

Policy Management

The following topics describe how to manage various policies on the Firepower Management Center:

- Requirements and Prerequisites for Policy Management, on page 1
- Policy Deployment, on page 2
- Policy Comparison, on page 13
- Policy Reports, on page 14
- Out-of-Date Policies, on page 15
- Performance Considerations for Limited Deployments, on page 16

Requirements and Prerequisites for Policy Management

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Network Admin
- Security Approver

Policy Deployment

Caution Do NOT push the FMC deployments over a VPN tunnel that is terminating directly on the Firepower Threat Defense. Pushing the FMC deployments can potentially inactivate the tunnel and disconnect the FMC and the Firepower Threat Defense.

Recovering the device from this situation can be very disruptive and require executing the disaster recovery procedure. This procedure resets the Firepower Threat Defense configuration to factory defaults by changing manager from FMC to local and configuring the device from beginning. For more information, see Deploying the FMC Policy Configuration over VPN Tunnel, on page 2.

After you configure your deployment, and any time you change that configuration, you must deploy the changes to affected devices. You can view deployment status in the Message Center.

Deploying updates the following components:

- Device and interface configurations
- Device-related policies: NAT, VPN, QoS, platform settings
- Access control and related policies: DNS, file, identity, intrusion, network analysis, prefilter, SSL
- Network discovery policy
- Intrusion rule updates
- Configurations and objects associated with any of these elements

You can configure the system to deploy automatically by scheduling a deploy task or by setting the system to deploy when importing intrusion rule updates. Automating policy deployment is especially useful if you allow intrusion rule updates to modify system-provided base policies for intrusion and network analysis. Intrusion rule updates can also modify default values for the advanced preprocessing and performance options in your access control policies.

In a multidomain deployment, you can deploy changes for any domain where your user account belongs:

- Switch to an ancestor domain to deploy changes to all subdomains at the same time.
- Switch to a leaf domain to deploy changes to only that domain.

Best Practices for Deploying Configuration Changes

The following are guidelines for deploying configuration changes.

Deploying the FMC Policy Configuration over VPN Tunnel

You can deploy the FMC policy configuration over a VPN tunnel, only if the deployment is for a device that does not terminate the tunnel. The FMC to Firepower Threat Defense management traffic should be its own secure transport SF tunnel and does not need to be over S2S VPN tunnel for any connectivity.

For policy-based VPN tunnel, choose the protected networks on both side to exclude the FMC to Firepower Threat Defense management traffic. For route-based VPN tunnel, configure the routing to exclude the FMC to Firepower Threat Defense management traffic to the VTI interface.

When you push the FMC deployments over the VPN tunnel with the management traffic that is also passing through the tunnel, in the event of any VPN misconfiguration, it inactivates the tunnel and results in disconnecting the FMC and the Firepower Threat Defense.

To reinstantiate the tunnel configuration, you can either:

- Remove the sensor from the Firepower Threat Defense and the FMC (resulting in losing all of its configuration), and then add the sensor again to the FMC.

  Or

- Contact Cisco TAC.

Note Reinstantiating the tunnel configuration requires overhauling of the system.

Inline vs Passive Deployments

Do not apply inline configurations to devices deployed passively, and vice versa.

Time to Deploy and Memory Limitations

The time it takes to deploy depends on multiple factors, including (but not limited to):

- The configurations you send to the device. For example, if you dramatically increase the number of Security Intelligence entries you block, deploy can take longer.
- Device model and memory. On lower-memory devices, deploying can take longer. For example, it can take up to five minutes to deploy to a Firepower 7010, 7020, or 7030 device.

Do not exceed the capability of your devices. If you exceed the maximum number of rules or policies supported by a target device, the system displays a warning. The maximum depends on a number of factors—not only memory and the number of processors on the device, but also on policy and rule complexity. For information on optimizing policies and rules, see Best Practices for Access Control Rules.

Interruptions to Traffic Flow and Inspection During Deploy

When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 8 and Configurations that Restart the Snort Process When Deployed or Activated, on page 10.

Caution We *strongly* recommend you deploy in a maintenance window or at a time when interruptions will have the least impact.

Auto-Enabling Application Detectors

If you are performing application control but disable required detectors, the system automatically enables the appropriate system-provided detectors upon policy deploy. If none exist, the system enables the most recently modified user-defined detector for the application.

