



Cisco FTD v7.0 on ASA 5500 and ISA 3000 and FTDv with FMC/FMCv

Common Criteria User Guide Supplement

IPS & VPN Functionality

Version 0.3
January 23, 2023

Prepared by:



Cisco Systems, Inc.,
170 West Tasman Drive, San Jose,
CA 95134-1706 USA

Network Traffic Control and IPS Functionality

Passive vs Inline Mode and Default Traffic Flows

FTD interfaces can be configured in either a passive or inline IPS deployment. In a passive deployment, an FTD interface is configured to receive traffic, but not forward that traffic (as an IDS), out-of-band from the flow of network traffic. In an inline deployment, network traffic flows across the FTD. One FTD can have multiple interface, where two interfaces are configured as inline pair, and a third interface is configured as passive.

Traffic policies are defined in terms of network “zones”, which in turn are associated with FTD interfaces. So, a policy may be defined to allow traffic from “zone0” to “zone1”, though those zones may be mapped to the “outside” and “inside” interfaces on one FTD and also be mapped to “int1” and “int2” interfaces of another FTD which enforces the same policy.

There are multiple types of policies that can be layered to apply to the same traffic flows (same zone-to-zone mappings). Having one type of policy applied to a zone/interface is sufficient to allow traffic flow. Traffic flow policy types include Prefilter, Access Control, and Intrusion policies.

Prefilter policies are sub-policies of Access Control policies, and every Access Control policy has an associated Prefilter policy, which is used to define rules for encapsulated traffic. There is no default action for nonencapsulated traffic; if a nonencapsulated connection does not match any prefilter rules, the system continues with applying rules in the Access Control policy. A Prefilter policy can contain multiple rules, which are enforced in the sequence they appear in the policy (the first rule that matches the traffic is the one that’s applied).

No FTD interface will forward traffic until policies have been configured and applied to that interface. Traffic will not be forwarded unless it’s explicitly permitted by at least one policy rule, thus an implicit “deny-all” rule is applied to all interfaces to which any traffic filtering rule has been applied. The implicit deny-all rule is executed after all admin-defined rules have been executed, and will result in dropping all traffic that has not been explicitly permitted, or explicitly denied. If an administrator wants to log all denied traffic, a rule entry should be added that denies all traffic and logs it, e.g. by either adding a rule at the end of a policy to explicitly drop and log all traffic, or by setting the Default Action for the policy to block all traffic, and enabling logging for the default rule, as show in this example:

The screenshot shows the Cisco FTD GUI for configuring a policy named 'my-new-policy'. The 'Rules' tab is selected, showing a table of rules. Below the table, the 'Default Action' is set to 'Access Control: Block All Traffic'. A 'Logging' dialog box is open, allowing configuration of logging options.

#	Name	Src Zone	Dest Zone	Src Net	Dest Net	VL	Use	App	Src	Dest	UR	ISE	Act	Att						
▼ Mandatory - my-new-policy (1-2)																				
1	icmp-any-any-log	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any
2	tcp-any-any-log	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any	Any
▼ Default - my-new-policy (-)																				
There are no rules in this section. Add Rule or Add Category																				
Default Action										Access Control: Block All Traffic										

Logging

- Log at Beginning of Connection
- Log at End of Connection

Send Connection Events to:

- Event Viewer
- Syslog: syslog-11
- SNMP Trap: Select an SNMP Alert Configuration...

Buttons: OK, Cancel



Passive Deployment

In a passive IPS deployment, the FTD device monitors traffic flowing across a network using a switch SPAN or mirror port. The SPAN or mirror port allows for traffic to be copied from other ports on the switch. This provides the system visibility within the network without being in the flow of network traffic. When configured in a passive deployment, the system cannot take certain actions such as blocking or shaping traffic. Passive interfaces receive all traffic unconditionally, and no traffic received on these interfaces is retransmitted.

You can configure one or more physical ports on a managed Device as passive interfaces.

IMPORTANT! When you disable a passive interface, users can no longer access it for security purposes.

1. Login with Administrator Role.

2. Select Device > Device Management.
3. Next to the Device where you want to configure the passive interface, click the edit icon ().
4. Next to the interface where you want to configure it as a passive interface, click the edit icon ().
5. Click **Passive**.
6. Associate a security zone with the passive interface
7. Check the **Enabled** check box.
8. Click **Save**.




Audit Record:

2016-11-23 17:50:36 admin Devices > Device Management > Device Edit > Interfaces  Save:82 10.128.120.41


Inline Deployment

In an inline IPS deployment, you configure the FTD device transparently on a network segment by binding two ports together. This allows the system to be installed in any network environment without the configuration of adjacent network Devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

You can configure one or more physical ports on a managed Device as inline interfaces. You must assign a pair of inline interfaces to an inline set before they can handle traffic in an inline deployment.

1. Login with Administrator Role.
2. Select Device > Device Management.
3. Next to the Device with the inline interfaces, click the edit icon ().
4. Next to the interface to be configured as an inline interface, click the edit icon ().
5. Click **Inline**.
6. Associate a security zone with the inline interface
7. Check the **Enabled** check box.
8. Click **Save**.
9. Select Device > Device Management.
10. Next to the Device where you want to add the inline set, click the edit icon ().
11. Click the **Inline Sets** tab.
12. Click Add Inline Set.

13. Enter a **Name**.

14. Next to **Interfaces**, choose one or more inline interface pairs, then click the icon ().

15. To specify that traffic is allowed to bypass detection and continue through the Device, choose Failopen (default). To use Failsafe, click on the **FailSafe** check box.

IMPORTANT! Failsafe option will prevent traffic from flowing through the appliance if a failure occurs for inline deployment. This can potentially cause a Denial of Service (DoS) attack on the monitored network.

16. Click **OK**.

Audit Record:

2016-11-23 18:07:40

admin

Devices > Device Management > Device Edit > Interfaces

 Save:83

10.128.120.41

Required Default Settings for the CC Evaluated Configuration

To configure traffic flow control as required for the CC-evaluated configuration follow guidance indicated in the list below, using the subsequent sections of this document for further details as needed to configure the Access Control Policies and Intrusion Policies.

- a) *packets which are invalid fragments:*
 - *To drop such traffic: Follow guidance in the Verify Enabled Preprocessors section of this document, then in the Intrusion Policy set the Rule State for all the IP defragmentation rules (GID:123) set to either Drop or "Drop and Generate Events".*
 - *To log such actions: Set the Rule State to "Drop and Generate Events".*
- b) *fragmented packets which cannot be re-assembled completely:*
 - *To drop such traffic: Follow guidance in the Verify Enabled Preprocessors section of this document, then in the Intrusion Policy set the Rule State for all the IP defragmentation rules (GID:123) set to either Drop or "Drop and Generate Events".*
 - *To log such actions: Set the Rule State to "Drop and Generate Events".*
- c) *packets where the source address of the network packet is defined as being on a broadcast network:*
 - *To drop such traffic: Ensure the Intrusion Policy has the Rule State for DECODE_IP4_SRC_BROADCAST (GID:116, SID:413) set to either Drop or "Drop and Generate Events".*
 - *To log such actions: Set the Rule State to "Drop and Generate Events".*
- d) *packets where the source address of the network packet is defined as being on a multicast network:*
 - *To drop such traffic: Ensure the Intrusion Policy has the Rule State for DECODE_IP4_SRC_MULTICAST (GID:116, SID:410) and DECODE_IPV6_SRC_MULTICAST (GID:116, SID:277) set to either Drop or "Drop and Generate Events".*
 - *To log such actions: Set the Rule State to "Drop and Generate Events".*
- e) *network packets where the source address of the network packet is defined as being a loopback address:*
 - *To drop such traffic: Ensure the Intrusion Policy has the Rule State for DECODE_BAD_TRAFFIC_LOOPBACK (GID:116, SID:150) set to either Drop or "Drop and Generate Events".*

- *To log such actions: Set the Rule State to “Drop and Generate Events”.*
- f) *network packets where the source or destination address of the network packet is defined as being unspecified (i.e. 0.0.0.0) or an address “reserved for future use” (i.e. 240.0.0.0/4) as specified in RFC 5735 for IPv4:*
 - *To drop such traffic: This is default behavior and cannot be disabled.*
 - *To log such actions: Enable logging in the Default Action of the Access Control Policy.*
- g) *network packets where the source or destination address of the network packet is defined as an “unspecified address” or an address “reserved for future definition and use” (i.e. unicast addresses not in this address range: 2000::/3) as specified in RFC 3513 for IPv6:*
 - *To drop such traffic: This is default behavior and cannot be disabled.*
 - *To log such actions: Enable logging in the Default Action of the Access Control Policy.*
- h) *network packets with the IP options: Loose Source Routing, Strict Source Routing, or Record Route specified:*
 - *To drop such traffic: Ensure the Intrusion Policy has the Rule State set to either Drop or “Drop and Generate Events” for POLICY-OTHER IP option loose source routing attempt (GID:1, SID:51034), POLICY-OTHER IP option strict source routing attempt (GID:1, SID:51035), and PROTOCOL-ICMP record route rr denial of service attempt (GID:1, SID:8730).*
 - *To log such actions: Set the Rule State to “Drop and Generate Events”.*
- i) *Other types of traffic:*
 - *These packets be dropped by default:*
 - i. *Slowpath Security Checks:*
 1. *In routed mode when receiving a through-the-box:*
 - a. *L2 broadcast packet (MAC address FF:FF:FF:FF:FF:FF)*
 - b. *IPv4 packet with destination IP address equal to 0.0.0.0*
 - c. *IPv4 packet with source IP address equal to 0.0.0.0*
 2. *In routed or transparent mode when receiving a through-the-box IPv4 packet with any of:*
 - a. *first octet of the source IP address equal to zero*
 - b. *network part of the source IP address equal to all 0's*
 - c. *network part of the source IP address equal to all 1's*
 - d. *source IP address host part equal to all 0's or all 1's*
 - ii. *ICMP Error Inspect and ICMPv6 Error Inspect (ICMP error packets when the ICMP error messages are not related to any session already established)*
 - iii. *ICMPv6 condition (when the appliance is not able to find any established connection related to the frame embedded in the ICMPv6 error message)*
 - iv. *ICMP Inspect bad icmp code (when an ICMP echo request/reply packet was received with a malformed code(non-zero))*
- *To log such actions: Enable logging in the Default Action of the Access Control Policy.*

The following traffic will be dropped by default, and auditing of those event can be enabled by enabling logging on the Default Action of the Access Control Policy.

- *Packets where the source address is equal to the address of the network interface where the network packet was received*
- *Packets where the source or destination address of the network packet is a link-local address*
- *Packets where the source address does not belong to the networks associated with the network interface where the network packet was received, including a description of how the TOE*

determines whether a source address belongs to a network associated with a given network interface

Configure Access Control Policy

An access control policy determines how the system handles traffic on the monitored network. Administrators can configure one or more access control policies, which they can then apply to one or more managed Devices. Each Device can have only one applied policy though. Access control rules can be added to a policy to provide granular control how traffic is handled and logged. To associate the access control policy and all rules under the policy to an interface, you first need to create the interface sets for the Device using “Configure Inline Interface” and “Configure Inline Set” sections from the general System User Guide. Then you can target the policy to a certain Device using the target tab.

For each rule, administrator can specify a rule *action*, that is, whether to trust, block, or inspect matching traffic with an intrusion policy. Each rule contains a set of conditions that identify the specific traffic you want to control. Rules can be simple or complex, matching traffic by any combination of security zone, IP address, application, protocols, ports, etc.

The system matches traffic to access control rules in order; the first matched rule handles the traffic.

Access Control Policy

On the Access Control Policy page ([Policies > Access Control](#)) administrator can view all the current access control policies by name and optional description and the following status information:








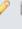












- When a policy is up to date on targeted Devices, in green text.
- When a policy is out of date on targeted Devices, in red text.

The default access control policy blocks all traffic from entering your network.

Creating Access Control Policy

When you create a new access control policy you must, at minimum, give it a unique name and specify a default action. Although you are not required to identify the policy targets at policy creation time, you must perform this step before you can apply the policy.

1. Login with Administrator Role or Access Admin.
2. Select Policies > Access Control.

Access Control Policy	Status	
Default Intrusion Prevention	Targeting 0 devices Up-to-date on all targeted devices	   
Default Network Discovery	Targeting 0 devices Up-to-date on all targeted devices	   
dp	Targeting 0 devices Up-to-date on all targeted devices	   
IDS Custom Policy	Targeting 0 devices Up-to-date on all targeted devices	   
Sarah Test	Targeting 0 devices Up-to-date on all targeted devices	   

3. Click New Policy.






New Access Control Policy ? x

Name:

Description:

Default Action: Block all traffic Intrusion Prevention Network Discovery

Targeted Devices

Available Devices	Selected Devices
<input type="text" value="Search"/>  birch  xiramat  tamarix  diana  phoebus	<input type="button" value="Add to Policy"/>

- In the **Name:** field, type a unique name for the new policy. Optionally, type a description in the **Description:** field.
- Specify the default action.

WARNING! Leave the default **Block all traffic** in the evaluated configuration.


- Select the Devices where you want to apply the policy. Click on the managed Device(s) you want the policy to applied to. Then click on **Add to Policy** button.
- Specify the initial **Default Action:**
 - Block all traffic creates a policy with the Access Control: Block All Traffic default action.
 - Intrusion Prevention** creates a policy with the **Intrusion Prevention: Balanced Security and Connectivity** default action, associated with the default intrusion variable set.
- Click **Save**.

- Click **Deploy** and select the Device(s) you want to deploy the setting to and click **Deploy** again.

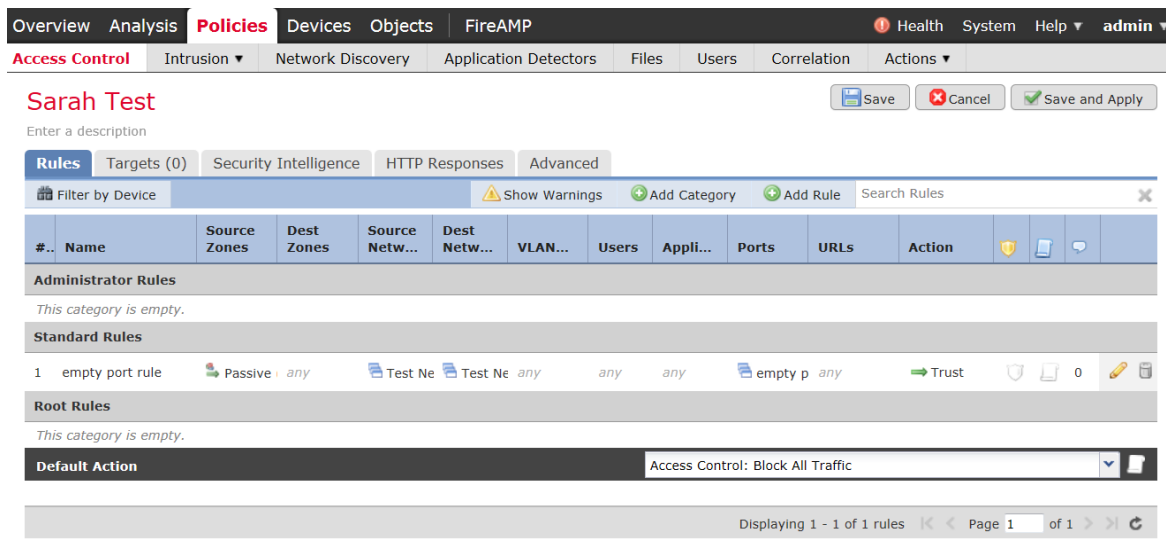
Audit Record:

2013-07-02 15:52:56 admin Policies > Access Control > Access Control Policy > cctest Create Policy 10.4.10.26

Editing Access Control Policy

- Login with Administrator Role.
- Select Policies > Access Control.
- Click the edit icon () next to the access control policy you want to configure.

The Policy Edit page appears.




The screenshot shows the 'Policy Edit' page for a policy named 'Sarah Test'. The page has a navigation bar with tabs for Overview, Analysis, Policies (selected), Devices, Objects, and FireAMP. Below the navigation bar, there are sub-tabs for Access Control, Intrusion, Network Discovery, Application Detectors, Files, Users, Correlation, and Actions. The main content area shows the policy name 'Sarah Test' with 'Save', 'Cancel', and 'Save and Apply' buttons. Below this, there are tabs for Rules, Targets (0), Security Intelligence, HTTP Responses, and Advanced. A 'Filter by Device' section is visible, followed by a table of rules. The table has columns for #, Name, Source Zones, Dest Zones, Source Netw..., Dest Netw..., VLAN..., Users, Appli..., Ports, URLs, Action, and a shield icon. The table shows one rule: 'empty port rule' with a 'Passive' status and 'Trust' action. At the bottom, there is a 'Default Action' dropdown set to 'Access Control: Block All Traffic' and a pagination bar showing 'Displaying 1 - 1 of 1 rules'.

- Make changes to the policy and click **Save**.
- Click **Deploy** and select the Device(s) you want to deploy the setting to and click **Deploy** again.

Audit Record:

2013-07-02 16:07:47 admin Policies > Access Control > Access Control Policy > cctest Save Policy 10.4.10.26
2017-02-24 17:30:41 System Task Queue Successful task completion : Policy Deployment to 172.18.152.193

Delete Access Control Policy

- Login with Administrator Role.
- Select Policies > Access Control.
- Click the delete icon () next to the policy you want to delete.
- Click **OK** to confirm.

