridden configuration file. See [Customize the Network Analysis Policy, on page 71](#).

Before you choose any of these options, review all the following details and examples that will help you in defining the network analysis policy overrides successfully. You must read and understand the examples for various scenarios explained here to avoid any risks and errors.

If you choose to override an inspector configuration from the **Actions** drop-down menu, you need to construct a JSON file for the network analysis policy overrides and upload the file.

For overriding an inspector configuration in the network analysis policy, you must upload only the changes that you require. You should not upload the entire configuration because it makes the overrides sticky in nature and therefore, any subsequent changes to the default values or configuration as part of the LSP updates would not be applied.

Here are the examples for various scenarios:

Enabling a Singleton Inspector when the Default State in the Base Policy is Disabled

```
{
  "rate_filter": {
    "enabled": true,
    "type": "singleton",
    "data": []
  }
}
```

Disabling a Singleton Inspector when the Default State in the Base Policy is Enabled

```
{
  "rate_filter": {
    "enabled": false,
    "type": "singleton",
    "data": []
  }
}
```

Enabling a Multiton Inspector when the Default State in the Base Policy is Disabled

```
{
  "ssh": {
    "enabled": true,
    "type": "multiton",
    "instances": []
  }
}
```

Disabling a Multiton Inspector when the Default State in the Base Policy is Enabled

```
{
  "ssh": {
```

```

        "enabled": false,
        "type": "multiton",
        "instances": []
    },
    "iecl04": {
        "type": "multiton",
        "enabled": false,
        "instances": []
    }
}

```

Overriding the Default Value of Specific Setting(s) for Singleton Inspector

```

{
  "normalizer": {
    "enabled": true,
    "type": "singleton",
    "data": {
      "tcp": {
        "block": true
      },
      "ip6": true
    }
  }
}

```

Overriding Specific Setting(s) of a Default Instance (where Instance Name Matches with Inspector Type) in Multiton Inspector

```

{
  "http_inspect": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "data": {
          "unzip": false
        },
        "name": "http_inspect"
      }
    ]
  }
}

```

Adding Binder Rule for a Default Instance with Required Changes



Note Default binder rules can't be edited, they are always appended at the end.

```

{
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "use": {
          "type": "http_inspect"
        },
        "when": {
          "role": "server",

```

```

        "service": "http",
        "dst_nets": "10.1.1.0/24"
      }
    ]
  }
}

```

Adding a New Custom Instance



Note Corresponding binder rule entry must be defined in the binder inspector.

```

{
  "telnet": {
    "enabled": true,
    "type": "multiton",
    "instances": [
      {
        "name": "telnet_my_instance",
        "data": {
          "encrypted_traffic": true
        }
      }
    ]
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_my_instance"
        }
      }
    ]
  }
}

```

Overriding a Singleton Instance, Multiton Default Instance, and Creating a New Multiton Instance in a Single JSON Override

Example to show the following in a single JSON override:

- Overriding a Singleton instance (**normalizer** inspector)
- Overriding a Multiton default instance (**http_inspect** inspector)
- Creating a new Multiton instance (**telnet** inspector)

```

{
  "normalizer": {
    "enabled": true,
    "type": "singleton",
    "data": {

```

```

        "tcp": {
            "block": true
        },
        "ip6": true
    }
},
"http_inspect": {
    "enabled": true,
    "type": "multiton",
    "instances": [
        {
            "data": {
                "unzip": false,
                "xff_headers": "x-forwarded-for true-client-ip x-another-forwarding-header"
            },
            "name": "http_inspect"
        }
    ]
},
"telnet": {
    "enabled": true,
    "type": "multiton",
    "instances": [
        {
            "name": "telnet_my_instance",
            "data": {
                "encrypted_traffic": true
            }
        }
    ]
},
"binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
        {
            "when": {
                "role": "any",
                "service": "telnet"
            },
            "use": {
                "type": "telnet",
                "name": "telnet_my_instance"
            }
        },
        {
            "use": {
                "type": "http_inspect"
            },
            "when": {
                "role": "server",
                "service": "http",
                "dst_nets": "10.1.1.0/24"
            }
        }
    ]
}
}

```



Note You don't need to give the **name** attribute for the default instance in binder rules.