Asset Rediscovery with Network Discovery Policy Changes

When you deploy changes to a network discovery policy, the system deletes and then rediscovers MAC address, TTL, and hops information from the network map for the hosts in your monitored networks. Also, the affected managed devices discard any discovery data that has not yet been sent to the FMC.

Related Topics

Snort® Restart Scenarios, on page 6

Deploy Configuration Changes



Caution Do NOT push the FMC deployments over a VPN tunnel that is terminating directly on the Firepower Threat Defense. Pushing the FMC deployments can potentially inactivate the tunnel and disconnect the FMC and the Firepower Threat Defense.

Recovering the device from this situation can be very disruptive and require executing the disaster recovery procedure. This procedure resets the Firepower Threat Defense configuration to factory defaults by changing manager from FMC to local and configuring the device from beginning. For more information, see Deploying the FMC Policy Configuration over VPN Tunnel, on page 2.

After you change configurations, deploy them to the affected devices. We *strongly* recommend that you deploy in a maintenance window or at a time when any interruptions to traffic flow and inspection will have the least impact.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 8 and Configurations that Restart the Snort Process When Deployed or Activated, on page 10.

Before you begin

- Review the guidelines described in Best Practices for Deploying Configuration Changes, on page 2.
- Be sure all managed devices use the same revision of the Security Zones object. If you have edited security zone objects: Do not deploy configuration changes to any device until you edit the zone setting for interfaces on *all* devices you want to sync. You must deploy to all managed devices at the same time. See Synchronizing Security Zone Object Revisions.

Note Policy deployment process fails if the sensor configuration is being read by the system during deployment. Executing commands such as `show running-config` from the sensor CLI disturbs the deployment, which results in deployment failure.

Procedure

Step 1 On the FMC menu bar, click **Deploy**.

The Deploy Policies dialog lists devices with out-of-date configurations. The **Version** at the top of the dialog specifies when you last made configuration changes. The **Current Version** column in the device table specifies when you last deployed changes to each device.

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

- Sort—Sort the device list by clicking a column heading.
- Expand—Click **Plus** to expand a device listing to view the configuration changes to be deployed. The system marks out-of-date policies with an **Index**.
- Filter—Filter the device list. Click the arrow in the upper-right corner of any column heading in the display, enter text in the **Filters** text box, and press Enter. Check or uncheck the check box to activate or deactivate the filter.
- Arrange—Place the mouse on a column heading to drag and drop the column in your preferred order.

Step 3 (Optional) Enable **Check for rule conflicts** to check for rule conflicts before deployment begins.

In policies that evaluate rules in a specific order, conflicts occur when a rule will never match traffic because an earlier rule always matches the traffic first. Although rule conflicts do not prevent deployment, properly creating and ordering rules is essential to building an effective deployment. The deployment process can check for rule conflicts in access control policies. However, in complex deployments, checking for rule conflicts can take some time.

Step 4 Click **Deploy**.

Step 5 If the system identifies errors or warnings in the changes to be deployed, it displays them in the **Errors and Warnings for Requested Deployment** window.

You have the following choices:

- Proceed—Continue deploying without resolving warning conditions. If you enabled rule conflict checks, warnings include detected rule conflicts. You cannot proceed if the system identifies errors.
- Cancel—Exit without deploying. Resolve the error and warning conditions, and attempt to deploy the configuration again.

What to do next

- (Optional) Monitor deployment status; see Viewing Deployment Messages.
- If deploy fails, see Best Practices for Deploying Configuration Changes, on page 2.

Related Topics

Snort® Restart Scenarios, on page 6

Redeploy Existing Configurations to a Device

You can force-deploy existing (unchanged) configurations to a single managed device. We *strongly* recommend you deploy in a maintenance window or at a time when any interruptions to traffic flow and inspection will have the least impact.

Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 8 and Configurations that Restart the Snort Process When Deployed or Activated, on page 10.

Before you begin

Review the guidelines described in Best Practices for Deploying Configuration Changes, on page 2.

Procedure

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

Step 2 Click **Edit** ( ) next to the device where you want to force deployment.

In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.

Step 3 Click **Device**.

Step 4 Click **Edit** ( ) next to the **General** section heading.

Step 5 Click **Force Deploy** ( ).

Step 6 Click **Deploy**.