Audit Record:

2013-07-02 16:10:29 admin Policies > Access Control > Access Control Policy > cctest Delete Policy 10.4.10.26



Access Control Rule

A set of access control rules is a key component of an access control policy. Access control rules allow administrator to manage, in a granular fashion, which traffic can enter the network, exit it, or cross from within without leaving it. Within an access control policy, the system matches traffic to rules in top-down order by rule number. In addition to its rule order and some other basic attributes, each rule has the following major components:

- A set of rule *conditions* that identifies the specific traffic you want to control.
- A rule *action*, which determines how the system handles traffic that meets the rule's conditions.
- Intrusion *inspection* option, which allow you to examine allowed traffic with intrusion policy.
- The *logging* option, which allow you to keep a record (event log) of the matching traffic.

The access control policy's default action defines the default action (for example, block all traffic) for the policy.

Creating and Editing Access Control Rules

1. Login with Administrator Role or Access Admin.
2. Select Policies > Access Control.
3. Click the edit icon () next to the access control policy you want to configure.
4. Add a new rule or edit an existing rule:
 - To add a new rule, click **Add Rule**.
 - To edit an existing rule, click the edit icon () next to the rule you want to edit.

Either the Add Rule or Editing Rule page appears.

Add Rule ? | X

Name Enabled Insert

Action Allow **IPS:** no policies **Files:** no inspection **Logging:** no logging

Zones Networks VLAN Tags Users Applications Ports URLs Inspection Logging Comments

Available Zones

- External
- External-1
- Internal
- Internal-1

Source Zones (0) **Destination Zones (0)**

any any

5. Configure the following rule components:


- You must provide a unique rule **Name**.
- Specify whether the rule is **Enabled**.
- Specify the rule position.
- Select a rule **Action**¹.
- Configure the rule's conditions².
- Configure the rule's **Inspection** option.
- Specify **Logging** option.
- Add **Comments**.

6. Click **Add** or **Save**.

Your changes are saved. You must apply the access control policy for your changes to take effect.

Audit Record:

2013-07-02 16:07:47	admin	Policies > Access Control > Access Control Policy > cctest	Save Policy	10.4.10.26
2013-07-02 16:27:29	admin	Policies > Access Control > Access Control Policy > test	Save Policy	10.4.10.26

Click on the compare () icon to see what rule(s) were added, removed, or modified and how.

For example, the following AC rule “cc rule” has been added to AC policy “test” by admin.

¹ The evaluated actions are Allow and Block.

² The evaluated conditions are Zones, Networks, Applications, and Ports. The other conditions are presented for completeness only.

test (2013-07-01 09:31:02 by ahepburn)		test (2013-07-02 16:12:57 by admin)	
Policy Information		Policy Information	
Modified	2013-07-01 09:31:02 by ahej	Modified	2013-07-02 16:12:57 by adrr
Rules		Security Intelligence	
Category 2		Blacklist Logging	
Name	Standard Rules	Send to Defense Center Disabled	
Rules		Rules	
Category 2		Category 2	
Name	Standard Rules	Name Standard Rules	
Rule 1		Rule 1	
Name	Standard Rules	Name cc rule	
Action	Standard Rules	Action allow	
Destination Ports		Destination Ports	
"TCP (6):123"		"TCP (6):123"	
Source Ports		Source Ports	
"TCP (6):123"		"TCP (6):123"	
Logging		Logging	
Log at beginning	Disabled	Log at beginning	Disabled
Log at end	Disabled	Log at end	Disabled
Log Files	Disabled	Log Files	Disabled
Send to Defense Center	Disabled	Send to Defense Center	Disabled

For example, the following AC rule “cc rule” has the new action set to block, from allow.

cctest (2013-07-02 15:54:53 by admin)		cctest (2013-07-02 16:07:47 by admin)	
Policy Information		Policy Information	
Modified	2013-07-02 15:54:53 by adrr	Modified	2013-07-02 16:07:47 by adrr
Rules		Rules	
Category 2		Category 2	
Name	Standard Rules	Name	Standard Rules
Rule 1		Rule 1	
Name	cc rule	Name	cc rule
Action	allow	Action	block



Understanding Rule Conditions

Administrator can set an access control rule to match traffic meeting any of the conditions described in the following table:

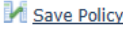
Zones	A configuration of one or more interfaces where you can apply policies. Zones provide a mechanism for classifying traffic on source and destination interfaces, and you can add source and destination zone conditions to rules.
Networks	Any combination of individual IPv4 and IPv6 addresses, CIDR blocks, and/or networks (by default, any). The system also supports Network Objects as described in Section 4, page 148 in the Cisco 3D System User Guide.
VLAN Tags	A number from 0 to 4094 that identifies traffic on your network by VLAN.
Applications	Applications provided by Cisco, user-defined applications, and application filters you create using the object manager.
Ports	Source and Destination ports. ICMPv4 and ICMPv6 type and code. Transport protocol ports, including individual and group port objects you create based on transport protocols ³ . The system supports Port Objects as described in Section 4, page 170 in the Cisco 3D System User Guide.
URLs	Cisco-provided URLs grouped by category and reputation, literal URLs, and any individual and group URL objects you create using the object manager.


To support the dynamic session establishment capability for FTP, you first need to create an access control rule that allows both FTP and FTP data. You can also configure the logging for this rule. This will enable the FTP application detector which has understanding of the application-level protocol so that FTP data connection will be allowed without additional rule.

Deleting Access Control Rules

1. Login with Administrator Role.
2. Select Policies > Access Control.
3. Click the edit icon () next to the access control policy you want to configure.
4. Click the delete icon () next to the access control rule you want to delete.
5. Click **OK** to confirm.
6. Click **Save**.

Audit Record:


2013-07-02 16:07:47	admin	Policies > Access Control > Access Control Policy > cctest	 Save Policy	10.4.10.26
2013-07-02 16:27:29	admin	Policies > Access Control > Access Control Policy > test	 Save Policy	10.4.10.26


Click on the compare () icon to see what rule was added, deleted, or modified and how. For example, the following AC rule “cc rule” has been deleted in AC policy “test” by admin.

³ We support all the protocol-specific attributes required in the FWPP.

test (2013-07-02 16:12:57 by admin)		test (2013-07-02 16:27:29 by admin)	
Policy Information		Policy Information	
Modified	2013-07-02 16:12:57 by adm	Modified	2013-07-02 16:27:29 by adm
Rules		Rules	
Category 2		Category 2	
Name	Standard Rules	Name	Standard Rules
Rule 1			
Name	cc rule		
Action	allow		
Destination Ports	"TCP (6):123"		
Source Ports	"TCP (6):123"		
Logging			
Log at beginning	Disabled		
Log at end	Disabled		
Log Files	Disabled		
Send to Defense Center	Disabled		

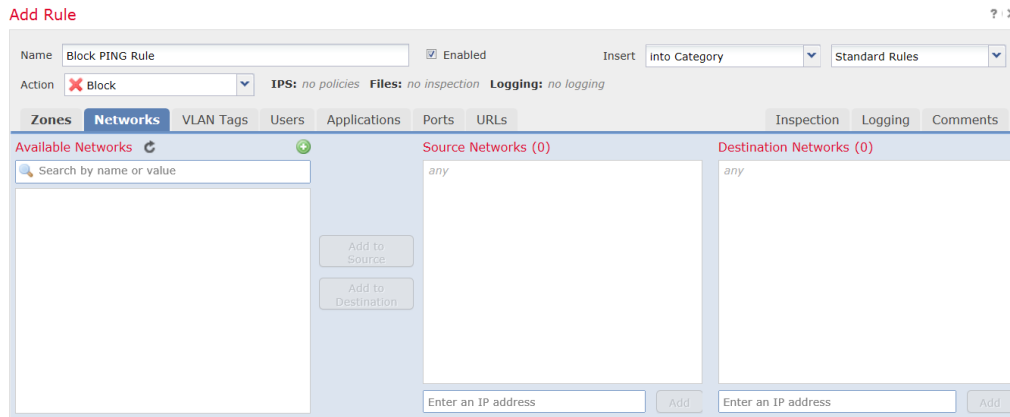
The following example demonstrates how to block all Ping (ICMP echo request) from the external network to internal network and log the connection attempt.

1. Login with Administrator Role.
2. Select Policies > Access Control.
3. Click the edit icon () next to the access control policy you want to configure.
4. Click **Add Rule**.
5. Type a name for the rule.
6. Leave the **Enabled** checkbox selected.
7. Let the rule get inserted into standard rules.
8. Select **Block** from drop-down list for the rule action.
9. On the **Zones** tab, select the **External** zone as the source zone and the **Internal** zone as the destination zone. You can click and drag or use the buttons.




10. On the **Networks** tab, select **any** as the source network and **any** as the destination network.

For granular control, you can enter IP address or range of IP addresses for source and destination networks. The system also supports IPv6 addresses as well.




11. On the **Ports** tab, in the second **Protocol** fields, select **ICMP(1)**.



The Select ICMP type and code pop-up window appears.

12. In the **Type:** field, select **8 (Echo Request)**.



13. Click **Add**.

14. On the Logging tab, check Log at Beginning of Connection.

15. In the Send Connection Events to: field, check the FMC.

16. Click **Add**.

#	Name	Sou... Zones	Dest Zones	Sou... Net...	Dest Net...	VLA...	U...	Ap...	Src...	Dest Ports	URLs	Action	Icons
Administrator Rules													
<i>This category is empty.</i>													
Standard Rules													
1	Block PING Rule	Exterr	Intern	any	any	any	any	any	any	ICMP (1):8	any	Block	Icons
Root Rules													
<i>This category is empty.</i>													
Default Action										Access Control: Block All Traffic			

17. Click **Save**.

The Intrusion and Network Analysis Policy (NAP) policies are associated with the Access Control (AC) policy which is then assigned to one or more sensors. However, only one AC policy can be assigned to any one sensor at a time (for example, if admin assigns AC policy 'XYZ' to a sensor with another policy assigned, the old AC policy will be unassigned automatically). Finally, when an AC policy is assigned to a sensor, that policy will be active on **all** the enabled interfaces on the sensor.

Modification of which Intrusion Policy is Active on Device's Interfaces

Create an IPS Policy and associate it with an AC Policy

1. Login with Administrator Role.
2. Select Policies > Access Control > Intrusion.
3. Click **Create Policy** and create the Intrusion policy.
4. Select Policies > Access Control.
5. Assign a Device (i.e., sensor) to the AC policy. Select the Device and click on **Add to Policy**.
6. Click **Save**.
7. Associate the Intrusion policy with the AC policy either through the default action or AC rule.
8. Click **Save**.

Audit Record:

2017-07-07 17:57:38	admin	Policies > Access Control > Access Control > Firewall Policy Editor	Save Policy AC policy 1:367
2017-07-07 17:56:28	admin	Policies > Access Control > Access Control > Firewall Policy Editor	Page View
2017-07-07 17:56:24	admin	Policies > Access Control > Access Control > Firewall Policy Editor	Create Policy AC policy 1;Assigned to device(s) : 172.18.152.1
2017-07-07 17:55:02	admin	Policies > Access Control > Access Control	Page View
2017-07-07 17:54:39	admin	Policies > Access Control > Intrusion	Page View
2017-07-07 17:54:38	admin	Policies > Intrusion > Intrusion Policy > IPS policy 1	Policy Committed - "Create initial policy"
2017-07-07 17:54:06	admin	Policies > Access Control > Intrusion	Page View

Click on the compare () icon to see what change.

AC policy 1 (2017-07-07 21:56:24/admin)	AC policy 1 (2017-07-07 21:57:37/admin)
Policy Information	Policy Information
Last Modified: 2017-07-07 21:56:24	Last Modified: 2017-07-07 21:57:37
Default Action	Default Action
Action: Access Control: Block All Traffic	Action: PERMIT
	Intrusion Policy: IPS policy 1
	Variable Set: Default-Set

Assign a Different AC Policy to the Device

1. Login with Administrator Role.
2. Select Policies > Access Control.
3. Edit a different AC policy.
4. Click on Policy Assignments.
5. Assign a Device to the AC policy. Select the Device and click on **Add to Policy**. Click **OK** and confirm.
6. Click **Save**.

Audit Record:

2017-07-07 18:08:00 admin Policies > Access Control > Access Control > Firewall Policy Editor  Save Policy New CC Policy:368


Click on the compare () icon to see what change.

New CC Policy (2017-03-03 00:50:45/Firepower System)	New CC Policy (2017-03-03 00:50:45/Firepower System)
Policy Information	Policy Information
Last Modified: 2017-03-03 00:50:45	Last Modified: 2017-07-07 22:07:59
	Applied To: 172.18.152.193

Associate AC Policy with Different Intrusion Policy

1. Login with Administrator Role.
2. Select Policies > Access Control.
3. Edit a AC policy.
4. Associate a different Intrusion policy either through the default action or AC rule.
5. Click **Save**.

Audit Record:

2017-07-07 18:12:07 admin Policies > Access Control > Access Control > Firewall Policy Editor  Save Policy AC policy 1:369

Click on the compare () icon to see what change.

AC policy 1 (2017-07-07 21:57:37/admin)	AC policy 1 (2017-07-07 22:12:07/admin)
Policy Information	Policy Information
Last Modified: 2017-07-07 21:57:37	Last Modified: 2017-07-07 22:12:07
Default Action	Default Action
Intrusion Policy: IPS policy 1	Intrusion Policy: SBD.1.3 Intrusion Policy

Enabling/Disabling a Device Interface with Intrusion Policy Applied

1. Login with Administrator Role.
2. Select Device > Device Management.
3. Edit an interface (e.g., eth1).

Edit Interface ? x

None Passive **Inline**

Security Zone: CC1

Inline Set: CC Inline Set


Enabled:

Save Cancel

4. To disable an interface, change the interface from **Inline** to **None**.
5. Click **Save**.

Audit Record:

2017-07-07 17:29:42 admin Devices > Device Management > Device Edit > Interfaces  Save:365 10.128.120.109

Click on the compare () icon to see what change.

SensorPolicy (2017-07-07 17:23:20 by admin from 10.128.120.10)	SensorPolicy (2017-07-07 17:29:42 by admin from 10.128.120.10)
SensorPolicy	SensorPolicy
firepower	firepower
Interfaces	Interfaces
eth1	eth1
Type: Inline	Type: None
Security Zone: CC1	Security Zone: CC1
Enabled: Yes	Enabled: No
Mode: Auto Negotiate	Mode: Auto Negotiate
MDI/MDIX: Auto	MDI/MDIX: Auto

Modification of which Mode(s) is/are Active on Device Interface

1. Login with Administrator Role.
2. Select Device > Device Management.
3. Edit an interface (e.g., eth1).

Edit Interface ? x

None Passive **Inline**

Security Zone: CC1

Inline Set: CC Inline Set

Enabled:

Save Cancel

- To change an interface mode, change the interface from **Inline** to **Passive**.
- Click **Save**.

Audit Record:

2017-07-07 17:14:57 admin Devices > Device Management > Device Edit > Interfaces Save:362 10.128.120.109

Click on the compare () icon to see what change.

SensorPolicy (2017-06-22 17:41:21 by admin from 127.0.0.1)	SensorPolicy (2017-07-07 17:14:57 by admin from 10.128.120.109)
SensorPolicy	
firepower	
Interfaces	
eth1	
Type Inline	Type Passive
Security Zone CC1	Security Zone CC1
eth2	
Type Inline	Type None
Security Zone CC2	Security Zone CC2
Enabled Yes	Enabled No
Mode Auto Negotiate	Mode Auto Negotiate
MDI/MDIX Auto	MDI/MDIX Auto
Inline Sets	
CC Inline Set	CC Inline Set
Interfaces eth1<->eth2	Interfaces eth1<->eth2

Note: eth1 and eth2 used to be inline and now eth1 is passive and eth2 is not active (i.e., disabled).

- Change eth1 and eth2 back to inline mode. Doing this also enables eth2.

Audit Record:

2017-07-07 17:23:20 admin Devices > Device Management > Device Edit > Interfaces Save:364 10.128.120.109

Click on the compare () icon to see what change.

SensorPolicy (2017-07-07 17:14:57 by admin from 10.128.120.10)		SensorPolicy (2017-07-07 17:23:08 by admin from 10.128.120.10)	
SensorPolicy		SensorPolicy	
firepower		firepower	
Interfaces		Interfaces	
eth1		eth1	
Type	Passive	Type	Inline
Security Zone		Security Zone	CC1
MTU	1518		
Load Balancing Mode	Use Inner IP Headers		
eth2		eth2	
Type	None	Type	Inline
Enabled	No	Enabled	Yes
		Mode	Auto Negotiate
		MDI/MDIX	Auto
		Inline Sets	
		CC Inline Set	
		Interfaces	
		eth1<->eth2	

Configure Security Intelligence

If you want to whitelist, blacklist, or monitor specific IP addresses, URLs, or domain names, you must configure custom objects, lists, or feeds. For your convenience, Cisco provides feeds containing IP addresses, domain names, and URLs with poor reputation, as determined by Talos:

- The *Intelligence Feed*, which comprises several regularly updated collections of IP addresses.
- The *DNS and URL Intelligence Feed*, which comprises several regularly updated collection of domain names and URLs.


You can also customize the feature to suit the unique needs of your organization, for example:

- **Global blacklist and custom blacklists**—the system allows you to manually blacklist specific IP addresses, URLs, or domain names in many ways depending on your needs.
- **Whitelisting to eliminate false positives**—when a blacklist is too broad in scope, or incorrectly blocks traffic that you want allow (for example, to vital resources), you can override a blacklist with a custom whitelist.
- **Monitoring instead of blacklisting**—especially useful in passive deployments and for testing feeds before you implemented them; you can merely monitor and log the violating sessions instead of blocking them.