Configuring arp_spoof

Example for configuring **arp_spoof**:

The **arp_spoof** inspector does not have any default configurations for any attributes. This demonstrates the case where you can provide the overrides.

```
{
  "arp_spoof": {
    "type": "singleton",
    "data": {
      "hosts": [
        {
          "ip": "1.1.1.1",
          "mac": "ff:0f:f1:0f:0f:ff"
        },
        {
          "ip": "2.2.2.2",
          "mac": "ff:0f:f2:0f:0f:ff"
        }
      ]
    },
    "enabled": true
  }
}
```

Configuring rate_filter

```
{
  "rate_filter": {
    "data": [
      {
        "apply_to": "[10.1.2.100, 10.1.2.101]",
        "count": 5,
        "gid": 135,
        "new_action": "alert",
        "seconds": 1,
        "sid": 1,
        "timeout": 5,
        "track": "by_src"
      }
    ],
    "enabled": true,
    "type": "singleton"
  }
}
```

Configuring Binder Rules when Multi-Hierarchy Network Analysis Policy is Used

This example illustrates adding a new custom instance in child policy and the way binder rules should be written. Binder rules are defined as a list and therefore, it is important to pick up the rules defined in the parent policy and build the new rules on top of it as rules will not be merged automatically. The binder rules available in child policy are a source of truth in totality.

On the threat defense, the default Cisco Talos policy rules are appended on these user-defined overrides.

Parent Policy:

We have defined a custom instance by the name **telnet_parent_instance** and the corresponding binder rule.

```
{
  "telnet": {
```

```

    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": true
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}

```

Child Policy:

This network analysis policy has the aforementioned policy as its base policy. We have defined a custom instance by the name **telnet_child_instance** and have also defined the binder rules for this instance. The binder rules from parent policy need to be copied here, and then child policy binder rules can be prepended or appended on top of it based on the nature of the rule.

```

{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": false
        },
        "name": "telnet_child_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet",
          "nets": "10.2.2.0/24"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_child_instance"
        }
      }
    ]
  }
}

```

```

    }
  },
  {
    "when": {
      "role": "any",
      "service": "telnet"
    },
    "use": {
      "type": "telnet",
      "name": "telnet_parent_instance"
    }
  }
]
}
}

```

Configuring List Inspector Attribute in General

While changing overrides for any attribute of type list, it is important to pass the full contents rather than partial override. This means if a base policy attributes are defined as:

```

{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1"
      }
    },
    {
      "entry2": {
        "key2": "value2"
      }
    }
  ]
}

```

If you want to modify **value1** to **value1-new**, the override payload must look like the following:

Correct Way:

```

{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1-new"
      }
    },
    {
      "entry2": {
        "key2": "value2"
      }
    }
  ]
}

```

Incorrect Way:

```

{
  "list-attribute": [
    {
      "entry1": {
        "key1": "value1-new"
      }
    }
  ]
}

```

```
]
}
```

You can understand this configuration by taking the trimmed values of the `alt_max_command_line_len` attribute in the `smtp` inspector. Suppose the default (base) policy configuration for `smtp` inspector is as follows:

```
{
  "smtp": {
    "type": "multiton",
    "instances": [
      {
        "name": "smtp",
        "data": {
          "decompress_zip": false,
          "normalize_cmds": "ATRN AUTH BDAT CHUNKING DATA DEBUG EHLO
EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL
NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SOML STARTTLS TICK
TIME TURN TURNME VERB VRFY X-ADAT XADR XAUTH XCIR X-DRCP X-
ERCP XEXCH50 X-EXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE
XSTA XTRN XUSR",
          "ignore_data": false,
          "max_command_line_len": 512,
          "max_header_line_len": 1000,
          "log_rcptto": false,
          "decompress_swf": false,
          "max_response_line_len": 512,
          "b64_decode_depth": -1,
          "max_auth_command_line_len": 1000,
          "log_email_hdrs": false,
          "xlink2state": "alert",
          "binary_data_cmds": "BDAT XEXCH50",
          "auth_cmds": "AUTH XAUTH X-EXPS",
          "log_filename": false,
          "uu_decode_depth": -1,
          "ignore_tls_data": false,
          "data_cmds": "DATA",
          "bitenc_decode_depth": -1,
          "alt_max_command_line_len": [
            {
              "length": 255,
              "command": "ATRN"
            },
            {
              "command": "AUTH",
              "length": 246
            },
            {
              "length": 255,
              "command": "BDAT"
            },
            {
              "length": 246,
              "command": "DATA"
            }
          ],
          "log_mailfrom": false,
          "decompress_pdf": false,
          "normalize": "none",
          "email_hdrs_log_depth": 1464,
          "valid_cmds": "ATRN AUTH BDAT CHUNKING DATA DEBUG EHLO
EMAL ESAM ESND ESOM ETRN EVFY EXPN HELO HELP IDENT MAIL
NOOP ONEX QUEU QUIT RCPT RSET SAML SEND SOML STARTTLS TICK
TIME TURN TURNME VERB VRFY X-ADAT XADR XAUTH XCIR X-DRCP X-
```

```

        ERCP XEXCH50 X-EXCH50 X-EXPS XGEN XLICENSE X-LINK2STATE XQUE
        XSTA XTRN XUSR",
        "qp_decode_depth": -1
    }
}
],
"enabled": true
}
}