The system identifies any errors or warnings with the configurations you are deploying. You can click **Proceed** to continue without resolving warning conditions. However, you cannot proceed if the system identifies an error.

Related Topics

Snort® Restart Scenarios, on page 6

Snort® Restart Scenarios

When the traffic inspection engine referred to as *the Snort process* on a managed device restarts, inspection is interrupted until the process resumes. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 8

for more information. Additionally, resource demands may result in a small number of packets dropping without inspection when you deploy, regardless of whether the Snort process restarts.

Any of the scenarios in the following table cause the Snort process to restart.

Table 1: Snort Restart Scenarios

Restart Scenario	More Information
Deploying a specific configuration that requires the Snort process to restart.	Configurations that Restart the Snort Process When Deployed or Activated, on page 10
Modifying a configuration that immediately restarts the Snort process.	Changes that Immediately Restart the Snort Process, on page 12
Traffic-activation of the currently deployed Automatic Application Bypass (AAB) configuration.	Configure Automatic Application Bypass

Related Topics

Access Control Policy Advanced Settings

Configurations that Restart the Snort Process When Deployed or Activated, on page 10

Inspect Traffic During Policy Apply

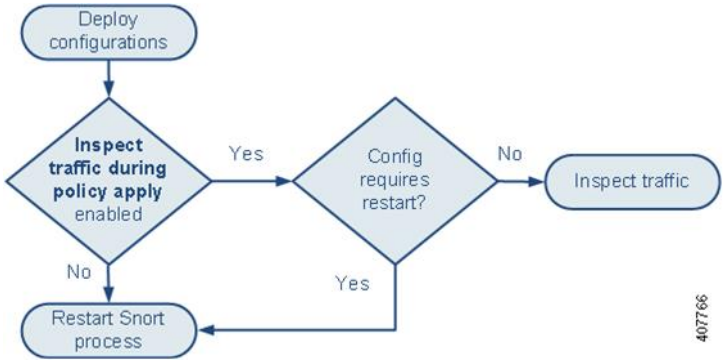
Inspect traffic during policy apply is an advanced access control policy general setting that allows managed devices to inspect traffic while deploying configuration changes; this is the case unless a configuration that you deploy requires the Snort process to restart. You can configure this option as follows:

- Enabled — Traffic is inspected during the deployment unless certain configurations require the Snort process to restart.

  When the configurations you deploy do not require a Snort restart, the system initially uses the currently deployed access control policy to inspect traffic, and switches during deployment to the access control policy you are deploying.

- Disabled — Traffic is not inspected during the deployment. The Snort process always restarts when you deploy.

The following graphic illustrates how Snort restarts can occur when you enable or disable **Inspect traffic during policy apply**.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 8 and Configurations that Restart the Snort Process When Deployed or Activated, on page 10.

Snort® Restart Traffic Behavior

The following tables explain how different devices handle traffic when the Snort process restarts.

Table 2: FTD and FTDv Restart Traffic Effects

Interface Configuration	Restart Traffic Behavior
inline: **Snort Fail Open: Down**: disabled	dropped
inline: **Snort Fail Open: Down**: enabled	passed without inspection Some packets can be delayed in buffer for several seconds before the system recognizes that Snort is down. This delay can vary depending upon the load distribution. However, the buffered packets are eventually passed.
routed, transparent (including EtherChannel, redundant, subinterface) when **preserve-connection** is enabled (**configure snort preserve-connection enable**; default) Note that both the FMC and Firepower Threat Defense must be running Version 6.2.0.2 or a subsequent 6.2.0.x patch. For more information, see Cisco Firepower Threat Defense Command Reference.	existing TCP/UDP flows: passed without inspection so long as at least one packet arrives while Snort is down new TCP/UDP flows and all non-TCP/UDP flows: dropped Note that the following traffic drops even when **preserve-connection** is enabled: • When Snort goes down before the system tags traffic that would have matched a **Do not decrypt** SSL rule or default policy action (tagging occurs after a few handshake packet exchanges) • plaintext, passthrough prefilter tunnel traffic that matches an **Analyze** rule action or an **Analyze all tunnel traffic** default policy action • connections that do not match an access control rule and are instead handled by the default action. • decrypted TLS/SSL traffic • a safe search flow • a captive portal flow