By default, Security Intelligence filtering is not constrained by zone, that is, Security Intelligence objects have an associated zone of Any. You can constrain by only one zone. To enforce Security Intelligence filtering for an object on multiple zones, you must add the object to the whitelist or black list separately for each zone. Also, the default whitelist or blacklist cannot be constrained by zone.

1. Login with Administrator, Intrusion Admin or Access Admin role

NOTE: You must be 'admin', 'intrusion admin', or 'access admin' role to configure this.

2. Select Policies > Access Control.
3. Click the edit icon () next to the access control policy you want to configure.
4. Click on the Security Intelligence tab.
5. You have the following options:
 - Click the **Networks** tab to add network objects.
 - Click the **URLs** tab to add URL objects.
6. Find the **Available Objects** you want to add to the whitelist or blacklist.
7. Select one or more **Available Objects** to add.
8. Optionally, constrain the selected objects by zone by selecting an **Available Zone**.

NOTE: You cannot constrain system-provided Security Intelligence lists by zone.

9. Click **Add to Whitelist** or **Add to Blacklist**, or click and drag the selected objects to either list.
10. Optionally, set blacklisted objects to monitor-only by right-clicking the object under **Blacklist**, then selecting **Monitor-only (do not black)**.
11. Choose a DNS policy from the **DNS Policy** drop-down list.
12. Click **Save**.

The policy hierarchy order is not configurable and follows this order: Security Intelligence (whitelist takes precedence over blacklist), anomaly-based rules, then signature-based rules.

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

For intrusion rules to affect traffic, you must correctly configure drop rules and rules that replace content, as well as well as correctly deploy managed Devices inline, that is, with inline

interface sets. Finally, you must enable the intrusion policy's *drop behavior*, or **Drop when Inline** setting.

Create Intrusion Policy

When you create a new intrusion policy you must give it a unique name, specify a base policy, and specify drop behavior.

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click Create Policy.
4. Enter a unique **Name** and, optionally, a **Description**.
5. Specify the initial **Base Policy**.

You can use either a system-provided or another custom policy as your base policy.


6. Set the policy's drop behavior:
 - Check the **Drop when Inline** check box to allow intrusion rules to affect traffic and generate events.
 - Clear the **Drop when Inline** check box to prevent intrusion rules from affecting traffic while still generating events.
7. Create the policy:
 - Click **Create Policy** to create the new policy and return to the Intrusion Policy page. The new policy has the same settings as its base policy.
 - Click **Create and Edit Policy** to create the policy and open it for editing in the advanced intrusion policy editor.

Audit Record:

2016-11-22 18:07:08 admin Policies > Intrusion > Intrusion Policy > Test Policy Committed - "Create initial policy" 10.128.120.41

Viewing Intrusion Rules in an Intrusion Policy

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

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click the edit icon () next to the intrusion policy.
4. Click **Rules** under **Policy Information** in the navigation panel.
5. Check the rule whose rule details you want to view.
6. Click **Show details** button.

Intrusion Rule States

Intrusion rule states allow 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.

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

Generate Events

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 via the event logging.

Drop and Generate Events


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 via the event logging.

Note that rules set to this rule state generate events but do not drop packets in a passive deployment, including deployments where a Device inline interface set is in tap mode. For the system to drop packets, you must also enable the **Drop when Inline** in your intrusion policy and deploy your Device inline.

Disable

You do not want the system to evaluate matching traffic.

NOTE: Choosing either the **Generate Events** or **Drop and Generate Events** options enables the rule. Choosing **Disable** disables the rule.

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click the edit icon () next to the intrusion policy.
4. Click **Rules** under **Policy Information** in the navigation panel.
5. Choose the rule or rules where you want to set the rule state.
6. Choose one of the following:
 - Rule State > Generate Events
 - Rule State > Drop and Generate Events
 - Rule State > Disable
7. To save changes you made in this policy since the last policy commit, click **Policy Information** in the navigation panel, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

Audit Record:

2016-11-22 18:11:26 admin


Intrusion Policy > default > rule_configs

Added "Drop and generate events" to APP-DETECT 12P.DNS request attempt (1:3706

Adding and Modifying Intrusion Event Thresholds

You can set a threshold for one or more specific rules in an intrusion policy. You can also separately or simultaneously modify existing threshold settings. You can set a single threshold for each. Adding a threshold overwrites any existing threshold for the rule.

You can also modify the global threshold that applies by default to all rules and preprocessor-generated events associated with the intrusion policy. Please see the “Global Rule Threshold” section for more details.

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click the edit icon () next to the intrusion policy.
4. Click **Rules** under **Policy Information** in the navigation panel.
5. Choose the rule or rules where you want to set a threshold.
6. Choose Event Filtering > Threshold. To remove the threshold, choose Event Filtering > Remove Thresholds.
7. Choose a threshold type from the **Type** drop-down list.
8. From the **Track By** drop-down list, choose whether you want the event instances tracked by **Source** or **Destination** IP address.
9. Enter a value in the **Count** field.
10. Enter a value in the **Seconds** field.
11. Click **OK**.
12. To save changes you made in this policy since the last policy commit, click **Policy Information** in the navigation panel, then click **Commit Changes**.

If you leave the policy without committing changes, changes since the last commit are discarded if you edit a different policy.

Audit Record:

2016-11-22 18:14:53	admin	Intrusion Policy > Test > rule_configs > APP-DETECT 12P.DNS request attempt (1:37062) > threshold	Added "Threshold" to Type
2016-11-22 18:14:53	admin	Intrusion Policy > Test > rule_configs > APP-DETECT 12P.DNS request attempt (1:37062) > threshold	Added "Source" to Track By
2016-11-22 18:14:53	admin	Intrusion Policy > Test > rule_configs > APP-DETECT 12P.DNS request attempt (1:37062) > threshold	Added "12" to Count
2016-11-22 18:14:53	admin	Intrusion Policy > Test > rule_configs > APP-DETECT 12P.DNS request attempt (1:37062) > threshold	Added "60" to Seconds

Intrusion Rules Editor

An *intrusion rule* is a set of keywords and arguments that the system uses to detect attempts to exploit vulnerabilities on your network. As the system analyzes network traffic, it compares packets against the conditions specified in each rule. If the packet data matches all the conditions specified in a rule, the rule triggers. If a rule is an *alert rule*, it generates an intrusion event. If it is a *pass rule*, it ignores the traffic. For a *drop rule* in an inline deployment, the system drops the packet and generates an event. You can view and evaluate intrusion events from the Firepower Management Center web interface.

All rules contain two logical sections: the rule header and the rule options. The rule header contains:

- the rule's action or type
- the protocol
- the source and destination IP addresses and netmasks
- direction indicators showing the flow of traffic from source to destination
- the source and destination ports

The rule options section contains:

- event messages
- keywords and their parameters and arguments
- patterns that a packet's payload must match to trigger the rule
- specifications of which parts of the packet the rules engine should inspect

The following diagram illustrates the parts of a rule:

For example,

Rule Header


```
alert tcp $EXTERNAL_NET any -> $HTTP_SERVERS $HTTP_PORTS
```

Rule Keywords and Arguments

```
(msg:"WEB-IIS newdsn.exe access";
flow:to_server.established; uricontent:"/scripts/
tools/newdsn.exe"; nocase; metadata:service http;
reference:bugtraq,1818; reference:cve,1999-0191;
reference:nessus,10360; classtype:web-application-
activity; sid:1024; rev:10; )
```

Intrusion Rule Header

Every rule has a rule header containing parameters and arguments. The following illustrates parts of a rule header:



Action (*alert*) – Generates an intrusion event when triggered. Other actions include pass or drop.

Protocol (*tcp*) – Tests TCP traffic only. ICMP, IP, TCP, and UDP protocols are also supported.

Source IP (*\$EXTERNAL_NET*) – Tests traffic coming from any host that is not on your internal network.

Source Port (*any*) – Tests traffic coming from any port on the originating host.

Operate (*->*) – Tests external traffic destined for the web servers on your network.

Destination IP (*\$HTTP_SERVERS*) - Tests traffic to be delivered to any host specified as a web server on your internal network. Both IP and IPv6 addresses and ranges are supported.

Destination Port (*\$HTTP_PORTS*) - Tests traffic delivered to an HTTP port on your internal network.

Intrusion Rule Options and Keywords

Rule options follow the rule header and are enclosed inside a pair of parentheses. There may be one option or many and the options are separated with a semicolon. If you use multiple options, these options form a logical AND. The action in the rule header is invoked only when all criteria in the options are true. In general, an option may have two parts: a keyword and an argument.

The *message* keyword: Specify meaningful text that appears as a message when the rule triggers.

The *ack* keyword: Specify the acknowledgement value. For example, (flags: A; ack: 0; msg: "TCP ping detected"); means receive a TCP packet with the A flag set and the acknowledgement contains a value of 0.

The *content* keyword: Specify data pattern inside a packet. The pattern may be presented in the form of an ASCII string or as binary data in the form of hexadecimal characters.

The *offset* keyword: Specify a certain offset from the start of the data part of the packet to search.

The *dsize* keyword: Specify the length of the data part of a packet.

The *flags* keyword: Find out which flag bits are set inside the TCP header of a packet.

The *fragbits* keyword: Find out which three frag bits (Reserved, Don't Frag, More Frag) in the IP headers.

The *fragoffset* keyword: Tests the offset of a fragmented packet.

The *itype* keyword: Specify the ICMP type.

The *icode* keyword: Specify the ICMP code.

The *icmp_id* keyword: Specify the ICMP identification number.

The *icmp_seq* keyword: Specify the ICMP sequence number.

The *ipopts* keyword: Specify the IP Options. Record Route, Loose Source Routing, Strict Source Routing.

The *ip_proto* keyword: Specify the IP protocol number.

The *id* keyword: Specify the IP header fragment identification field

The *nocase* keyword: Its only purpose is to make a case insensitive search of a pattern within the data part of a packet. It is used in conjunction with the *content* keyword.

The *seq* keyword: Specify the sequence number of a TCP packet.

The *window* keyword: Specify the TCP window size.

The *flow* keyword: Apply a rule on TCP sessions to packets flowing in a particular direction.

The *tos* keyword: Detect a specific value in the Type of Service (TOS) field of the IP header.

The *ttl* keyword: Detect Time to Live value in the IP header of the packet.

Required Header Field Inspection	Intrusion Rule Keyword or Rule
IPv4:	
Version	alert (msg:"DECODE_NOT_IPV4_DGRAM"; sid:1; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
Header Length	alert (msg:"DECODE_IPV4_INVALID_HEADER_LEN"; sid:2; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
Packet Length	alert (msg:"DECODE_IPV4_DGRAM_LT_IPHDR"; sid:3; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;) alert (msg:"DECODE_IPV4_DGRAM_GT_CAPLEN"; sid:6; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
ID	id
IP Flags	fragbits
Fragment Offset	fragoffset
Time to Live (TTL)	ttl

Protocol	ip_proto
Header Checksum	Inspected by "Checksum Verification" preprocessor.
Source Address	Source IP OR alert (msg:"DECODE_IP4_SRC_MULTICAST"; sid:410; gid:116; rev:1; metadata:rule-type decode; classtype:misc-activity;)
Destination Address	Destination IP OR alert (msg:"DECODE_IP4_DST_RESERVED"; sid:412; gid:116; rev:1; metadata:rule-type decode; classtype:misc-activity;)
IP Options.	ipopts
IPv6:	
Version	alert (msg:"DECODE_IPV6_IS_NOT"; sid:271; gid:116; rev:1; metadata:rule-type decode; classtype: protocol-command-decode;)
payload length	dsize OR alert (msg:"DECODE_IPV6_TRUNCATED_EXT"; sid:272; gid:116; rev:1; metadata:rule-type decode; classtype:bad-unknown;)
next header	alert (msg:"DECODE_IPV6_BAD_NEXT_HEADER"; sid:281; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
hop limit	alert (msg:"DECODE_IPV6_MIN_TTL"; sid:270; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
source address	Source IP OR

	alert (msg:"DECODE_IPV6_SRC_MULTICAST"; sid:277; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
destination address	Destination IP OR alert (msg:"DECODE_IPV6_DST_RESERVED_MULTICAST"; sid:278; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;) alert (msg:"DECODE_IPV6_DST_ZERO"; sid:276; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
routing header	alert (msg:"DECODE_IPV6_ROUTE_AND_HOPBYHOP"; sid:282; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;) alert (msg:"DECODE_IPV6_TWO_ROUTE_HEADERS"; sid:283; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
ICMP:	
Type	itype
Code	icode
Header Checksum	Inspected by "Checksum Verification" preprocessor.
Rest of Header(varies based on the ICMP type and code)	icmp_id, icmp_seq
ICMPv6:	
Type	itype
Code	icode
Header Checksum	Inspected by "Checksum Verification" preprocessor.
TCP:	
source port	Source Port
destination port	Destination Port
sequence number	seq
acknowledgement number	ack
offset	alert (msg:"DECODE_TCP_INVALID_OFFSET"; sid:46; gid:116; rev:1; metadata:rule-type

	decode; reference:cve,2004-0816; classtype:bad-unknown;)
reserved	Inspected and normalized by preprocessor, if configured.
TCP flags	flags
window	window
checksum	Inspected by "Checksum Verification" preprocessor.
urgent pointer	alert (msg:"DECODE_TCP_BAD_URP"; sid:419; gid:116; rev:1; metadata:rule-type decode; classtype: misc-activity;) OR Inspected and normalized by preprocessor, if configured.
TCP options	alert (msg:"DECODE_TCPOPT_TRUNCATED"; sid:55; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
UDP:	
Source port	Source Port
destination port	Destination Port
length;	alert (msg:"DECODE_UDP_DGRAM_INVALID_LENGTH"; sid:96; gid:116; rev:1; metadata:rule-type decode; classtype:protocol-command-decode;)
UDP checksum	Inspected by "Checksum Verification" preprocessor.

Writing New Rules

1. Login with Administrator Role or Intrusion Admin.
2. Access the intrusion rules using either of the following methods:
 - Choose Policies > Access Control > Intrusion then click [Intrusion Rules](#).
 - Choose Objects > Intrusion Rules.
3. Click Create Rule.
4. Enter a value in the **Message** field.
5. Choose a value from each of the following drop-down lists:
 - Classification
 - Action
 - Protocol
 - Direction
6. Enter values in the following fields:

- Source IPs
- Destination IPs
- Source Port
- Destination Port

NOTE: The system uses the value 'any' if you do not specify a value for these fields.

7. Click Add Option.
8. Enter any arguments for the keyword you added.
9. Optionally, repeat steps 6 to 8.
10. If you added multiple keywords, you can:
 - Reorder keywords – Click the up or down arrow next to the keyword you want to move.
 - Delete a keyword – Click the **X** next to that keyword.
11. Click Save As New.

Audit Record:

2016-11-17 18:40:22 admin Policies > Intrusion > Rule Editor > Create save 1.1000000.1 10.128.120.41

Intrusion Rules Import

As new vulnerabilities become known, the Cisco Talos Security Intelligence and Research Group (Talos) releases intrusion rule updates that you can import onto your Firepower Management Center, and then implement by deploying the changed configuration to your managed Devices. These updates affect intrusion rules, preprocessor rules, and the policies that use the rules.

Intrusion rule updates are cumulative, and Cisco recommends you always import the latest update.

For changes made by an intrusion rule update to take effect, you must redeploy configurations. When importing a rule update, you can configure the system to automatically redeploy to affected Devices. This approach is especially useful if you allow the intrusion rule update to modify system-provided base intrusion policies.

1. Manually download the update from the Cisco Support Site (<http://www.cisco.com/cisco/web/support/index.html>).
2. Login with Administrator Role.
3. Choose **System > Updates**, then click the **Rule Updates** tab.
4. If you want to move all user-defined rules that you have created or imported to the deleted folder, you must click **Delete All Local Rules** in the toolbar, then click **OK**.
5. Choose **Rule Update or text rule file to upload and install** and click **Browse** to navigate to and choose the rule update file.

6. If you want to automatically re-deploy policies to your managed Devices after the update completes, choose **Reapply all policies after the rule update import completes**.
7. Click **Import**. The system installs the rule update and displays the Rule Update Log detailed view.


NOTE: Contact Support if you receive an error message while installing the rule update.

TCP Stream Reassembly

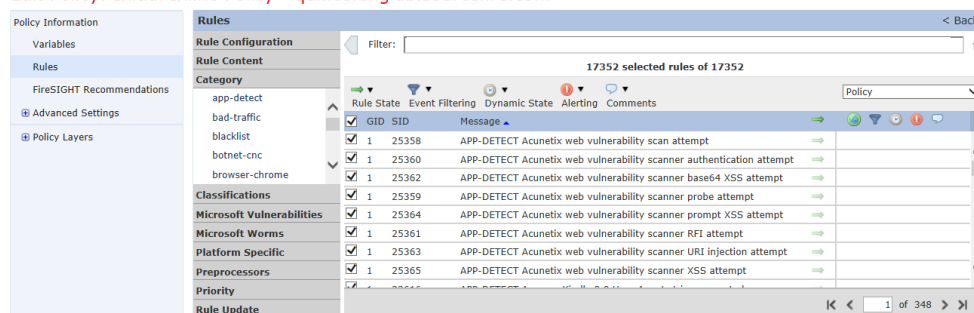
Please refer to the [TCP Stream Reassembly](#) section of Firepower Management Center Configuration Guide, Version 7.0

Configure Dynamic Rule State

The administrator can configure traffic bandwidth control at the policy level to stop excessive traffic from a specific source or network, to a specific destination or network, or all detected traffic.