```

Now, if you want to add two more objects to the `alt_max_command_line_len` list:

```

{
    "length": 246,
    "command": "XEXCH50"
},
{
    "length": 246,
    "command": "X-EXPS"
}

```

Then the custom network analysis policy override JSON would look like the following:

```

{
    "smtp": {
        "type": "multiton",
        "instances": [
            {
                "name": "smtp",
                "data": {
                    "alt_max_command_line_len": [
                        {
                            "length": 255,
                            "command": "ATRNR"
                        },
                        {
                            "command": "AUTH",
                            "length": 246
                        },
                        {
                            "length": 255,
                            "command": "BDAT"
                        },
                        {
                            "length": 246,
                            "command": "DATA"
                        },
                        {
                            "length": 246,
                            "command": "XEXCH50"
                        },
                        {
                            "length": 246,
                            "command": "X-EXPS"
                        }
                    ]
                }
            }
        ]
    },
    "enabled": true
}

```

Configuring Overrides when Multi-Hierarchy Network Analysis Policy is used in Multiton Inspector

This example illustrates overriding attributes in child policy and how the merged configuration will be used in the child policy for any instance. Any overrides defined in the child policy will be merged with the parent policy. Thus, if attribute1 and attribute2 are overridden in parent policy and attribute2 and attribute3 are overridden in the child policy, the merged configurations are for child policy. This means that attribute1 (defined in parent policy), attribute2 (defined in child policy), and attribute3 (defined in child policy) will be configured on the device.

Parent Policy:

Here we have defined a custom instance by the name **telnet_parent_instance** and overridden 2 attributes namely, **normalize** and **encrypted_traffic** in the custom instance.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": false
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}
```

Child Policy:

This network analysis policy has the aforementioned policy as its base policy. We have overridden attribute **encrypted_traffic** from parent policy and also overridden new attribute **ayt_attack_thresh**.

```
{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "encrypted_traffic": true,
          "ayt_attack_thresh": 1
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  }
```

```

    }
  }

```

With the above policy JSON, when you deploy the network analysis policy the following merged JSON will be configured on the device.

```

{
  "telnet": {
    "type": "multiton",
    "instances": [
      {
        "data": {
          "normalize": true,
          "encrypted_traffic": true,
          "ayt_attack_thresh": 1
        },
        "name": "telnet_parent_instance"
      }
    ],
    "enabled": true
  },
  "binder": {
    "enabled": true,
    "type": "binder",
    "rules": [
      {
        "when": {
          "role": "any",
          "service": "telnet"
        },
        "use": {
          "type": "telnet",
          "name": "telnet_parent_instance"
        }
      }
    ]
  }
}

```

This example illustrates details for the custom network analysis policy. The same behavior is also exhibited in the default instance. Also, a similar merging would be done for Singleton inspectors.

Removing all the Inspector Overrides for the Network Analysis Policy:

Whenever you want to remove all the overrides for a specific network analysis policy, you can upload an empty JSON. While uploading the overrides, choose the option **Replace inspector overrides**.