Interface Configuration	Restart Traffic Behavior
routed, transparent (including EtherChannel, redundant, subinterface) when either of the following occurs: • **preserve-connection** is disabled (**configure snort preserve-connection disable**) • the FMC or Firepower Threat Defense is not running version 6.2.0.2 or a subsequent 6.2.0.x patch	dropped
inline: tap mode	egress packet immediately, copy bypasses Snort
passive	uninterrupted, not inspected

Table 3: 7000 and 8000, NGIPSv Restart Traffic Effects

Interface Configuration	Restart Traffic Behavior
inline: **Failsafe** enabled or disabled	passed without inspection A few packets might drop if **Failsafe** is disabled and Snort is busy but not down.
inline: tap mode	egress packet immediately, copy bypasses Snort
passive	uninterrupted, not inspected
routed, switched (7000 and 8000 Series only)	dropped

Table 4: ASA FirePOWER Restart Traffic Effects

Interface Configuration	Restart Traffic Behavior
routed or transparent with fail-open	passed without inspection
routed or transparent with fail-close	dropped

Note In addition to traffic handling when the Snort process is down while it restarts, traffic can also pass without inspection or drop when the Snort process is busy, depending on the configuration of the Failsafe option (see Inline Sets) or the Snort Fail Open **Busy** option (see Configure an Inline Set). A device supports either the Failsafe option or the Snort Fail Open option, but not both.

Warning Do not reboot the system while the Snort Rule Update is in progress.

Snort-busy drops happen when snort is not able to process the packets fast enough. Lina does not know whether Snort is busy due to processing delay, or if is stuck or due to call blocking. When transmission queue is full, snort-busy drops occur. Based on Transmission queue utilization, Lina will try to access if the queue is being serviced smoothly.



Note When the Snort process is busy but not down during configuration deployment, some packets may drop on routed, switched, or transparent interfaces if the total CPU load exceeds 50 percent.

Configurations that Restart the Snort Process When Deployed or Activated

Deploying any of the following configurations except AAB restarts the Snort process as described. Deploying AAB does not cause a restart, but excessive packet latency activates the currently deployed AAB configuration, causing a partial restart of the Snort process.

Access Control Policy

- Add the first or remove the last URL category/reputation condition in an access control rule.
- Change the total number of active intrusion policies by adding an intrusion policy that is not currently used, or by removing the last instance of an intrusion policy. You can use an intrusion policy in an access control rule, as the default action, or as the default intrusion policy.

Access Control Policy Advanced Settings

- Deploy when **Inspect Traffic During Policy Apply** is disabled.
- Configure a non-default value under Files and Malware Settings.
- Add or remove an SSL policy.
- Enable or disable adaptive profiles.
- Enable or disable the **Log Session/Protocol Distribution** troubleshooting option.

Security Intelligence

Add or delete multiple Security Intelligence whitelist or blacklist networks or network objects. Changes can be to custom or system-provided lists, and whether the Snort process restarts can vary by device, depending on the memory available for inspection.

SSL Policy

- Add the first or remove the last category/reputation condition in an SSL rule.

File Policy

Deploy the first or last of any one of the following configurations; note that while otherwise deploying these file policy configurations does not cause a restart, deploying non-file-policy configurations can cause restarts.

- Enable or disable **Inspect Archives**.
- Take either of the following actions:

  - Enable or disable **Inspect Archives** when the deployed access control policy includes at least one file policy.
  - Add the first or remove the last file policy rule when **Inspect Archives** is enabled (note that at least one rule is required for **Inspect Archives** to be meaningful).

- Select **Detect Files** or **Block Files** in a file rule.
- Enable or disable **Store files** in a **Detect Files** or **Block Files** rule.
- Add the first or remove the last active file rule that combines the **Malware Cloud Lookup** or **Block Malware** rule action with an analysis option (**Spero Analysis or MSEXE**, **Dynamic Analysis**, or **Local Malware Analysis**) or a store files option (**Malware, Unknown**, **Clean**, or **Custom**).

Note that access control rules that deploy these file policy configurations to security zones or tunnel zones cause a restart only when your configuration meets the following conditions:

- Source or destination security zones in your access control rule must match the security zones associated with interfaces on the target devices.
- Unless the destination zone in you access control rule is *any*, a source tunnel zone in the rule must match a tunnel zone assigned to a tunnel rule in the prefilter policy.

Identity Policy

- When SSL decryption is disabled (that is, when the access control policy does not include an SSL policy), add the first or remove the last active authentication rule.