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click the edit icon () next to the policy you want to configure.
4. Click **Rules** under **Policy Information** in the navigation panel.
5. Select the rule or rules where you want to add a dynamic rule state. You have the following options:
 - To select a specific rule, select the check box next to the rule.
 - To select all the rules, select the check box at the top of the column.

Edit Policy: Initial Inline Policy - quince.englab.sourcefire.com



6. Select Dynamic State > Add Rate-Based Rule State.

The Add Rate-Based Rule State dialog box appears.

Add Rate-Based Rule State for 17352 rules ? x

Track By	<input type="text" value="Destination"/>
Network	<input type="text"/>
Rate	<input type="text"/> Count / <input type="text"/> Seconds
New State	<input type="text" value="Drop and Generate Events"/>
Timeout	<input type="text"/>

7. Select the appropriate **Track By** option to indicate how you want the rule matches tracked:
 - Select **Source** to track the number of hits for that rule from a specific source or set of sources.

- Select **Destination** to track the number of hits for that rule to a specific destination or set of destinations.
 - Select **Rule** to track all matches for that rule.
8. When you set **Track By** to **Source** or **Destination**, enter the address of each host you want to track in the **Network** field.

You can specify a single IP address, address block, variable, or a comma-separated list comprised of any combination of these.

9. Indicate the number of rule matches per time period to set the attack rate:
- In the **Count** field, using an integer between 1 and 2147483647, specify the number of rule matches you want to use as your threshold.
 - In the **Seconds** field, using an integer between 1 and 2147483647, specify the number of seconds that make up the time period for which attacks are tracked.
10. Select a **New State** radio button to specify the action to be taken when the conditions are met:
- Select **Drop and Generate Events** to generate an event and drop the packet that triggered the event in inline deployments or generate an event in passive deployments.
11. In the **Timeout** field, type the number of seconds you want the action to remain in effect. After the timeout occurs, the rule reverts to its original state. Specify 0 or leave the field blank to prevent the action from timing out.
12. Click **OK**.
13. Select Commit Changes.
14. Deploy the policy.

Audit Record:

2013-03-19 18:03:50	admin	Intrusion Policy > CC Test > advanced_configs > rate_based_attacks > control_connections > connect > Control Simultaneous Connections 1	Added "Source" to Track By
2013-03-19 18:03:50	admin	Intrusion Policy > CC Test > advanced_configs > rate_based_attacks > control_connections > connect > Control Simultaneous Connections 1	Added "10.1.1.1" to Network
2013-03-19 18:03:50	admin	Intrusion Policy > CC Test > advanced_configs > rate_based_attacks > control_connections > connect > Control Simultaneous Connections 1	Added "1000" to Count
2013-03-19 18:03:50	admin	Intrusion Policy > CC Test > advanced_configs > rate_based_attacks > control_connections > connect > Control Simultaneous Connections 1	Added "No" to Drop
2013-03-19 18:03:50	admin	Intrusion Policy > CC Test > advanced_configs > rate_based_attacks > control_connections > connect > Control Simultaneous Connections 1	Added "1" to Timeout

Global Rule Threshold



The global rule threshold sets limits for event logging by an intrusion policy. You can set a global rule threshold across all traffic to limit how often the policy logs events from a specific source or destination and displays those events per specified time period. You can also set thresholds per rule, or preprocessor rule in the policy. When you set a global threshold, that threshold applies for each rule in the policy that does not have an overriding specific threshold. Thresholds can prevent you from being overwhelmed with a large number of events.

Every intrusion policy contains a default global rule threshold that applies by default to all intrusion rules and preprocessor rules. This default threshold limits the number of events on traffic going to a destination to one event per 60 seconds.

You can:

- Change the global threshold.
- Disable the global threshold.
- Override the global threshold by setting individual thresholds for specific rules.

For example, you might set a global limit threshold of five events every 60 seconds, but then set a specific threshold of ten events for every 60 seconds for SID1315. All other rules generate no more than five events in each 60-second period, but the system generates up to ten events for each 60-second period for SID1315.

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click the edit icon () next to the policy you want to configure.
4. Click **Advanced Setting** in the navigation panel.
5. If Global Rule Thresholding under Intrusion Rule Thresholds is disabled, click Enabled.
6. Click the edit icon () next to **Global Rule Thresholding**.
7. Using the **Type** radio buttons, specify the type of threshold that will apply over the time you specify in the **Seconds** field.

- 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.

- 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 2, and **Seconds** to 10, the following event counts result:

- 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 triggered the second time and following events are ignored).
8. Using the **Track By** radio buttons, specify the tracking method. This determines whether the event in stance count is calculated per source or destination IP address.
 9. Enter a value in the **Count** field.
 10. Enter a value in the **Seconds** field.
 11. To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

Audit Record:

2016-11-17 19:48:46	admin	Intrusion Policy > default > advanced_configs > threshold_global	Changed Count to "2" (from "1")	10.128.120.41
2016-11-17 19:48:46	admin	Intrusion Policy > default > advanced_configs > threshold_global	Changed Seconds to "65" (from "60")	10.128.120.41

Stateful Session Behaviors

The system implements packet decoders and preprocessors to detect anomalous traffic that might signal an intrusion attempt and, when the appropriate enabled accompanying decoder and preprocessor rules, report on detected anomalies. Next, intrusion rules examine the decoded packets for attacks based on patterns. Used together, intrusion rules and preprocessors provide broader and deeper packet inspection than a signature-based system and help to identify intrusions more effectively.

Before packets can be inspected, the packets must be captured from the network. As the system captures packets, it sends them to the packet decoder. The packet decoder converts the packet headers and payloads into a format that can be easily used by the preprocessors and the rules engine. 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, as described in the following table.

TCP/IP Layer	Decoded Packets
Data Link	Ethernet
	Virtual local area network (VLAN)
Network	Internet Protocol version 4 (IPv4)

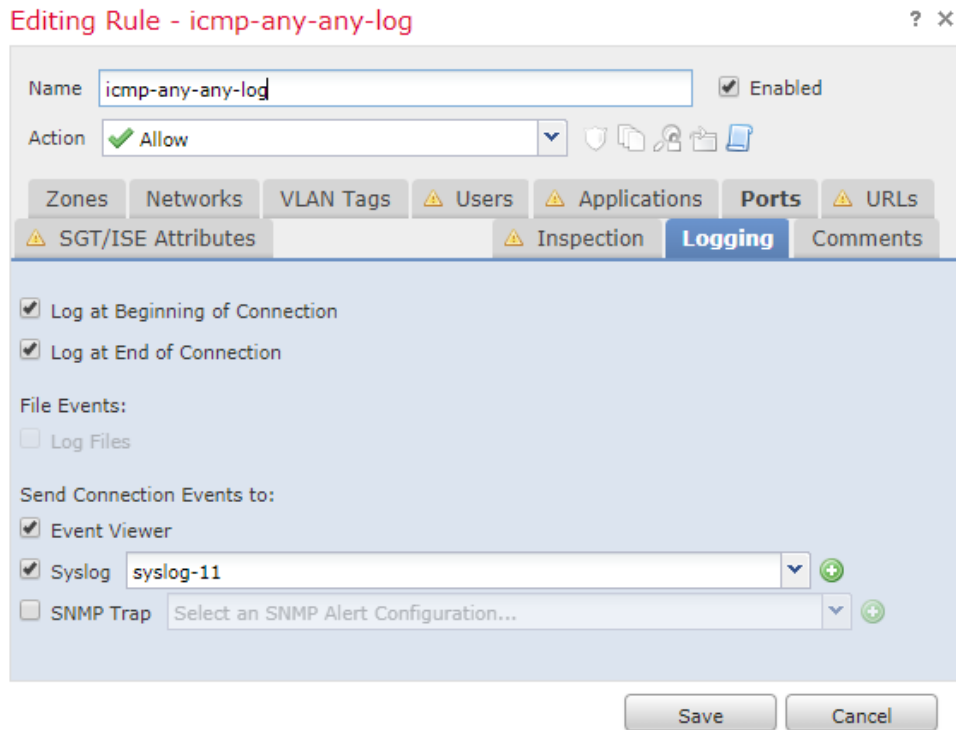
	Internet Protocol version 6 (IPv6)
	Internet Control Message Protocol version 4 (ICMPv4)
	Internet Control Message Protocol version 6 (ICMPv6)
Transport	Transmission Control Protocol (TCP)
	User Datagram Protocol (UDP)

After the packets are decoded through the first three TCP/IP layers, they are sent to preprocessors, which normalize traffic at the application layer and detect protocol anomalies. The following three preprocessors must be enabled and configured in the evaluated configuration (by default, all three preprocessors are enabled):

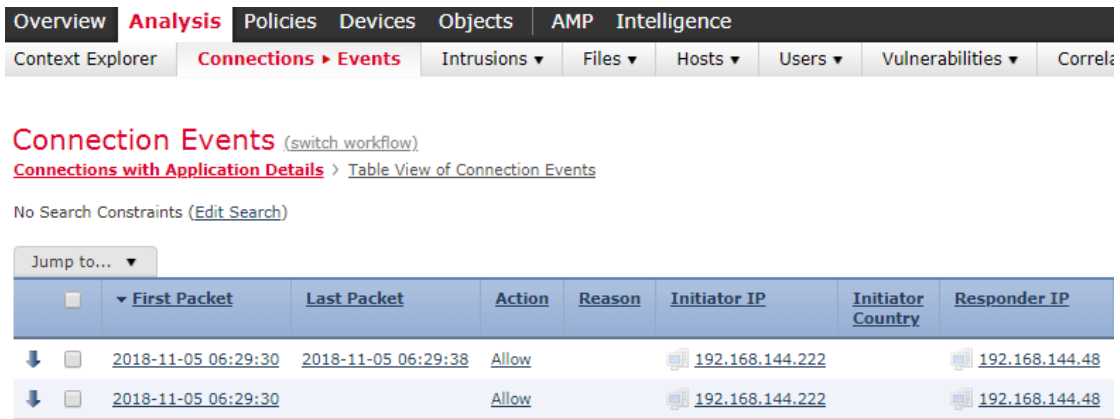
- TCP Streaming Preprocessor - Administrators can configure the system so that the preprocessor detects any TCP traffic that cannot be identified as part of an established TCP session. Stateful inspection allows administrators to ignore these packets because they are not part of an established TCP session and do not provide meaningful information.
- UDP Streaming Preprocessor - UDP data streams are not typically thought of in terms of sessions. However, the stream preprocessor uses the source and destination IP address fields in the encapsulating IP datagram header and the port fields in the UDP header to determine the direction of flow and identify a session.
- IP Defragmentation Preprocessor - When an IP datagram is broken into two or more smaller IP datagrams because it is larger than the maximum transmission unit (MTU), it is fragmented. A single IP datagram fragment may not contain enough information to identify a hidden attack. Attackers may attempt to evade detection by transmitting attack data in fragmented packets or attempt to crash the system when reassembling the fragmented packets. The IP defragmentation preprocessor reassembles fragmented IP datagrams, and if fragmented datagrams cannot be reassembled, it will be rejected (i.e., dropped) and logged with certain intrusion rules enabled.

To view audit records related to these traffic events:

- 1) Configure Access Control Policy rules to send log messages to Send Connection Events to “Event Viewer”, as in this examples:



2) View the events via FMC by navigating to Analysis > Connections > Events.



3) To filter events, click “Edit Search.”

4) To sort events, click any column heading, e.g. “Initiator IP”.

Verify Enabled Preprocessors

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click Create Policy.

Create Intrusion Policy

4. In the **Name** field, enter a unique name and optionally a description.
5. Click **Create and Edit Policy**.

Edit Policy: CC Test

6. Click **Advanced Settings**.
7. Verify that IP Defragmentation, TCP Stream and UDP Stream are enabled, all of which are enabled by default regardless of which base policy template is used when creating the custom policy.

Edit Policy: CC Test

8. Click on Policy Information and Commit Changes.
9. Optionally, enter a comment and click **OK**.
10. Associate the intrusion policy with the access control policy (via the **Advanced** tab under Policies > Access Control > [Edit or Add]).

NOTE! You cannot apply the intrusion policy until it is associated with an access control policy or rule.

Audit Record:

2013-03-19 19:26:14	admin	Policies > Intrusion > Intrusion Policy > CC Test	Policy Committed: "Create initial policy"	10.4.11.248
2013-03-19 19:28:10	admin	Intrusion Policy > CC Test > advanced_configs > normalize	Added "Enabled" to Normalize TCP	10.4.11.248
2013-03-19 19:28:10	admin	Intrusion Policy > CC Test > advanced_configs > normalize	Added "Enabled" to Normalize TCP Payload	10.4.11.248

Configure Anomaly Detection

Preprocessors prepare traffic to be further inspected by normalizing traffic and identifying protocol anomalies. Preprocessors can generate preprocessor events when packets trigger preprocessor options that you configure. The base policy for your network analysis policy determines which preprocessors are enabled by default and the default configuration for each.

The FTP/Telnet decoder analyzes FTP and telnet data streams, normalizing FTP and telnet commands before processing by the rules engine. You can enable rule126:3 to generate an event when this anomaly is detected in Telnet traffic, and rule125:9 when it is detected on the FTP command channel.

The inline normalization preprocessor normalizes traffic to minimize the chances of attackers evading detection in inline deployments. You can specify normalization of any combination of IPv4, IPv6, ICMPv4, ICMPv6, and TCP traffic. When the packet decoding **Detect Protocol Header Anomalies** option is enabled, you can enable the following rules in the decoder rule category to generate events for this option:

- You can enable rule 116:428 to generate an event when the system detects an IPv4 packet with a TTL less than the specified minimum.
- You can enable rule 116:270 to generate an event when the system detects an IPv6 packet with a hop limit that is less than the specified minimum.

The system can detect, drop, and log anomaly fragmented packets if the IP Defragmentation Preprocessor is enabled and certain intrusion rules are enabled.

1. Login with Administrator Role or Intrusion Admin.
2. Select Policies > Access Control > Intrusion.
3. Click Create Policy.

Create Intrusion Policy ? x

Policy Information

Name *

Description

Drop when Inline

Base Policy

Variables


Use the system default value

Networks to protect


* Required

4. In the **Name** field, enter a unique name and optionally a description.
5. Click Create and Edit Policy.

Edit Policy: CC Test



6. Click Manage Rules.



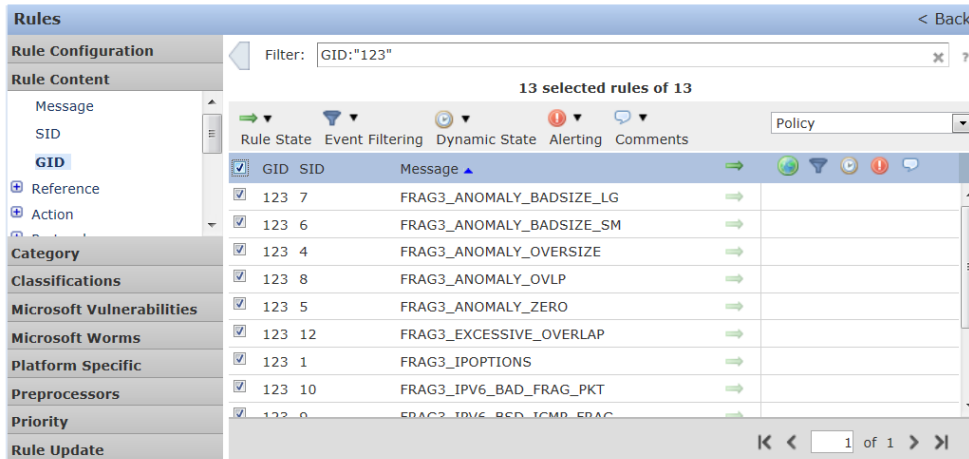
7. Click **Rule Content** and select **GID**.

The Enter the GID filter pop-up window appears.

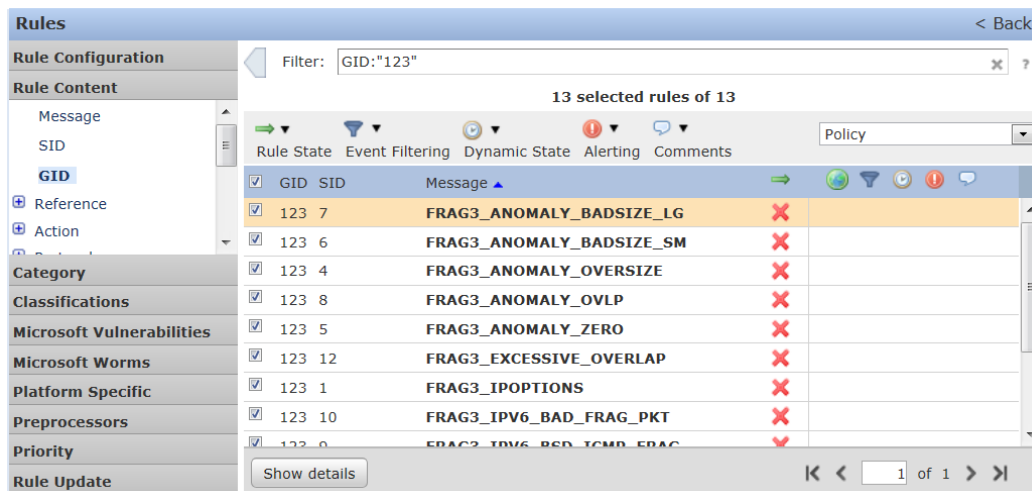


8. Enter **123** and click **OK**.


9. Select all the rules. Hint: Click the top checkbox.



10. In the **Rule State** field, click and select **Drop and Generate Events**.



11. For more details on each rule, click on a rule and select **Show details**.



11. Click on Policy Information and Commit Changes.
12. Optionally, enter a comment and click **OK**.
13. Associate the intrusion policy with the access control policy.