```

{
}

```

Related Topics

[Snort 3 Definitions and Terminologies for Network Analysis Policy](#), on page 62

[Network Analysis Policy Mapping](#), on page 68

[Custom Network Analysis Policy Creation for Snort 3](#), on page 65

[Search for an Inspector on the Network Analysis Policy Page](#), on page 70

[Copy the Inspector Configuration](#), on page 70

[Customize the Network Analysis Policy](#), on page 71

[View the List of Inspectors with Overrides](#), on page 75

Network Analysis Policy Settings and Cached Changes

When you create a new network analysis policy, it has the same settings as its base policy.

When tailoring a network analysis policy, especially when disabling inspectors, keep in mind that some inspectors and intrusion rules require that traffic first be decoded or preprocessed in a certain way. If you disable a required inspector, the system automatically uses it with its current settings, although the inspector remains disabled in the network analysis policy web interface.



Note Because preprocessing and intrusion inspection are so closely related, the network analysis and intrusion policies examining a single packet **must** complement each other. Tailoring preprocessing, especially using multiple custom network analysis policies, is an **advanced** task.

The system caches one network analysis policy per user. While editing a network analysis policy, if you select any menu or other path to another page, your changes stay in the system cache even if you leave the page.



PART III

Encrypted Visibility Engine for Snort 3

- [Encrypted Visibility Engine, on page 93](#)



CHAPTER 7

Encrypted Visibility Engine

Encrypted Visibility Engine (EVE) is used to identify client applications and processes utilizing TLS encryption. It enables visibility and allows administrators to take actions and enforce policy within their environments. The EVE technology can also be used to identify and stop malware.

- [Overview of Encrypted Visibility Engine, on page 93](#)
- [How EVE Works, on page 94](#)
- [Configure EVE, on page 95](#)
- [Configure EVE Exception Rules, on page 96](#)

Overview of Encrypted Visibility Engine

The encrypted visibility engine (EVE) is used to provide more visibility into the encrypted sessions without the need to decrypt them. These insights into encrypted sessions are obtained by Cisco's open-source library that is packaged in Cisco's vulnerability database (VDB). The library fingerprints and analyzes incoming encrypted sessions and matches it against a set of known fingerprints. This database of known fingerprints is also available in the Cisco VDB.



Note The encrypted visibility engine feature is supported only on management center-managed devices running Snort 3. This feature is not supported on Snort 2 devices and device manager-managed devices.

Some of the important features of EVE are the following:

- You can take access control policy actions on the traffic using information derived from EVE.
- The VDB included in Cisco Secure Firewall has the ability to assign applications to some processes detected by EVE with a high confidence value. Alternatively, you can create custom application detectors to:
 - Map EVE-detected processes to new user-defined applications.
 - Override the built-in value of process confidence that is used to assign applications to EVE-detected processes.

See the **Configuring Custom Application Detectors** and **Specifying EVE Process Assignments** sections in the **Application Detection** chapter of the [Cisco Secure Firewall Management Center Device Configuration Guide](#).

- EVE can detect the operating system type and version of the client that created a Client Hello packet in the encrypted traffic.
- EVE supports fingerprinting and analysis of Quick UDP Internet Connections (QUIC) traffic too. The server name from the Client Hello packet is displayed in the URL field of the **Connection Events** page.

**Attention**

To use EVE on management center, you must have a valid Threat license on your device. In the absence of a Threat license, the policy displays a warning and deployment is not allowed.

**Note**

- EVE can detect the operating system type and version of SSL sessions. Normal usage of the operating system, such as running applications and package management software, can trigger OS detection. To view client OS detection, in addition to enabling the EVE toggle button, you must enable **Hosts** under **Policies > Network Discovery**. To view a list of possible operating systems on the host IP address, click **Analysis > Hosts > Network Map**, and then choose the required host.
- EVE will not provide visibility or insights for encapsulated traffic.

Related Links

[Configure EVE, on page 95](#)

How EVE Works

The Encrypted Visibility Engine (EVE) inspects the Client Hello portion of the TLS handshake to identify client processes. The Client Hello is the initial data packet that is sent to the server. This gives a good indication of the client process on the host. This fingerprint, combined with other data such as destination IP address, provides the basis for EVE's application identification. By identifying specific application fingerprints in the TLS session establishment, the system can identify the client process and take appropriate action (allow/block).