  An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive authentication cannot identify user** selected.

Network Analysis Policy

- Change the total number of network analysis policies by adding a network analysis policy that is not currently used, or by removing the last instance of a network analysis policy. You can use a network analysis policy with network analysis rules or as the default network analysis policy.
- Change the value for the IMAP, POP, or SMTP preprocessor **Base64 Decoding Depth**, **7-Bit/8-Bit/Binary Decoding Depth**, **Quoted-Printable Decoding Depth**, or **Unix-to-Unix Decoding Depth**.

Network Discovery

- Enable or disable non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy.

Device Management

- Routing: Add a routed interface pair or virtual router to a 7000 or 8000 Series device.
- VPN: Add or remove a VPN on a 7000 or 8000 Series device.

Caution The system does not warn you that the Snort process restarts when you add or remove a VPN on a 7000 or 8000 Series device.

- MTU: Change the highest MTU value among all non-management interfaces on a device.
- 7000/8000 series high availability: Change a high-availability state sharing option.
- Automatic Application Bypass (AAB): The currently deployed AAB configuration activates when a malfunction of the Snort process or a device misconfiguration causes a single packet to use an excessive amount of processing time. The result is a partial restart of the Snort process to alleviate extremely high latency or prevent a complete traffic stall. This partial restart causes a few packets to pass without inspection, or drop, depending on how the device handles traffic.

Updates

- System update: Deploy configurations the first time after a software update that includes a new version of the Snort binary or data acquisition library (DAQ).
- VDB: Deploy configurations the first time after installing a vulnerability database (VDB) .
- Intrusion rule update: Deploying configurations the first time after importing an intrusion rule update (also known as a *Snort Rule Update* or *SRU*).

Caution Intrusion rule updates are cumulative. Any shared object rule that is added or modified since your last update causes a restart when you deploy, even if the current update has no shared object rule changes.

Related Topics

Deploy Configuration Changes, on page 4
Snort® Restart Scenarios, on page 6

Changes that Immediately Restart the Snort Process

The following changes immediately restart the Snort process without going through the deploy process. How the restart affects traffic depends on how the target device handles traffic. See Snort® Restart Traffic Behavior, on page 8 for more information.

- Take any of the following actions involving applications or application detectors:
  - Activate or deactivate a system or custom application detector.
  - Delete an activated custom detector.
  - **Save and Reactivate** an activated custom detector.
  - Create a user-defined application.

  The Snort process restarts on all managed devices.
- Create or break a Firepower Threat Defense high availability pair—Restarts the Snort process on the primary and secondary devices.

- Install a vulnerability database (VDB) update.
- Restart the Snort process in the 7000 or 8000 Series user interface (**System > Configuration > Process**)—The system prompts you for confirmation and allows you to cancel.

Policy Comparison

To review policy changes for compliance with your organization's standards or to optimize system performance, you can examine the differences between two policies or between a saved policy and the running configuration.

You can compare the following policy types:

- DNS
- File
- Health
- Identity
- Intrusion (Only Snort 2 policies)
- Network Analysis
- SSL

The comparison view displays both policies in a side-by-side format. Differences between the two policies are highlighted:

- Blue indicates that the highlighted setting is different in the two policies, and the difference is noted in red text.
- Green indicates that the highlighted setting appears in one policy but not the other.

Comparing Policies

You can compare policies only if you have access rights and any required licenses for the specific policy, and you are in the correct domain for configuring the policy.

Procedure

Step 1 Access the management page for the policy you want to compare:

- DNS—**Policies** > **Access Control** > **DNS**
- File—**Policies** > **Access Control** > **Malware & File**
- Health—**System** > **Health** > **Policy**
- Identity—**Policies** > **Access Control** > **Identity**
- Intrusion—**Policies** > **Access Control** > **Intrusion**

Note You can compare only Snort 2 policies.

- Network Analysis—**Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**

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

- SSL—**Policies** > **Access Control** > **SSL**

Step 2 Click **Compare Policies**.

Step 3 From the **Compare Against** drop-down list, choose the type of comparison you want to make:

- To compare two different policies, choose **Other Policy**.
- To compare two revisions of the same policy, choose **Other Revision**.
- To compare another policy to the currently active policy, choose **Running Configuration**.