NOTE! You cannot apply the intrusion policy until it is associated with an access control policy or rule.

Audit Record:

2013-03-01 13:36:21	admin	Intrusion Policy > test2 > advanced_configs > normalize	Added "Enabled" to Normalize TCP	10.4.10.223
2013-03-01 13:36:21	admin	Intrusion Policy > test2 > advanced_configs > normalize	Added "Enabled" to Normalize TCP Payload	10.4.10.223
2013-03-01 13:36:20	admin	Policies > Intrusion > Intrusion Policy	Page View	10.4.10.223
2013-03-01 13:36:19	admin	Policies > Intrusion > Intrusion Policy > test2	Policy Committed: ""	10.4.10.223

Portscan Detection

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

By itself, a portscan is not evidence of an attack. In fact, some of the port scanning techniques used by attackers can also be employed by legitimate users on your network. Cisco's portscan detector is designed to help you determine which portscans might be malicious by detecting patterns of activity.

Protocol Types



TCP	Detects TCP probes such as SYN scans, ACK scans, TCP connect() scans, and scans with unusual flag combinations such as Xmas tree, FIN, and NULL.
UDP	Detects UDP probes such as zero-byte UDP packet.
ICMP	Detects ICMP echo requests (pings).
IP	Detects IP protocol scans. These scans differ from TCP and UDP scans because the attacker, instead of looking for open ports, is trying to discover which IP protocols are supported on a target host.

When portscan detection is enabled, you must enable rules with GeneratorID (GID)122 and a SnortID (SID) from among SIDs 1 through 27 to generate events for each enabled portscan type.

Portscan Event Packet View

When you enable the accompanying preprocessor rules, the portscan detector generates intrusion events that you can view just as you would any other intrusion event. However, the information presented on the packet view is different from the other types of intrusion events.



Begin by using the intrusion event views to drill down to the packet view for a ports can event. Note that you cannot download a portscan packet because single port scan events are based on multiple packets; however, the portscan packet view provides all usable packet information.

For any IP address, you can click the address to view the context menu and select [whois](#) to perform a lookup on the IP address or [View Host Profile](#) to view the host profile for that host.

Portscan Packet View

Device	The Device that detected the event.
Time	The time when the event occurred.
Message	The event message generated by the preprocessor.
Source IP	The IP address of the scanning host.
Destination IP	The IP address of the scanned host.
Port/Proto Count	For TCP and UDP portscans, the number of times that the port being scanned changes. For example, if the first port scanned is 80, the second port scanned is 8080, and the third port scanned is again 80, then the port count is 3. For IP protocol portscans, the number of times that the protocol being used to connect to the scanned host changes.
Port/Proto Range	For TCP and UDP portscans, the range of the ports that were scanned. For IP protocol portscans, the range of IP protocol numbers that were used to attempt to connect to the scanned host.

Open Ports	The TCP ports that were open on the scanned host. This field appears only when the portscan detects one or more open ports.
------------	---

- 1 Login with Administrator Role or Intrusion Admin.
- 2 Select Policies > Access Control > Intrusion then click on [Network Analysis Policy](#).
- 3 Click the edit icon () next to the policy you want to edit.
- 4 Click [Settings](#).
- 5 If Portscan Detection under Specific Threat Detection is disabled, click Enabled.
- 6 Click the edit icon () next to [Portscan Detection](#).
- 7 In the [Protocol](#) field, specify protocols to enable.

NOTE! You must ensure TCP stream processing is enabled to detect scans over TCP, and that UDP stream processing is enabled to detect scans over UDP. Also make sure you do not enable “Packet Size Performance Boost” and “Packet Type Performance Boost”.

- 8 In the [Scan Type](#) field, specify portscan types you want to detect.
- 9 Choose a level from the [Sensitivity Level](#) list.

NOTE! If you are encountering inconsistent detection (especially on the virtual Sensor), try disabling the “Latency-based performance setting”.

- 10 To save changes you made in this policy since the last policy commit, click [Policy Information](#), then click [Commit Changes](#).

Audit Record:

2016-11-29 13:20:15	admin	Network Analysis Policy > default > settings	Changed Portscan Detection to "Enabled" (from "Disabled")	10.128.120.136
2016-11-29 13:20:15	admin	Network Analysis Policy > default > settings > portscan	Changed Protocol to "TCP" (from "TCP, UDP")	10.128.120.136

Rate-Based Attack Prevention

Rate-based attacks (i.e., flooding attacks) are attacks that depend on frequency of connection or repeated attempts to perpetrate the attack. You can use rate-based detection criteria to detect a rate-based attack as it occurs and respond to it when it happens, then return to normal detection settings after it stops.

You can configure your network analysis policy to include rate-based filters that detect excessive activity directed at hosts on your network. You can use this feature on managed Devices deployed in inline mode to block rate-based attacks for a specified time, then revert to only generating events and not drop traffic.

The SYN attack prevention option helps you protect your network hosts against SYN floods. You can protect individual hosts or whole networks based on the number of packets seen over a period of time. If your Device is deployed passively, you can generate events. If your Device is placed inline, you can also drop the malicious packets. After the timeout period elapses, if the rate condition has stopped, the event generation and packet dropping stops.

For example, you could configure a setting to allow a maximum of 10 SYN packets from anyone IP address, and block further connections from that IP address for 60 seconds.

You can also limit TCP/IP connections to or from hosts on your network to prevent denial of service (DoS) attacks or excessive activity by users. When the system detects the configured number of successful connections to or from a specified IP address or range of addresses, it generates events on additional connections. The rate-based event generation continues until the timeout period elapses without the rate condition occurring. In an inline deployment you can choose to drop packets until the rate condition times out.

For example, you could configure a setting to allow a maximum of 10 successful simultaneous connections from anyone IP address, and block further connections from that IP address for 60 seconds.


Rate-based attack prevention identifies abnormal traffic patterns and attempts to minimize the impact of that traffic on legitimate requests. Rate-based attacks usually have one of the following characteristics:


- Any traffic containing excessive incomplete connections to hosts on the network, indicating a SYN flood attack
- Any traffic containing excessive complete connections to hosts on the network, indicating a TCP/IP connection flood attack
- Excessive rule matches in traffic going to a particular destination IP address or addresses or coming from a particular source IP address or addresses.
- Excessive matches for a particular rule across all traffic.

In a network analysis policy, you can either configure SYN flood or TCP/IP connection flood detection for the entire policy; in an intrusion policy, you can set rate-based filters for individual intrusion or preprocessor rules. Note that you cannot manually add a rate-based filter to GID135 rules or modify their rule state. Rules with GID135 use the client as the source value and the server as the destination value.

The *detection_filter* keyword prevents a rule from triggering until a threshold number of rule matches occur within a specified time. When a rule includes the *detection_filter* keyword, the system tracks the number of incoming packets matching the pattern in the rule per timeout period. The system can count hits for that rule from particular source or destination IP addresses. After the rate exceeds the rate in the rule, event notification for that rule begins.

Rate-based attack prevention can be configured at the policy level to stop SYN flood attacks. Excessive connections can also be blocked from a specific source or to a specific destination. These controls are configured through Network Analysis Policies (NAP) as described here:

- 1 Login with Administrator Role or Intrusion Admin.
- 2 Select Policies > Access Control > Intrusion then click on [Network Analysis Policy](#).
- 3 Click the edit icon () next to the policy you want to edit.
- 4 Click [Settings](#).

- 5 If Rate-Based Attack Prevention under Specific Threat Detection is disabled, click Enabled.
- 6 Click the edit icon () next to **Rate-Based Attack Prevention**.
- 7 You have two choices:
 - To prevent incomplete (half-open, ‘embryonic’) connections intended to flood a host, click **Add** under **SYN Attack Prevention**.
 - To prevent excessive numbers of connections, click **Add** under **Control Simultaneous Connections**.
- 8 Specify how you want to track traffic:
 - To track all traffic from a specific source or range of sources, choose **Source** from the **Track By** drop-down list, and enter a single IP address or address block in the **Network** field.
 - To track all traffic to a specific destination or range of destinations, choose **Destination** from the **Track By** drop-down list, and enter an IP address or address block in the **Network** field.

NOTE! To load-balance the traffic for maximum performance, the source and destination address and port are used to determine which Snort Instance the traffic is sent to.

- 9 Specify the triggering rate for the rate tracking setting:
 - For SYN attack (half-open, ‘embryonic’ connections) configuration, enter the number of SYN packets per number of seconds in the **Rate** fields.
 - For simultaneous connection configuration, enter the number of connections in the **Count** field.

NOTE! The recommended setting is between 600 - 6,000 TCP SYN/connection requests per minute per IP address. However, the exact number will vary and will depend on the host(s) and/or network configuration.

- 10 To drop packets matching the rate-based attack prevention settings, check the **Drop** check box.
- 11 In the **Timeout** field, enter the time period after which to stop generating events (and if applicable, dropping) for traffic with the matching pattern of SYNs or simultaneous connections.
- 12 Click **OK**.
- 13 To save changes you made in this policy since the last policy commit, click **Policy Information**, then click **Commit Changes**.

Audit Record:

2016-11-30 18:53:29 admin Network Analysis Policy > default > settings Changed Rate-Based Attack Prevention to "Enabled" (from "Disabled") 10.128.120.136

Specific Attacks

To detect specific attacks, please enable the following rules below:

Attack	Attack Category	Rule
Teardrop	IP Attack	Rule 123:2 "FRAG2_TEARDROP"
Bonk	IP Attack	Rule 123:4 "FRAG3_ANOMALY_OVERSIZE"
Boink	IP Attack	Rule 123:4 "FRAG3_ANOMALY_OVERSIZE"
Land	IP Attack	Rule 116:151 "DECODE_BAD_TRAFFIC_SAME_SRCDEST"
Nuke	ICMP Attack	alert tcp \$EXTERNAL_NET any -> \$HOME_NET 135:139 (msg:"SERVER-OTHER Winnuke attack"; flow:stateless; flags:U+; metadata:ruleset community; reference:bugtraq,2010; reference:cve,1999-0153; classtype:attempted-dos; sid:1257000; rev:15; gid:1001;)
Ping of Death	ICMP Attack	Rule 123:7 "FRAG3_ANOMALY_BADSIZE_LG"
Null flags	TCP Attack	alert tcp \$EXTERNAL_NET any -> \$HOME_NET any (msg:"Null TCP attack"; flags:0; classtype:attempted-dos; sid:269; rev:3;)
SYN+FIN flags	TCP Attack	Rule 116:420 "DECODE_TCP_SYN_FIN"
FIN only flags	TCP Attack	alert tcp any any -> any any (sid:1000003; gid:1; flags:F*; msg:"FIN only"; classtype:attempted-dos; rev:3;)
SYN+RST flags	TCP Attack	Rule 116:421 "DECODE_TCP_SYN_RST"
Bomb	UDP Attack	Rule 116:98 "DECODE_UDP_DGRAM_LONG_PACKET"
Chargen DoS	UDP Attack	Rule 1:271 "SERVER-OTHER echo+chargen bomb"

The FTP/Telnet decoder analyzes FTP and telnet data streams, normalizing FTP and telnet commands before processing by the rules engine.

You can set options for decoding on multiple FTP servers. Each server profile you create contains the server IP address and the ports on the server where traffic should be monitored. You can specify which FTP commands to validate and which to ignore for a particular server, and set maximum parameter lengths for commands. You can also set the specific command syntax the decoder should validate against for particular commands and set alternate maximum command parameter lengths.

Networks

Use this option to specify one or more IP addresses of FTP servers.

Ports

Use this option to specify the ports on the FTP server where the managed Device should monitor traffic. In the interface, list multiple ports separated by commas. Port21 is the well-known port for FTP traffic.

File Get Commands

Use this option to define the FTP commands used to transfer files from server to client. Do not change these values unless directed to do so by Support.

File Put Commands

Use this option to define the FTP commands used to transfer files from client to server. Do not change these values unless directed to do so by Support.

Additional FTP Commands

Use this line to specify the additional commands that the decoder should detect. Separate additional commands by spaces.

The HTTP Inspect preprocessor is responsible for:

- Decoding and normalizing HTTP requests sent to and HTTP responses received from web servers on your network.
- Separating messages sent to web servers into URI, non-cookie header, cookie header, method, and message body components to improve performance of HTTP-related intrusion rules.

Networks

Use this option to specify the IP address of one or more servers. You can specify a single IP address or address block, or a comma-separated list comprised of either or both.

Ports

The ports whose HTTP traffic the preprocessor engine normalizes. Separate multiple port numbers with commas.

HTTP Methods

Specifies HTTP request methods in addition to GET and POST that you expect the system to encounter in traffic. Use a comma to separate multiple values.

Intrusion rules use the *content* or *protected_content* keyword with the HTTP Method argument to search for content in HTTP methods. You can enable rule 119:31 to generate events when a method other than GET, POST, or a method configured forth is option is encountered in traffic.

The SMTP preprocessor instructs the rules engine to normalize SMTP commands. The preprocessor can also extract and decode email attachments in client-to-server traffic and, depending on the software version, extract email filenames, addresses, and header data to provide context when displaying intrusion events triggered by SMTP traffic.

Ports

Specifies the ports whose SMTP traffic you want to normalize. You can specify a value greater than or equal to 0. Separate multiple ports with commas.

Stateful Inspection

When selected, causes SMTP decoder to save state and provide session context for individual packets and only inspects reassembled sessions. When cleared, analyze each individual packet without session context.

Custom Commands

When **Normalize** is set to *Cmds*, normalizes the listed commands.

Detect Unknown Commands

Detects unknown commands in SMTP traffic.

You can enable rules124:5 to generate events for this option.

Checksum Verification

The system can verify all protocol-level checksums to ensure that complete IP, TCP, UDP, and ICMP transmissions are received and that, at a basic level, packets have not been tampered with or accidentally altered in transit. A checksum uses an algorithm to verify the integrity of a protocol in the packet. The packet is considered to be unchanged if the system computes the same value that is written in the packet by the end host.

Disabling checksum verification may leave your network susceptible to insertion attacks. Note that the system does not generate checksum verification events. In an inline deployment, you can configure the system to drop packets with invalid checksums.

NOTE! Do not disable checksum verification in the evaluated configuration.

Portscan Event Packet View

You can set any of the following options to **Enabled** or **Disabled** in a passive or inline deployment, or to **Drop** in an inline deployment:



- ICMP Checksums
- IP Checksums
- TCP Checksums
- UDP Checksums

To drop offending packets, in addition to setting an option to **Drop** you must also enable **Inline Mode** in the associated network analysis policy and ensure that the Device is deployed inline.

Setting these options to **Drop** in a passive deployment, or in an inline deployment in tap mode, is the same as setting them to **Enabled**.

The default for all checksum verification options is **Enabled**.

- 1 Login with Administrator Role or Intrusion Admin.

- 2 Select Policies > Access Control > Intrusion then click on [Network Analysis Policy](#).
- 3 Click the edit icon () next to the policy you want to edit.
- 4 Click [Settings](#).
- 5 If Checksum Verification under Transport/Network Layer Preprocessors is disabled, click Enabled.
- 6 Click the edit icon () next to [Checksum Verification](#).
- 7 For each protocol, click **Drop**.
- 8 To save changes you made in this policy since the last policy commit, click [Policy Information](#), then click [Commit Changes](#).

Audit Record:

2017-02-28 20:19:19	admin	Network Analysis Policy > default > settings > checksum_mode	Changed ICMP Checksums to "Drop and Generate Events" (from "Enabled")	10.128.120.150
2017-02-28 20:19:19	admin	Network Analysis Policy > default > settings > checksum_mode	Changed IP Checksums to "Drop and Generate Events" (from "Enabled")	10.128.120.150
2017-02-28 20:19:19	admin	Network Analysis Policy > default > settings > checksum_mode	Changed TCP Checksums to "Drop and Generate Events" (from "Enabled")	10.128.120.150
2017-02-28 20:19:19	admin	Network Analysis Policy > default > settings > checksum_mode	Changed UDP Checksums to "Drop and Generate Events" (from "Enabled")	10.128.120.150

VPN Functionality

FTD VPN Overview

A virtual private network (VPN) connection establishes a secure tunnel between endpoints over a public network such as the Internet. This chapter applies to Remote Access and Site-to-site VPNs on FTD Devices only. It describes the Internet Protocol Security (IPsec), and the Internet Key Exchange (IKE) standards that are used to build site-to-site and remote access VPNs.

VPN Licensing

There is no specific licensing for enabling FTD VPN, it is available by default.

The FMC determines whether to allow or block the usage of strong crypto on a FTD based on attributes provided by the smart licensing server. This is controlled by whether you selected the option to allow export-controlled functionality on the Device when you registered with Cisco Smart License Manager. If you are using the evaluation license, or you did not enable export-controlled functionality, you cannot use strong encryption.

VPN Types

The FMC supports the following types of VPN connections:

Remote Access VPNs on FTD

Remote access VPNs are secure, encrypted connections, or tunnels, between remote users and your company's private network. The connection consists of a VPN endpoint device, which is a workstation or mobile device with VPN client capabilities, and a VPN headend device, or secure gateway, at the edge of the corporate private network.