EVE can identify over 5,000 client processes. The system maps a number of these processes to client applications for use as criteria in access control rules. This gives the system the ability to identify and control these applications without enabling TLS decryption. By using fingerprints of known malicious processes, EVE technology can also be used to identify and block encrypted malicious traffic without outbound decryption.

Through machine learning (ML) technology, Cisco processes over one billion TLS fingerprints and over 10000 malware samples daily to create and update EVE fingerprints. These updates are then delivered to customers using Cisco Vulnerability Database (VDB) package.

If EVE does not recognize a fingerprint, it identifies client application and estimates the threat score of the first flow using the destination details, such as IP address, port, and server name. At this point, the status of the fingerprints are randomized and the status can be viewed in the debug logs. For subsequent flows with the same fingerprint, EVE skips reanalysis and marks the fingerprint status as unlabeled. If you intend to block traffic based on EVE's Low or Very Low score thresholds, the initial flow is blocked. However, future flows will be allowed once the application's fingerprint is cached.

Configure EVE

Procedure

-
- Step 1** Choose **Policies > Access Control**.
- Step 2** Click **Edit** (✎) next to the access control policy you want to edit.
- Step 3** Choose **Advanced Settings** from the **More** drop-down arrow at the end of the packet flow line.
- Step 4** Click **Edit** (✎) next to **Encrypted Visibility Engine**.
- Step 5** On the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.
- Step 6** Click **OK**.
- Step 7** Click **Save**.
-

What to do next

Deploy configuration changes.

View EVE Events

After enabling the **Encrypted Visibility Engine** and deploying your access control policy, you can start sending live traffic through your system. You can view the logged connection events on the **Connection Events** page or on the **Unified Events** page.

Do the following steps to access the connection events, in the management center.

Procedure

-
- Step 1** Click **Analysis > Connections Header > Events**.
- Step 2** Click the **Table View of Connection Events** tab.
- You can also view the connection event fields in the **Unified Events** page. Click **Analysis > Unified Events** to access the **Unified Events** page.
- Encrypted Visibility Engine can identify the client process that initiated a connection, the OS on the client, and if the process contains malware or not.
-

On the **Connection Events** page, you must explicitly enable the following columns that are added for Encrypted Visibility Engine.

- Encrypted Visibility Process Name
- Encrypted Visibility Process Confidence Score
- Encrypted Visibility Threat Confidence

- Encrypted Visibility Threat Confidence Score
- Detection Type

For information about these fields, see *Connection and Security Intelligence Event Fields* in [Cisco Secure Firewall Management Center Administration Guide](#).



Note On the **Connection Events** page, if processes are assigned applications, the **Detection Type** column displays **Encrypted Visibility Engine** indicating that the client application was identified by EVE. Without application assignments to process names, the **Detection Type** column displays **AppID** indicating that the engine that identified the client application was AppID.

View EVE Dashboard

You can view the EVE analysis information in the following dashboards:

Before you begin

- In an access control policy, the **Encrypted Visibility Engine (EVE)** must be enabled under the **Advanced Settings**.

Procedure

-
- Step 1** Go to **Overview > Dashboards**, and click **Dashboard**.
- Step 2** In the **Summary Dashboard** window, click the **switch dashboard** link and choose **Application Statistics** from the dropdown box.
- Step 3** Choose the **Encrypted Visibility Engine** tab to view the following two dashboards:
- **Top Encrypted Visibility Engine Discovered Processes**—Displays top client processes used in your network and the connection count. You can click the process name in the table to see the filtered view of the **Connection Events** page, which is filtered by the process name.
 - **Connections by Encrypted Visibility Engine Threat Confidence**—Displays connections by the confidence levels (Very High, Very Low, and so on). You can click the Threat confidence level in the table to see the filtered view of the **Connection Events** page, which is filtered by the confidence level.
-

Configure EVE Exception Rules

You can create an encrypted visibility engine (EVE) exception rule to ensure the continuity of trusted connections and services by bypassing the EVE's block action. You can add attributes such as process names and destination IP address to the exception rule. For example, you may want to bypass EVE's block verdict for trusted networks. All the connections in the bypassed networks are exempted from EVE's block verdict based on the threat confidence level.

Procedure

-
- Step 1** Choose **Policies > Access Control**.
- Step 2** Click **Edit** (✎) next to the access control policy you want to edit.
- Step 3** From the **More** drop-down arrow at the end of the packet flow line, choose **Advanced Settings**.
- Step 4** Next to **Encrypted Visibility Engine (EVE)**, click **Edit** (✎).
- Step 5** On the **Encrypted Visibility Engine** page, click the **Encrypted Visibility Engine (EVE)** toggle button to enable EVE.
- Step 6** Enable the **Block Traffic Based on EVE Score** toggle button to block traffic based on EVE's threat confidence level.
- Step 7** Click **Add Exception Rule** and add one or more of the following attributes.
- Under the **Process Name** tab, enter an EVE-identified process name, and click **Add to Process** on the right side of the window.

You can add multiple process names to the same exception rule. EVE exception list based on process names works only with EVE-identified process names, which are case- and space-sensitive.
 - Under the **Network Objects** tab, perform one of the following:
 - Choose one or more IP addresses from the list and add to the **Selected Networks** list.
 - Under **Selected Networks**, manually enter the IP address and click the + icon to add it to the list of selected networks.
 - (Optional) In the **Comment** field available on all the tabs, you can enter a reason for adding the required attributes to the EVE exception rule.
- Step 8** Click **Save** to save the EVE exception rule.
- Step 9** Save and deploy the access control policy on the devices.
-



Note When a connection matches an exception rule, it bypasses the EVE's block verdict. You can view EVE's action in the **Connection Events** or **Unified Events** page. The **Reason** column header displays **EVE Exempted** for identification of such EVE-bypassed traffic.


Add Exception Rule from Unified Events

You can use the **Unified Events** page to add exception rules for connections that are blocked by EVE.

Before you begin

Exception list is supported only from threat defense Version 7.6.0 or later.

Procedure

-
- Step 1** Click **Analysis > Unified Events**.
- Step 2** In the **Reason** column with **Encrypted Visibility Block** as the reason, click the **Ellipsis**  icon inside the cell.
- Step 3** Choose **Add EVE Exception Rule** from the drop-down list.
- Step 4** In the **Encrypted Visibility Engine** window that is displayed, the rule is automatically added to the bottom of the exception list. You can review and make changes to the added rule before saving and deploying the configuration.
-



PART **IV**

Elephant Flow Detection for Snort 3

- [Elephant Flow Detection](#), on page 101



CHAPTER 8

Elephant Flow Detection

Elephant flows are extremely large (in total bytes), continuous flows set up by a TCP (or other protocols) flow measured over a network link. By default, elephant flows are those larger than 1 GB/10 seconds. They can cause performance duress in Snort cores. Elephant flows are not numerous, but they can occupy a disproportionate share of the total bandwidth over a period of time. They can lead to problems, such as high CPU utilization, packet drops, and so on.

From management center 7.2.0 onwards (Snort 3 devices only), you can use the elephant flow feature to detect and remediate elephant flows, which helps to reduce system stress and resolve the mentioned issues.

- [About Elephant Flow Detection and Remediation, on page 101](#)
- [Elephant Flow Upgrade from Intelligent Application Bypass, on page 101](#)
- [Configure Elephant Flow, on page 102](#)

About Elephant Flow Detection and Remediation

You can use the elephant flow detection feature to detect and remediate elephant flows. The following remediation actions can be applied:

- **Bypass elephant flow**—You can configure elephant flow to bypass Snort inspection. If this is configured, Snort does not receive any packet from that flow.
- **Throttle elephant flow**—You can apply rate-limit to the flow and continue to inspect flows. The flow rate is calculated dynamically and 10% of the flow rate is reduced. Snort sends the verdict (QoS flow with 10% less flow rate) to the firewall engine. If you choose to bypass all applications including unidentified applications, you cannot configure the throttle action (rate-limit) for any flow.



Note For the elephant flow detection to work, Snort 3 must be the detection engine.

Elephant Flow Upgrade from Intelligent Application Bypass

Intelligent Application Bypass (IAB) is deprecated from version 7.2.0 onwards for Snort 3 devices.

For devices running 7.2.0 or later, you must configure elephant flow settings under the **Elephant Flow Settings** section in the AC policy (Advanced settings tab).

Post-upgrade to 7.2.0 (or later), if you are using a Snort 3 device, the elephant flow configuration settings will be picked and deployed from the **Elephant Flow Settings** section and not from the **Intelligent Application Bypass Settings** section, so if you have not migrated to Elephant Flow configuration settings, your device will lose the elephant flow configuration upon the next deployment.

The following table shows the IAB or elephant flow configurations that can be applied to version 7.2.0 or later and to version 7.1.0 or earlier that are running Snort 3 or Snort 2 engines.

Management Center	Threat Defense	Elephant Flow or IAB Configuration