FTD can be configured to support Remote Access VPNs over TLS⁴ or IPsec IKEv2 by the FMC. Functioning as secure gateways in this capacity, they authenticate remote users, authorize access, and encrypt data to provide secure connections to your network. No other types of appliances, managed by the FMC, support Remote Access VPN connections.

FTD secure gateways support the AnyConnect Secure Mobility Client full tunnel client. This client is required to provide secure IPsec IKEv2 connections for remote users

Site-to-site VPNs on FTD

A site-to-site VPN connects networks in different geographic locations. You can create site-to-site IPsec connections between managed Devices, and between managed Devices and other Cisco or third-party peers that comply with all relevant standards. These peers can have any mix of inside and outside IPv4 and IPv6 addresses. Site-to-site tunnels are built using the Internet Protocol Security (IPsec) protocol suite and IKEv2. After the VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel.

⁴ TLS VPN is out of scope. In the evaluated configuration, only IPsec VPN is allowed.

VPN Basics

Tunneling makes it possible to use a public TCP/IP network, such as the Internet, to create secure connections between remote users and private corporate networks. Each secure connection is called a tunnel.

IPsec-based VPN technologies use the Internet Key Exchange (IKE) and IPsec tunneling standards to build and manage tunnels. IKE and IPsec accomplish the following:

- Negotiate tunnel parameters.
- Establish tunnels.
- Authenticate users and data.
- Manage security keys.
- Encrypt and decrypt data.
- Manage data transfer across the tunnel.
- Manage data transfer inbound and outbound as a tunnel endpoint or router.

A Device in a VPN functions as a bidirectional tunnel endpoint. It can receive plain packets from the private network, encapsulate them, create a tunnel, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets from the public network, decapsulate them, and send them to their final destination on the private network.

After the site-to-site VPN connection is established, the hosts behind the local gateway can connect to the hosts behind the remote gateway through the secure VPN tunnel. A connection consists of the IP addresses and hostnames of the two gateways, the subnets behind them, and the method the two gateways use to authenticate to each other.

On FTD, the system does not send VPN traffic until it has passed through the access control policy. Incoming tunnel packets are decrypted before being sent to the Snort process. Outgoing packets are processed by Snort before encryption. Identifying the protected networks for each endpoint node of a VPN tunnel determines which traffic is allowed to pass through the FTD and reach the internal hosts. In addition, the system does not send tunnel traffic to the public source when the tunnel is down.

The secret keys used for symmetric encryption, private keys, and CSPs used to generate keys, are zeroized immediately after use (for IPsec VPN functions, within FTD only), or on system shutdown (for all other functions). For plaintext keys unrelated to IPsec VPN like the TLS and SSH-related keys: the TOE destroys the reference to the keys stored in volatile memory directly followed by a request for garbage collection; the TOE destroys the abstraction that represents the key for keys stored in non-volatile storage the TSF.

Internet Key Exchange (IKE)

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2

negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection.

An IKE policy is a set of algorithms that two peers use to secure the IKE negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters protect subsequent IKE negotiations. In an IKEv2 policy, you can select multiple algorithms and modulus groups from which peers can choose during the Phase 1 negotiation. It is possible to create a single IKE policy, although you might want different policies to give higher priority to your most desired options. For site-to-site VPNs, you can create a single IKE policy.

To define an IKE policy, specify:

- A unique priority (1 to 65,543, with 1 the highest priority).
- An encryption method for the IKE negotiation, to protect the data and ensure privacy.
- A Hashed Message Authentication Codes (HMAC) method (called integrity algorithm in IKEv2) to ensure the identity of the sender, and to ensure that the message has not been modified in transit.
- For IKEv2, a separate pseudorandom function (PRF) used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. The options are the same as those used for the hash algorithm.
- A Diffie-Hellman group to determine the strength of the encryption-key-determination algorithm. The Device uses this algorithm to derive the encryption and hash keys.
- An authentication method, to ensure the identity of the peers.
- A limit to the time the Device uses an encryption key before replacing it.

When IKE negotiation begins, the peer that starts the negotiation sends all of its policies to the remote peer, and the remote peer searches for a match with its own policies, in priority order. A match between IKE policies exists if they have the same encryption, hash (integrity and PRF for IKEv2), authentication, and Diffie-Hellman values, and an SA lifetime less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter lifetime—From the remote peer policy—Applies.

IPsec

IPsec is one of the most secure methods for setting up a VPN. IPsec provides data encryption at the IP packet level, offering a robust security solution that is standards-based. With IPsec, data is transmitted over a public network through tunnels. A tunnel is a secure, logical communication path between two peers. Traffic that enters an IPsec tunnel is secured by a combination of security protocols and algorithms.

An IPsec Proposal policy defines the settings required for IPsec tunnels. An IPsec proposal is a collection of one or more crypto-maps that are applied to the VPN interfaces on the Devices. A crypto-map combines all the components required to setup IPsec security associations, including:

- A proposal (or transform set) is a combination of security protocols and algorithms that secure traffic in an IPsec tunnel. During the IPsec security association (SA) negotiation, peers search for a proposal that is the same at both peers. When it is found, it is applied

to create an SA that protects data flows in the access list for that crypto-map, protecting the traffic in the VPN. For IKEv2 proposals, you can configure multiple encryption and integration algorithms for a single proposal.

- A crypto map, combines all components required to setup IPsec security associations (SA), including IPsec rules, proposals, remote peers, and other parameters that are necessary to define an IPsec SA. When two peers try to establish an SA, they must each have at least one compatible crypto-map entry.

Deciding Which Encryption Algorithm to Use

When deciding which encryption algorithms to use for the IKE policy or IPsec proposal, your choice is limited to algorithms supported by the Devices in the VPN.

For IKEv2, you can configure multiple encryption algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order.

For IPsec proposals, the algorithm is used by the Encapsulating Security Protocol (ESP), which provides authentication, encryption, and anti-replay services. ESP is IP protocol type 50.

To operate in the CC-evaluated configuration you can choose from the encryption algorithms listed in **bold** below, as described below (not using AES-GCM-192 nor AES-192). The other algorithms are supported by the product, but not allowed in the CC-evaluated configuration.

Allowed in the CC-evaluated configuration:

- **AES-GCM:** Advanced Encryption Standard in Galois/Counter Mode is a block cipher mode of operation providing confidentiality and data-origin authentication, and provides greater security than AES. AES-GCM offers three different key strengths: 128-, 192-, and 256-bit keys, but the CC-evaluated configuration only allows use of **128-bit and 256-bit**. A longer key provides higher security but a reduction in performance. GCM is a mode of AES that is required to support NSA Suite B. NSA Suite B is a set of cryptographic algorithms that Devices must support to meet federal standards for cryptographic strength.
- **AES:** Advanced Encryption Standard is a symmetric cipher algorithm that provides greater security than DES and is computationally more efficient than 3DES. AES offers three different key strengths: 128-, 192-, and 256-bit keys, but the CC-evaluated configuration only allows use of **128-bit and 256-bit**. A longer key provides higher security but a reduction in performance.

Not allowed in the CC-evaluated configuration:

- **AES-GMAC:** Advanced Encryption Standard Galois Message Authentication Code is a block cipher mode of operation providing only data-origin authentication. It is a variant of AES-GCM that allows data authentication without encrypting the data. AES-GMAC offers three different key strengths: 128-, 192-, and 256-bit keys.

- **3DES:** Triple DES, which encrypts three times using 56-bit keys, is more secure than DES because it processes each block of data three times with a different key. However, it uses more system resources and is slower than DES.
- **DES:** Data Encryption Standard, which encrypts using 56-bit keys, is a symmetric secret-key block algorithm. It is faster than 3DES and uses less system resources, but it is also less secure. If you do not need strong data confidentiality, and if system resources or speed is a concern, choose DES.
- **Null:** A null encryption algorithm provides authentication without encryption. This is typically used for testing purposes only.

WARNING! In the evaluated configuration, only AES and AES-GCM are allowed.

Deciding Which Hash Algorithm to Use

In IKEv2 policies, the hash algorithm creates a message digest, which is used to ensure message integrity. In IKEv2, the hash algorithm is separated into two options, one for the integrity algorithm, and one for the pseudo-random function (PRF).

In IPsec proposals, the hash algorithm is used by the Encapsulating Security Protocol (ESP) for authentication. In IKEv2 IPsec Proposals, this is called the integrity hash.

For IKEv2, you can configure multiple hash algorithms. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order.

Allowed in the CC-evaluated configuration:

- **SHA (Secure Hash Algorithm):** SHA is more resistant to brute-force attacks than MD5. However, it is also more resource intensive than MD5. For implementations that require the highest level of security, use the SHA hash algorithm.
 - **SHA (SHA-1)** produces a 160-bit digest.

The following SHA-2 options, which are even more secure, are available for IKEv2 configurations. Choose one of these if you want to implement the NSA Suite B cryptography specification.

- **SHA256** Specifies the Secure Hash Algorithm SHA2 with the 256-bit digest.
- **SHA384** Specifies the Secure Hash Algorithm SHA2 with the 384-bit digest.
- **SHA512** Specifies the Secure Hash Algorithm SHA2 with the 512-bit digest.

Not allowed in the CC-evaluated configuration:

- **MD5 (Message Digest 5)**—Produces a 128-bit digest. MD5 uses less processing time for an overall faster performance than SHA, but it is considered to be weaker than SHA.
- **Null or None (NULL, ESP-NONE)**—(IPsec Proposals only) A null Hash Algorithm; this is typically used for testing purposes only. However, you should choose the null integrity algorithm if you select one of the AES-GCM/GMAC options as the encryption algorithm.

Even if you choose a non-null option, the integrity hash is ignored for these encryption standards.

WARNING! In the evaluated configuration, only SHA-1 and SHA-2 are allowed.

Deciding Which Diffie-Hellman Modulus Group to Use

You can use the following Diffie-Hellman key derivation algorithms to generate IPsec security association (SA) keys. Each group has a different size modulus. A larger modulus provides higher security, but requires more processing time. You must have a matching modulus group on both peers.

If you select AES encryption, to support the large key sizes required by AES, you should use Diffie-Hellman (DH) Group 5 or higher.

To implement the NSA Suite B cryptography specification, use IKEv2 and select one of the elliptic curve Diffie-Hellman (ECDH) options: 19, 20, or 21. Elliptic curve options and groups that use 2048-bit modulus are less exposed to attacks such as Logjam.

For IKEv2, you can configure multiple groups. The system orders the settings from the most secure to the least secure and negotiates with the peer using that order.

Allowed in the CC-evaluated configuration:

- Diffie-Hellman Group **14**: 2048-bit modulus.
- Diffie-Hellman Group **19**: 256-bit elliptic curve.
- Diffie-Hellman Group **20**: 384-bit elliptic curve.

Not allowed in the CC-evaluated configuration:

- Diffie-Hellman Group 1: 768-bit modulus.
- Diffie-Hellman Group 2: 1024-bit modulus.
- Diffie-Hellman Group 5: 1536-bit modulus.
- Diffie-Hellman Group 21: 521-bit elliptic curve.

WARNING! In the evaluated configuration, **only DH Groups 14, 19 and 20** are allowed.

Deciding Which Authentication Method to Use

Pre-shared keys and digital certificates are the methods of authentication available for VPNs.

Site-to-site VPN, which uses IKEv2, can use both options.

Remote Access VPN, which uses TLS and IPsec IKEv2 only, supports digital certificate authentication only.

Pre-shared keys allow for a secret key to be shared between two peers and used by IKE during the authentication phase. The same shared key must be configured at each peer or the IKE SA cannot be established.

Digital certificates use RSA or ECDSA key pairs to sign and encrypt IKE key management messages. Certificates provide non-repudiation of communication between two peers, meaning that it can be proved that the communication actually took place. When using this authentication method, you need a Public Key Infrastructure (PKI) defined where peers can obtain digital certificates from a Certification Authority (CA). CAs manage certificate requests and issue certificates to participating network devices providing centralized key management for all of the participating devices.

Pre-shared keys do not scale well, using a CA improves the manage ability and scalability of your IPsec network. With a CA, you do not need to configure keys between all encrypting devices. Instead, each participating device is registered with the CA, and requests a certificate from the CA. Each device that has its own certificate and the public key of the CA can authenticate every other device within a given CA's domain.

PKI Infrastructure

A PKI provides centralized key management for participating network devices. It is a defined set of policies, procedures, and roles that support public key cryptography by generating, verifying, and revoking public key certificates commonly known as digital certificates.

In public key cryptography, each endpoint of a connection has a key pair consisting of both a public and a private key. The key pairs are used by the VPN endpoints to sign and encrypt messages. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other, securing the data flowing over the connection.

Generate a general purpose RSA (2048, or 3072 bits) or ECDSA (256, 384, or 521 bits) key pair, used for both signing and encryption, or you generate separate key pairs for each purpose. Separate signing and encryption keys help to reduce exposure of the keys. TLS uses a key for encryption but not signing, however, IKE uses a key for signing but not encryption. By using separate keys for each, exposure of the keys is minimized.

Digital Certificates

When you use Digital Certificates as the authentication method for VPN connections, peers are configured to obtain digital certificates from a Certificate Authority (CA). CAs are trusted authorities that “sign” certificates to verify their authenticity, there by guaranteeing the identity of the device or user.

CA servers manage public CA certificate requests and issue certificates to participating network devices as part of a Public Key Infrastructure (PKI), this activity is called Certificate Enrollment. These digital certificates, also called identity certificates contain:

- The digital identification of the owner for authentication, such as name, serial number, company, department, or IP address.
- A public key needed to send and receive encrypted data to the certificate owner.
- The secure digital signature of a CA.

Certificates also provide non-repudiation of communication between two peers, meaning that it they prove that the communication actually took place.

Certificate Authority Certificates

In order to validate a peer's certificate, each participating device must retrieve the CA's certificate from the server. A CA certificate is used to sign other certificates. It is self-signed and called a root certificate. This certificate contains the public key of the CA, used to decrypt and validate the CA's digital signature and the contents of the received peer's certificate. The CA certificate may be obtained by:

- Using the Simple Certificate Enrollment Protocol (SCEP) to retrieve the CA's certificate from the CA server
 - Manually copying the CA's certificate from another participating device
-
- **NOTE:** SCEP is out in scope of the Common Criteria evaluation.

Trustpoints

Once enrollment is complete, a trustpoint is created on the managed Device. It is the object representation of a CA and associated certificates. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.

PKCS#12 File

A PKCS#12, or PFX, file holds the server certificate, any intermediate certificates, and the private key in one encrypted file. This type of file may be imported directly into a Device to create a trustpoint.

Revocation Checking

A CA may also revoke certificates for peers that no longer participate in your network. Revoked certificates are either managed by an Online Certificate Status Protocol (OCSP) server or are listed in a certificate revocation list (CRL). A peer may check these before accepting a certificate from another peer.

FTD Site-to-site VPN

FTD site-to-site VPN supports the following features:

- IKEv2 (use of IKEv1 is not allowed in the CC-evaluated configuration)
- Certificates and automatic or manual pre-shared keys for authentication.
- IPv4 & IPv6. All combinations of inside and outside are supported.
- IPsec IKEv2 Site-to-Site VPN topologies provide configuration settings to comply with Security Certifications.
- VPN alerts when the tunnel goes down.
- Tunnel statistics available using the FTD Unified CLI.

IPsec and IKE

In the FMC, site-to-site VPNs are configured based on IKE policies and IPsec proposals that are assigned to VPN topologies. Policies and proposals are sets of parameters that define the characteristics of a site-to-site VPN, such as the security protocols and algorithms that are used to secure traffic in an IPsec tunnel. Several policy types may be required to define a full configuration image that can be assigned to a VPN topology.

Authentication

For authentication of VPN connections, configure a pre-shared key in the topology, or a trustpoint on each device. Pre-shared keys allow for a secret key, used during the IKE authentication phase, to be shared between two peers. A trustpoint includes the identity of the CA, CA-specific parameters, and an association with a single enrolled identity certificate.



Extranet Devices

Each topology type can include Extranet devices, devices that you do not manage in FMC. These include:

- Cisco devices that FMC supports, but for which your organization is not responsible. Such as spokes in networks managed by other organizations within your company, or a connection to a service provider or partner's network.
- Non-Cisco devices. You cannot use FMC to create and deploy configurations to non-Cisco devices

Add non-Cisco devices, or Cisco devices not managed by the FMC, to a VPN topology as "Other" devices. Also specify the IP address of each remote device.

Managing FTD Site-to-site VPN


1. Login with Administrator Role.
2. For certificate authentication for your VPNs, you must prepare the Devices by allocating trustpoints as described in “FTD Certificate Based Authentication” section below.
3. Select **Devices > VPN > Site To Site** to manage your FTD Site-to-site VPN configurations and deployments. Choose from the following:
 - Add—To create a new VPN topology, click **Add VPN > Firepower Threat Defense Device**, and continue as instructed in “Configuring FTD Site-to-site VPN” section below.
 - Edit—To modify the settings of an existing VPN topology, click the edit icon (). Modifying is similar to configuring, continue as instructed above.
 - Delete—To delete a VPN deployment, click the delete icon ().
 - Deploy—Click **Deploy**.

Configuring FTD Site-to-site VPN

1. Login with Administrator Role.
2. Choose **Devices > VPN > Site To Site**. Then **Add VPN > Firepower Threat Defense**, or edit a listed VPN Topology.

3. Enter a unique **Topology Name**. We recommend naming your topology to indicate that it is a FTD VPN, and its topology type.
4. Choose the **Network Topology** for this VPN. For example, point to point topology.
5. Choose the IKE versions to use during IKE negotiations, **IKEv2**. Default is IKEv2.

NOTE! Always select IKEv2 to conform to the Common Criteria certified configuration.

6. Add Endpoints for this VPN deployment by clicking the add icon () for each node in the topology.

Open the **Endpoint** tab.

Fields

Devices

Choose an endpoint node for your deployment:

- A FTD managed by this FMC.
- An **Extranet** Device (Cisco or third-party) not managed by FMC.

Device Name

For Extranet Devices only, provide a name for this device. We recommend naming it such that it is identifiable as an un-managed Device.

Interface

If you chose a managed Device as your endpoint, choose an interface on that managed Device.

IP Address

- If you choose a Device not managed by the FMC, specify an IP address for the endpoint.
- If you chose a managed Device as an endpoint, choose a single IPv4 address or multiple IPv6 addresses from the drop-down list (these are the addresses already assigned to this interface on this managed Device).
- All endpoints in a topology must have the same IP addressing scheme. IPv4 tunnels can carry IPv6 traffic and vice-versa. The Protected Networks define which addressing scheme the tunneled traffic will use.

This IP is Private

Check the check box if the endpoint resides behind a firewall with network address translation (NAT).

Public IP address

If you checked on the **This IP is Private** check box, specify a public IP address for the firewall.

Connection Type


Specify the allowed negotiation as bidirectional, answer-only, or originate-only.

Certificate Map

Choose a pre-configured certificate map object, or add a certificate map object that defines what information is necessary in the received peer certificate for it

to be valid for VPN connectivity. Please see “FTD Certificate Map Object” section below.

Protected Networks

Define a list of networks protected by this VPN endpoint. Click the add icon () to select from available Network Objects or add Network Objects inline. VPN endpoints cannot have the same or overlapping IP address and protected networks.

7. Specify non-default IKE options for this VPN deployment.

Open the **IKE** tab.

Fields

Policy

Choose a predefined IKEv2 policy object or create a new one to use. For more details, please see “Configure IKEv2 Policy Object” section below.

Authentication Type

- **Pre-shared Manual Key**—Manually assign the pre-shared key that is used for this VPN. Specify the **Key** and then re-enter it in **Confirm Key** to confirm.

When this option is chosen for IKEv2, the **Enforce hex-based pre-shared key only** check box appears, check if desired. If enforced, you must enter a valid hex value for the key, an even number of 2-256 characters, using numerals 0-9, or A-F. Pre-shared keys can be entered as text (ASCII character) strings, or HEX values. The text-based pre-shared keys can be composed of any combination characters from four character sets: upper-case letters, lower-case letters, numbers (0-9), and special characters including “!”, “@”, “#”, “\$”, “%”, “^”, “&”, “*”, “(”, and “)”. For enhanced key strength test-based keys should include at least one character from each of the four character sets, and HEX-based keys should include a majority of the 16 HEX characters. The TOE supports text-based keys that are from 1 character in length up to 127 in length, and HEX-based keys up to 256 characters, but for enhanced security it is recommended that text-based keys be at least 22 characters, and HEX-based keys be at least 44 characters.

- **Certificate**—When you use Certificates as the authentication method for VPN connections, peers obtain digital certificates from a CA server in your PKI, and trade them to authenticate each other. In the **Certificate** field, select a pre-configured PKI Enrollment Object. This enrollment object is used to generate a trustpoint with the same name on the managed Device. The trustpoint is created when the PKI enrollment object is associated with that Device.

8. Specify non-default IPsec options for this VPN deployment.

Fields

Crypto-Map Type

A crypto map combines all the components required to setup IPsec security associations (SA). When two peers try to establish an SA, they must each have at least one compatible crypto-map entry. The proposals defined in the crypto-map entry are used in the IPsec security negotiation to protect the data flows specified by that crypto-map's IPsec rules. Choose static or dynamic for this deployment's crypto-map:

- **Static**—Use a static crypto-map in a point-to-point or full mesh VPN topology.
- **Dynamic**—Dynamic crypto-maps essentially create a crypto-map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements.


Dynamic crypto-map policies apply only in a hub-and-spoke VPN configuration. In a point-to-point or full mesh VPN topology, you can apply only static crypto-map policies.

IKEv2 Mode

For IPsec IKEv2 only, specify the encapsulation mode for applying ESP encryption and authentication to the tunnel. This determines what part of the original IP packet has ESP applied.

- **Tunnel mode**—(default) Encapsulation mode is set to tunnel mode. Tunnel mode applies ESP encryption and authentication to the entire original IP packet (IP header and data), hiding the ultimate source and destination addresses and becoming the payload in a new IP packet.
- **Transport preferred**—Encapsulation mode is set to transport mode with an option to fallback to tunnel mode if the peer does not support it. In Transport mode only the IP payload is encrypted, and the original IP headers are left intact. Therefore, the admin must select a protected network that matches the VPN interface IP address.
- **Transport required**—Encapsulation mode is set to transport mode only, falling back to tunnel mode is not allowed. If the endpoints cannot successfully negotiate transport mode, due to one endpoint not supporting it, the VPN connection is not established.

Proposals

Click () to specify the proposals for your chosen IKEv2 method. Select from the available **IKEv2 IPsec Proposals** objects, or create and then select a new one. See “Configure IKEv2 IPsec Proposal Object” section for more details.

Enable Security Association (SA) Strength Enforcement

Enabling this option ensures that the encryption algorithm used by the child IPsec SA is not stronger (in terms of the number of bits in the key) than the parent IKE SA.

Enable Reverse Route Injection

Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.

Enable Perfect Forward Secrecy

Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the pre-shared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the Modulus Group list.

Modulus Group

The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group.

Lifetime (seconds)

The number of seconds a security association exists before expiring. The default is 28,800 seconds.

Lifetime (kbytes)

The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires. The default is 4,608,000 kilobytes. No specification allows infinite data.

9. Specify non-default Advanced options for this VPN deployment.

Select the Advanced > IKE > Settings.

Fields

IKE Keepalive

Enable or disables IKE Keepalives. Or set to Enable Infinite specifying that the Device never starts keep alive monitoring itself.

Threshold

Specifies the IKE keep alive confidence interval. This is the number of seconds allowing a peer to idle before beginning keepalive monitoring. The minimum and default is 10seconds; the maximum is 3600 seconds.

Retry Interval

Specifies number of seconds to wait between IKE keep alive retries. The default is 2 seconds, the maximum is 10 seconds.

Identity Sent to Peers:

Choose the Identity that the peers will use to identify themselves during IKE negotiations:

- **autoOrDN(default)**—Determines IKE negotiation by connection type: IP address for pre-shared key, or Cert DN for certificate authentication (not supported).
- **ipAddress**—Uses the IP addresses of the hosts exchanging IKE identity information.
- **hostname**—Uses the fully qualified domain name of the hosts exchanging IKE identity information. This name comprises the hostname and the domain name.

Enable Aggressive Mode

Available only in a hub-and-spoke VPN topology. Select this negotiation method for exchanging key information if the IP address is not known and DNS resolution might not be available on the Devices. Negotiation is based on hostname and domain name.

Open the **Advanced** Tab and select **Tunnel** in the navigation pane.

Fields

NAT Settings

Keepalive Messages Traversal—Elect whether to enable NAT keepalive message traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there is a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow.

If you select this option, configure the **Interval**, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 5 to 3600 seconds. The default is 20 seconds.

Certificate Map Settings

- **Use the certificate map configured in the Endpoints to determine the tunnel**—If this option is enabled (checked), the tunnel will be determined by matching the contents of the received certificate to the certificate map objects configured in the endpoint nodes.
 - **Use the certificate OU field to determine the tunnel**—Indicates that if a node is not determined based on the configured mapping (the above option) if selected, then use the value of the organizational unit (OU) in the subject distinguished name (DN) of the received certificate to determine the tunnel.

- **Use the IKE identity to determine the tunnel**—Indicates that if a node is not determined based on a rule matching or taken from the OU (the above options) if selected, then the certificate-based IKE sessions are mapped to a tunnel based on the content of the phase1 IKE ID.
- **Use the peer IP address to determine the tunnel**—Indicates that if a tunnel is not determined based on a rule matching or taken from the OU or IKE ID methods (the above options) if selected, then use the established peer IP address.

10. Click **Save**.

FTD Certificate Based Authentication

Installing a certificate using manual enrollment.

1. Login with Administrator Role.
2. On the Devices > Certificates screen, choose Add > Add New Certificate to open the Add New Certificate dialog.
3. Choose a Device from the **Device** drop down list.
4. Associate a certificate enrollment object with this Device in one of the following ways:
 - Choose a Certificate Enrollment Object of the appropriate type from the drop-down list.
 - Click (+), to add a new Certificate Enrollment Object. Please see “Adding Certificate Enrollment Object” section below.
5. Press **Install**, to initiate the manual enrollment process.

The **CA Certificate** status will go from *In Progress* to *Available* as the FMC installs the CA certificate (provided in the enrollment object) on the managed Device, authenticates the CA Server, and creates a trustpoint on the managed Device.

The **Identity Certificate** status will reach Pending state when the Certificate Signing Request (CSR) is generated by the managed Device and placed in the Identity Certificate field.

6. Execute the appropriate activity with your PKI CA Server to obtain an identity certificate.
 - a) Click the **Identity Certificate** magnifying glass to view and copy the CSR.
 - b) Execute the appropriate activity with your PKI CA Server to obtain an identity certificate using this CSR. This activity is completely independent of the FMC or the managed Device. When complete, You will have an Identity Certificate for the managed Device. You can copy it or place it in a file.
 - c) To finish the manual process, install the obtained identity certificate on to the managed Device.

Return to the FMC dialog to paste the Identity Certificate into its field. Or, select **Browse** to choose the identity certificate file.

7. Select **Import** to import the Identity Certificate.

The Identity Certificate status will be *Available* when the import complete.

8. Click the magnifying glass to view the **Identity Certificate** for this Device.

Installing a certificate by importing a PKCS12 file. A PKCS12 file size should not be larger than 24K.

1. Login with Administrator Role.
2. Go to Devices > Certificates, then click + Add > Import PKCS12 File to open the Import PKCS12 File dialog.
3. Choose a pre-configured managed Device from the **Device** drop down list.
4. Specify a Certificate Enrollment type of PKCS12.
5. Select **Browse** to find and choose your PKCS#12 Certificate file.
6. Enter the **Passphrase** for decryption.
7. Press **Add**

For file import, the **CA Certificate** and **Identity Certificate** status will go from *In Progress* to *Available* as it installs the PKCS12 file on the Device.

8. Once *Available*, click the magnifying glass to view the Identity Certificate for this Device.

Adding Certificate Enrollment Objects

1. Login with Administrator Role.
2. Open the Add Cert Enrollment dialog:
 - Directly from Object Management: In the **Objects > Object Management** screen, choose **PKI > Cert Enrollment** from the navigation pane, and press **Add Cert Enrollment**.
 - While configuring a managed Device: In the **Devices > Certificates** screen, choose **Add > Add New Certificate** and click (+) for the **Certificate Enrollment** field.
3. Enter the **Name**, and optionally, a **Description** of this enrollment object.

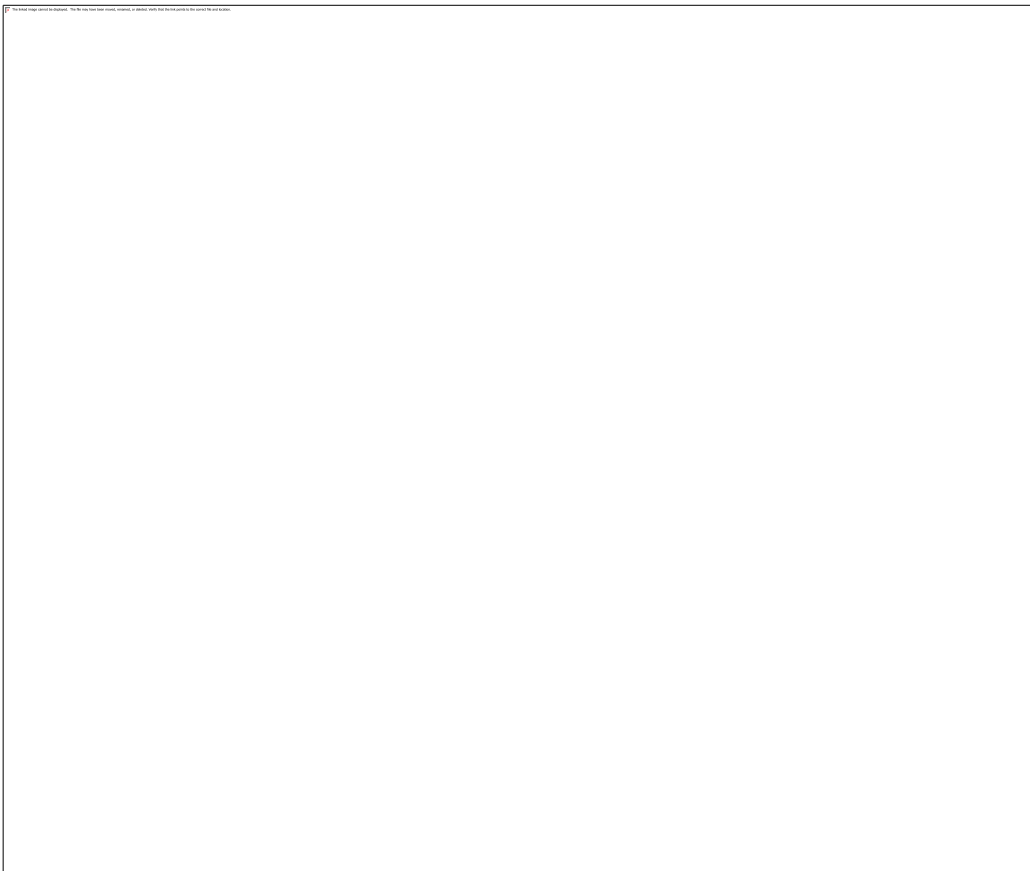
When enrollment is complete, this name is the name of the trustpoint on the managed Devices with which it is associated.

4. Open the **CA Information** tab and choose the **Enrollment Type**.
 - **Self-Signed Certificate**—The managed Device, acting as a CA, generates its own self-signed root certificate. No other information is needed in this pane.

- **SCEP**—(Default) Simple Certificate Enrollment Protocol. Specify the SCEP information.
- **Manual**—Paste an obtained CA certificate in the **CA Certificate** field. You can obtain a CA certificate by copying it from another device.
- **PKCS12 File**—Import a PKCS12 file on a FTD managed Device that supports VPN connectivity. A PKCS#12, or PFX, file holds a server certificate, intermediate certificates, and a private key in one encrypted file.

NOTE! Self-signed certificate and SCEP are out of scope.

5. Open the **Certificate Parameters** tab and specify the certificate contents, which will be included in the Certificate Signing Request (CSR). Specify information to be included in the CSR sent to the CA server, and specify at least: Country Code (C), Organization (O), Organization Unit (OU), and Common Name (CN).




6. Open the **Key** tab and follow guidance in the bullets below.

Fields

- **Key Type**—RSA, or ECDSA
- **Key Name**—If the key pair you want to associate with the certificate already exists, this field specifies the name of that key pair. If the key pair does not exist, this field

specifies the name to assign to the key pair that will be generated during enrollment. If you do not specify an RSA key pair, the fully qualified domain name (FQDN) key pair is used instead.

- **Key Size**—If the key pair does not exist, defines the desired key size (modulus), in bits. When the selected Key Type is RSA multiple key sizes are selectable, but in the CC- configuration the only allowed RSA key size 2048. When the selected Key Type is ECDSA multiple key sizes are selectable, all of which are allowed in the CC- configuration, so select any one of 256, 384, or 521.



NOTE: This key generation process on FMC results in private keys stored in binary format on FMC and also on the FTD that will use the certificate associated with the generated keys. The private keys are stored in FMC and FTD in locations that are not accessible via any administrative interface, and only the unique identification of each key's associated PKI certificate is visible (on FMC via **Devices > Certificates**; and on FTD by viewing output of **"show crypto ca certificates"**).

7. Click the **Revocation** tab, and specify the revocation options. Specify whether to check the revocation status of a certificate by choosing and configuring the method. Revocation checking is off by default, neither method (CRL or OCSP) is checked.

Fields

- **Enable Certificate Revocation Lists**—Check to enable CRL checking.

- **Use CRL distribution point from the certificate**—Check to obtain the revocation lists distribution URL from the certificate.
- **Use static URL configured**—Check this to add a static, pre-defined distribution URL for revocation lists. Then add the URLs.

CRL Server URLs—The URL of the server from which the CRL can be downloaded.

- Enable Online Certificate Status Protocol (OCSP)—Check to enable OCSP checking.

OCSP Server URL—The URL of the OCSP server checking for revocation if you require OCSP checks. This URL for the OCSP responder must start with http://.

- Consider the certificate valid if revocation information cannot be reached—Checked by default. Uncheck if you do not want to allow this.
- If IPsec sessions fail due to inability to contact the CRL or OCSP server, restore connectivity to the CRL or OCSP server before reattempting to establish the TLS sessions.

FTD Certificate Map Object

Certificate Map objects are a named set of certificate matching rules. These objects are used to provide an association between a received certificate and a Remote Access VPN connection profile. Connection Profiles and Certificate Map objects are both part of a remote access VPN policy. If a received certificate matches the rules contained in the certificate map, the connection is "mapped", or associated with the specified connection profile. The rules are in priority order, they are matched in the order they are shown in the UI. The matching ends when the first rule within the Certificate Map object results in a match.