Management Center 7.0 or 7.1	Snort 2 device	Configuration from IAB is applicable.
	Snort 3 device	Configuration from IAB is applicable.
Management Center 7.2.0	Snort 2 device	Configuration from IAB is applicable.
	Snort 3 device (7.1.0 and earlier)	Configuration from IAB is applicable.
	Snort 3 device (7.2.0 and later)	Configuration from Elephant Flow is applicable.

Configure Elephant Flow

You can configure elephant flow to take actions on elephant flows, which helps resolve issues, such as system duress, high CPU utilization, packet drops, and so on.



Attention

Elephant flow detection is not applicable for prefiltered, trusted, or fast-forwarded flows, which do not process through Snort. As elephant flows are detected by Snort, elephant flow detection is not applicable for encrypted traffic.

Procedure

- Step 1** If **View** (👁) appears instead, settings are inherited from an ancestor policy, or you do not have permission to modify the settings. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Figure 1: Configure Elephant Flow Detection**Figure 2: Configure Elephant Flow Detection**

Step 2 The **Elephant Flow Detection** toggle button is enabled by default. You can configure the values for flow bytes and flow duration. When they exceed your configured values, elephant flow events are generated.

Step 3 To remediate elephant flows, enable the **Elephant Flow Remediation** toggle button.

Step 4 To set the criteria for remediation of the elephant flow, configure the values for CPU utilization %, duration of fixed time windows, and packet drop %.

CPU utilization is calculated per elephant flow and is derived from the flow latency. If the CPU utilization crosses the configured threshold and other configurations, such as fixed time windows and packet drops, are also matched, the elephant flow remediation actions are applied. Similarly, packet drop calculation is based on the packets dropped per CPU. After the packet drop percentage exceeds the configured value on a specific CPU, the remediation actions are applied. For example, consider that configurations are set to the default, that is, CPU utilization of 40%, fixed time window of 30 seconds, and packet drop of 5%. On a specific CPU, if more than 5% of packet drops are detected and the CPU utilization per flow exceeds 40% in the fixed time frame of 30 seconds, then the flows are either bypassed or throttled.

Step 5 You can perform the following actions for elephant flow remediation when it meets the configured criteria:

- a. **Bypass the flow**—Enable this button to bypass Snort inspection for selected applications or filters. Choose from:
 - **All applications including unidentified applications**—Select this option to bypass all the application traffic. If you configure this option, you cannot configure the throttle action (rate-limit) for any flow.
 - **Select Applications/Filters**—Select this option to select the applications or filters whose traffic you want to bypass; see the topic **Configuring Application Conditions and Filters** in the **Access Control Rules** chapter in the [Cisco Secure Firewall Management Center Device Configuration Guide](#).
- b. **Throttle the flow**—Enable this button to apply rate-limit to the flow and continue to inspect flows. Note that you can select the applications or filters to bypass Snort inspection and throttle the remaining flows.

Note

Automatic removal of throttle from a throttled elephant flow occurs when the system is out of duress, that is, the percentage of Snort packet drops is lesser than your configured threshold. Consequently, rate limiting is also removed.

You can also manually remove throttling from a throttled elephant flow, using the following threat defense commands:

- **clear efd-throttle <5-tuple/all> bypass**—This command removes throttling from the throttled elephant flow and bypasses Snort inspection.
- **clear efd-throttle <5-tuple/all>**—This command removes throttling from the throttled elephant flow and Snort inspection continues. Elephant flow remediation is skipped after using this command.

For more information about these commands, see the [Cisco Secure Firewall Threat Defense Command Reference](#).

Note

Taking action on elephant flows (bypass and throttle the flow) is not supported on Cisco Firepower 2100 series devices.

Step 6 Click **OK** to save the elephant flow settings.

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

What to do next

Deploy configuration changes.

After configuring your elephant flow settings, monitor your connection events to see if any flows are detected, bypassed, or throttled. You can view this in the **Reason** field of your connection event. The three reasons for elephant flow connections are:

- Elephant Flow
- Elephant Flow Throttled
- Elephant Flow Trusted



Attention Enabling elephant flow detection alone does not cause generation of connection events for elephant flows. If a connection event is already logged for another reason and the flow is also an elephant flow, then the **Reason** field contains this information. However, to ensure that you are logging all elephant flows, you must enable connection logging in the applicable access control rules.

Refer to [Cisco Secure Firewall Elephant Flow Detection](#) for more information.