1. Login with Administrator Role.
2. Open the Objects > Object Management > VPN > Certificate Map dialog.

Fields


Map Name—Identify this object so it can be referred to from other configurations, such as a name or the location of the site-to-site VPN peer.

Mapping Rule—Specify the contents of the certificate to evaluate. If the certificate satisfies these rules, the site-to-site VPN peer will be mapped to the connection profile containing this object.

- **Field**—Select the field for the matching rule to Alternative Subject, which is the Subject Alternative Name (SAN).
- **Component**—When the Field is set to *Alternative Subject* the Component will be frozen as *Whole* Field.
- **Operator**—Set the operator for the matching rule to Equals.
- **Value**—Input the Fully Qualified Domain Name (FQDN) of the site-to-site VPN peer.

Configure IKEv2 Policy Object

Use the IKEv2 policy dialog box to create, delete, and edit an IKEv2 policy object. These policy objects contain the parameters required for IKEv2 policies.

1. Login with Administrator Role.
2. Choose **Objects > Object Management** and then **VPN > IKEv2 Policy** from the table of contents.
3. Choose  **Add IKEv2 Policy** to create a new policy.
4. Enter a **Name** for this policy.

The name of the policy object. A maximum of 128 characters is allowed.

5. Enter a **Description** for this policy.


A description of the policy object. A maximum of 1024 characters is allowed.

6. Enter the **Priority**.

The priority value of the IKE proposal. The priority value determines the order of the IKE proposals compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, it tries to use the parameters defined in the next lowest priority policy. Valid values range from 1 to 65535. The lower the number, the higher the priority. If you leave this field blank, FMC assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

7. Set the **Lifetime** of the security association (SA), in seconds. You can specify a value from 120 to 2,147,483,647 seconds. The default is 86400.
8. Choose the **Integrity Algorithms** portion of the Hash Algorithm used in the IKE policy. The Hash Algorithm creates a Message Digest, which is used to ensure message integrity.
9. Choose the **Encryption Algorithm** used to establish the Phase 1 SA for protecting Phase 2 negotiations.
10. Choose the **PRF Algorithm**.
11. Select and **Add** a **DH Group**.
12. Click **Save**.

Configure IKEv2 IPsec Proposal Object

1. Login with Administrator Role.
2. Choose **Objects > Object Management** and then **VPN > IKEv2 IPsec Proposal** from the table of contents.
3. Choose  **Add IKEv2 IPsec Proposal** to create a new Proposal.
4. Enter a **Name** for this Proposal.

The name of the proposal object. A maximum of 128 characters is allowed.

5. Enter a **Description** for this Proposal.

A description of the proposal object. A maximum of 1024 characters is allowed.

6. Choose the **ESP Hash** method, the hash or integrity algorithm to use in the Proposal for authentication.
7. Choose the **ESP Encryption** method. The Encapsulating Security Protocol (ESP) encryption algorithm for this Proposal.
8. Click **Save**.

FTD Remote Access VPN

FTD provides secure gateway capabilities that support remote access TLS⁵ and IPsec IKEv2 VPNs. The full tunnel client, AnyConnect Secure Mobility Client, provides secure TLS and IKEv2 IPsec connections to the security gateway for remote users. It is the only client supported on endpoint devices for remote VPN connectivity to FTD. The client gives remote users the benefits of an TLS or IKEv2 IPsec VPN client without the need for network administrators to install and configure clients on remote computers. The AnyConnect mobile client for Windows, Mac, and Linux is deployed from the secure gateway upon connectivity. The AnyConnect apps for Apple iOS and Android devices are installed from the platform app store.

Use the Remote Access VPN Policy wizard in the FMC to quickly and easily set up these two types of remote access VPNs with basic capabilities. Then, enhance the policy configuration if desired and deploy it to your FTD secure gateway devices.

Note: If a Remote Access VPN session is being used to tunnel SSH for remote administrative access to FTD (SSH over IPsec) and the IPsec connection between from the VPN client (AnyConnect) and FTD is unintentionally broken the AnyConnect client will automatically attempt to re-initiate the IPsec connection. If AnyConnect is able to automatically reestablish the IPsec session the tunneled SSH session, if previously established, may have remained active though it may have timed out and thus would need to be reinitiated from the client. In most cases no administrative action is required (on the AnyConnect client nor on FTD) though any connectivity issues will need to be resolved in the networks between AnyConnect and FTD. If IPsec connectivity has been lost for an extended period of time AnyConnect will discontinue automatic attempts to reestablish the tunnel, in which case the local administrative user of AnyConnect must re-initiate the IPsec tunnel, and reinitiate the SSH session through IPsec to FTD.

Managing FTD Remote Access VPN




1. Login with Administrator Role.
2. Choose Devices > VPN > Remote Access.

The policies displayed in the list were created using the VPN Configuration Wizard, and possibly already edited. Out of date status indicates there is an older version of the remote access VPN


⁵ Remote access TLS VPN is not in scope.

policy on the targeted Devices. Deploy the latest remote access VPN policy to update the policy configuration.

3. Choose from the following actions:

- Add ()—Creates a new Remote Access VPN Policy using a wizard that walks you through a basic policy configuration.
- Edit ()— Modify an existing Remote Access VPN policy. Click the edit icon or the VPN policy row to open the policy for editing.
- Delete ()— Delete a Remote Access VPN configuration.

Editing FTD Remote Access VPN Policy

1. Login with Administrator Role.
2. Choose **Devices > VPN > Remote Access**.
3. Select an existing Remote Access policy in the list and click the corresponding Edit icon ().
4. To add or edit a **Connection Profile**. Remote Access VPN policy contains the Connection Profiles targeted for specific devices. These policies pertain to creating the tunnel itself, such as, how authentication is accomplished, and how addresses are assigned (DHCP or Address Pools) to VPN clients. They also include user attributes, which are identified in group policies configured on the FTD. A Device also provides a default connection profile named *DefaultWEBVPNGroup*. The connection profile that is configured using the wizard appears in the list.

The Connection Profile page lists the profiles created under the Remote Access VPN policy. The table lists information about client address assignment, group policy, and authentication options.

To add a connection profile, choose the **Add** icon and specify the following in the **Add Connection Profile** window:

- **Connection Profile**—Provide a name that the remote users will use for VPN connections. Specifies a set of parameters that define how the remote users connect to the VPN device. For more information about Connection Profile, see the “Adding and Editing FTD Remote Access VPN Connection Profile” section below.
- **Group Policy**—A collection of user-oriented attributes which are applied to the client when the VPN connectivity is established. Group policies configure common attributes for groups of users. For more information about Group Policy, see the “Configuring Group Policies” section below.

Whether using the **Add Connection Profile** or the **Edit Connection Profile** window, use the **AAA** tab to set the Authentication Method. In the CC-evaluated configuration

authentication of remote access clients must include use of X.509v3 certificates, and must use the Distinguished Name (DN) in the client's certificate to uniquely identify the client. On the **AAA** tab:

- Set the **Authentication Method** to "Client Certificate Only".
- Expand the "**Map username from client certificate**" options.
- Click the "**Use entire DN (Distinguished Name) as username**" radio button.
- Make other changes as needed in the Connection Profile, then click **Save**.

5. From a Connection Profile, add or edit **Access Interfaces**.

- **Access Interface**—The interface group or security zone to which the interface belongs. Select the value from the drop-down list. The interface group or security zone must be a **Routed** type. Other interface types are not supported for Remote Access VPN connectivity. Associate the **Protocol** object with the access interface.
- **Enable IKEv2**—Select this option to enable IKEv2 settings.
- **Configure Interface Specific Identity Certificate**. Select Interface Identity Certificate from the drop-down list.

6. Choose the **Advanced** tab to complete the Remote Access VPN configuration:

- **Configuring the Any Connect Client Images**

The Cisco AnyConnect Secure Mobility client provides secure TLS or IPsec (IKEv2) connections to the FTD for remote users with full VPN profiling to corporate resources. Without a previously-installed client, remote users can enter the IP address of an interface configured to accept clientless VPN connections in their browser to download and install the AnyConnect client. The FTD downloads the client that matches the operating system of the remote computer. After downloading, the client installs and establishes a secure connection. In the case of a previously installed client, when the user authenticates, the FTD, examines the revision of the client, and upgrades the client as necessary.

- i. Click the **Add** icon in the **Available AnyConnect Images** portion of the **AnyConnect Images** dialog.
 - ii. Enter then **Name**, **File Name**, and **Description** for the available AnyConnect Image.
 - iii. Click **Browse** to navigate to the location for selecting the client image to be uploaded.
 - iv. Click **Save** to upload the image in the FMC.
- **Configuring the Address Assignment Policy**

The FTD can use IPv4 or IPv6 policy for assigning IP addresses to Remote Access VPN clients. If you configure more than one address assignment method, the FTD tries each of the options until it finds an IP address.

You can use the IPv4 or IPv6 policy to find an IP address to the Remote Access VPN clients. Firstly, you must try with the IPv4 policy and later followed by IPv6 policy.

- i. **Use DHCP**—Obtains IP addresses from a DHCP server configured in a connection profile. You can also define the range of IP addresses that the DHCP server can use by configuring DHCP network scope in the group policy. If you use DHCP, configure the server in the **Objects > Object Management > Network** pane. This method is available for IPv4 assignment policies.
 - ii. **Use an internal address pool**—Internally configured address pools are the easiest method of address pool assignment to configure. If you use this method, create the IP address pools in **Objects > Object Management > Address Pools** pane and select the same in the connection profile. This method is available for both IPv4 and IPv6 assignment policies.
 - iii. **Reuse an IP address so many minutes after it is released**—Delays the reuse of an IP address after its return to the address pool. Adding a delay helps to prevent problems firewall scan experience when an IP address is reassigned quickly. By default, the delay is set to zero, meaning the FTD does not impose a delay in reusing the IP address. If you want to extend the delay, enter the number of minutes in the range 0-480 to delay the IP address reassignment. This configurable element is available for IPv4 assignment policies.
- Configuring the **Certificate Maps**

Certificate to connection profile maps are used for certificate authentication on secure gateways.

- i. Set the General Settings for Certificate Group Matching

Select any, or all, of the following options to establish authentication and to determine to which connection profile (tunnel group) to map the client. Selections are priority-based, if a match is not found for the first selection matching continues down the list of options. When the rules are satisfied, the mapping is done. If the rules are not satisfied, the default connection profile (listed at the bottom) is used for this connection.

- Use Group URL if Group URL and Certificate Map match different Connection profiles

- Use the configured rules to match a certificate to a Connection Profile—Enable this to use the rules defined here in the Connection Profile Maps
- ii. Add Certificate to Connection Profile Mapping for this policy.
 1. Click **Add**.
 2. Choose or create a **Certificate Map** Object.
 3. Specify the **Connection Profile** that is used if the rules in the certificate map object are satisfied.
 4. Click **Save**.
- **Configuring Group Policies**

A Group Policy is a set of attribute and value pairs, stored in a group policy object, that define the remote access VPN experience. For example, in the group policy object, you configure general attributes such as addresses, protocols, and connection settings.

 - i. Select more group policies to associate with this Remote Access VPN policy. These are above and beyond the default group policy assigned at RAVPN policy creation time. Click **Add**. To create a group policy, please see “Configure Group Policy Object” section below.
 - ii. Click **OK** when you have the **Selected Group Policy** window set as desired.
- **Configuring the IPsec**
 - i. **Crypto Maps**—The Crypto Maps page lists the interface groups on which IKEv2 protocol is enabled. Crypto Maps are auto generated for the interfaces on which IKEv2 protocol is enabled.
 1. Select IPsec > Crypto Maps
 2. Click on **Interface Group**, the interface group on which IKEv2 protocol is enabled.
 3. On **IKEv2 IPsec Proposals**, click **Edit** to specify the proposals for your chosen IKEv2 method. On the IKEv2 IPsec Proposal dialog box, select from the available Transform Sets, or create a new IKEv2 IPsec proposal.
 4. **Enable Reverse Route Injection**—enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint.
 5. **Enable Perfect Forward Secrecy**—Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the

exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the pre-shared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the **Modulus Group** list.

6. **Lifetime Duration (seconds)**—The lifetime of the security association (SA), in seconds. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be setup more quickly than with shorter lifetimes. You can specify a value from 120 to 2,147,483,647 seconds. The default is 28,800 seconds, 8 hours.
 7. **Lifetime Size (kbytes)**—The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires. You can specify a value from 10 to 2,147,483,647 kilobytes. The default is 4,608,000 kilobytes. No specification allows infinite data.
- ii. IKE— Click the **Add** button to select from the available IKEv2 policies or add a new IKEv2 policy.
1. **Name**—Name of the IKEv2 policy.
 2. **Lifetime**—Specify the key lifetime in seconds, configurable from 120 - 2,147,483,647 seconds (the default 86,400 seconds, 24 hours).
 3. **Integrity**—The Integrity Algorithms portion of the Hash Algorithm used in the IKEv2 policy.
 4. **Encryption**—The Encryption Algorithm used to establish the Phase 1 SA for protecting Phase 2 negotiations.
 5. **PRF Hash**—The pseudorandom function (PRF) portion of the Hash Algorithm used in the IKE policy. With IKEv2, you can specify different algorithms for these elements.
 6. **DH Group**—The Diffie-Hellman group used for encryption.
- iii. IPsec
1. **Identity Sent to Peers**— Choose the Identity that the peers will use to identify themselves during IKE negotiations:

- a. **Auto**—Determines IKE negotiation by connection type: IP address for pre-shared key, or Cert DN for certificate authentication (not supported).
- b. **IPAddress**—Uses the IP addresses of the hosts exchanging IKE identity information.
- c. **Hostname**—Uses the fully qualified domain name of the hosts exchanging IKE identity information. This name comprises the hostname and the domain name.

2. NAT Settings

- a. **Keepalive Messages Traversal**—Select whether to enable NAT keepalive message traversal. NAT traversal keepalive is used for the transmission of keepalive messages when there is a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow. If you select this option, configure the interval, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 10 to 3600 seconds. The default is 20 seconds.
- b. **Interval**—Sets the NAT keepalive interval, from 10 to 3600 seconds. The default is 20 seconds.

7. Click **Save**.

8. Click **Deploy** to deploy the configuration changes on the FTD.

Configure Group Policy Object

1. Login with Administrator Role.
2. Choose **Objects > Object Management > VPN > Group Policy**. Click **Add Group Policy** or choose a current policy to edit, then select the **Advanced** tab.

Session Settings Fields

1. **Access Hours**—Choose or create a time range object. This object specifies the range of time this group policy is available to be applied to a remote access user.
2. **Idle Timeout/Alert Interval**—Specifies user's idle timeout period in minutes. If there is no communication activity on the user connection in this period, the system stops the connection. The minimum time is 1 minute. The default is 30 minutes. The Alert interval specifies the interval of time before idle time is reached to display a message to the user.

Restrict VPN Client Connections by Source IP

Optionally, it's possible to restrict which source IP addresses are permitted to attempt to negotiate IKE connections with the FTD. Note: Using these restrictions will impact Remote Access (VPN Client) IPsec connections, as well as peer-to-peer IPsec VPNs.

1. Create a flex config object with the access control list configuration.
 - a. Objects > Object Management > FlexConfig > FlexConfig Object > Add FlexConfig Object.
 - b. Example configuration:


```
access-li vpn-block extended deny udp host 192.168.1.2 any eq 4500
access-li vpn-block extended deny udp host 192.168.1.2 any eq 500
access-li vpn-block extended permit ip any any
access-group vpn-block in interface outside control-plane
```
 - c. Select 'Everytime' Deployment
 - d. Select 'Append' Type
2. Create a FlexConfig policy and add the object created in the step #1.
 - a. Devices > FlexConfig > New Policy > User Defined (select object) > Add (click the ">" button) > Save.
 - b. Deploy the policy.
3. Notes:
 - a. Make sure the access list name given in the FlexConfig object is not conflicting with any other access list name
 - b. When the VPN interface name is changed, remove the flex config object from the flex config policy and add it back once the name change is deployed.
 - c. In the text of the FlexConfig object, don't use the full command "access-list". If you do, the command will be rejected, so these examples intentionally use "access-li".
 - d. The example above would deny all types of connections from a specific source host address, but a specific subnet could be used instead, or a series of subnets can be added as separate lines in the access-list.
 - e. If desired, you variables in the FlexConfig object to substitute for IP address, subnets, interface names, or UDP ports.
 - f. End the ACL with "permit ip any any", which will allow all other traffic to be filtered by the FTD's other ACLs and intrusion rules.
 - g. This ACL must be applied to the interface's "control-plane" to block incoming IKE connections before they would otherwise get negotiated by the FTD.

Vulnerability Mitigation

The following steps need to be followed to ensure that the TOE is operating with all potential vulnerabilities mitigated -

- Cisco VDB Fingerprint Database version needs to be up to date. The “Update the Vulnerability Database (VDB) Manually” section of the FMC-CG lists the instructions to update the Cisco VDB Fingerprint Database.
- Follow all instructions listed in Section 4.2.8 of the “Cisco FTD 7.0 on ASA 5500 and ISA 3000 and FTDv with FMC/FMCv Common Criteria Supplemental User Guide” to add a rule to only access from a trusted subnet by listing them in the Access List.
- The “Allow External Database Access” option in the FMC should not be enabled to avoid the FMC connecting to an external SQL database. This connection is disabled by default.
- ESXi 6.7 and 7.0 should be updated in the FMCv TOE to include the latest patches – ESXi670-202210101-SG for ESXi 6.7 and ESXi70U3si-20841705 for ESXi 7.0 respectively.