Step 4 Depending on the comparison type you choose, you have the following choices:

- If you are comparing two different policies, choose the policies you want to compare from the **Policy A** and **Policy B** drop-down lists.
- If you are comparing the running configuration to another policy, choose the second policy from the **Policy B** drop-down list.

Step 5 Click **OK**.

Step 6 Review the comparison results:

- Comparison Viewer—To use the comparison viewer to navigate individually through policy differences, click **Previous** or **Next** above the title bar.
- Comparison Report—To generate a PDF report that lists the differences between the two policies, click **Comparison Report**.

Policy Reports

For most policies, you can generate two kinds of reports. A report on a single policy provides details on the policy's current saved configuration, while a comparison report lists only the differences between two policies. You can generate a single-policy report for all policy types except health.

Note Intrusion policy reports combine the settings in the base policy with the settings of the policy layers, and make no distinction between which settings originated in the base policy or policy layer.

Generating Current Policy Reports

You can generate policy reports only if you have access rights and any required licenses for the specific policy, and you are in the correct domain for configuring the policy.

Procedure

Step 1 Access the management page for the policy for which you want to generate a report:

- Access Control—**Policies** > **Access Control**
- DNS—**Policies** > **Access Control** > **DNS**
- File—**Policies** > **Access Control** > **Malware & File**
- Health—**System** > **Health** > **Policy**
- Identity—**Policies** > **Access Control** > **Identity**
- Intrusion—**Policies** > **Access Control** > **Intrusion**
- NAT for 7000 & 8000 Series devices—**Devices** > **NAT**
- Network Analysis—**Policies** > **Access Control**, then click **Network Analysis Policies** or **Policies** > **Access Control** > **Intrusion**, then click **Network Analysis Policies**

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

- SSL—**Policies** > **Access Control** > **SSL**

Step 2 Click **Report** () next to the policy for which you want to generate a report.

Out-of-Date Policies

The Firepower System marks out-of-date policies with red status text that indicates how many of its targeted devices need a policy update. To clear this status, you must re-deploy the policy to the devices.

Configuration changes that require a policy re-deploy include:

- Modifying an access control policy: any changes to access control rules, the default action, policy targets, Security Intelligence filtering, advanced options including preprocessing, and so on.
- Modifying any of the policies that the access control policy invokes: the SSL policy, network analysis policies, intrusion policies, file policies, identity policies, or DNS policies.
- Changing any reusable object or configuration used in an access control policy or policies it invokes:
  - network, port, VLAN tag, URL, and geolocation objects
  - Security Intelligence lists and feeds
  - application filters or detectors
  - intrusion policy variable sets
  - file lists
  - decryption-related objects and security zones
- Updating the system software, intrusion rules, or the vulnerability database (VDB).

Keep in mind that you can change some of these configurations from multiple places in the web interface. For example, you can modify security zones using the object manager (**Objects** > **Object Management**), but modifying an interface type in a device's configuration (**Devices** > **Device Management**) can also change a zone and require a policy re-deploy.

Note that the following updates do **not** require policy re-deploy:

- automatic updates to Security Intelligence feeds and additions to the Security Intelligence global Block or Do Not Block list using the context menu
- automatic updates to URL filtering data
- scheduled geolocation database (GeoDB) updates

Performance Considerations for Limited Deployments

Host, application, and user discovery data allow the system to create a complete, up-to-the-minute profile of your network. The system can also act as an intrusion detection and prevention system (IPS), analyzing network traffic for intrusions and exploits and, optionally, dropping offending packets.

Combining discovery and IPS gives context to your network activity and allows you to take advantage of many features, including:

- impact flags and indications of compromise, which can tell you which of your hosts are vulnerable to a particular exploit, attack, or piece of malware
- adaptive profile updates and Firepower recommendations, which allow you to examine traffic differently depending on the destination host
- correlation, which allows you to respond to intrusions (and other events) differently depending on the affected host

However, if your organization is interested in performing only IPS, or only discovery, there are a few configurations that can optimize the performance of the system.

Discovery Without Intrusion Prevention

The *discovery* feature allows you to monitor network traffic and determine the number and types of hosts (including network devices) on your network, as well as the operating systems, active applications, and open ports on those hosts. You can also configure managed devices to monitor user activity on your network. You can use discovery data to perform traffic profiling, assess network compliance, and respond to policy violations.

In a basic deployment (discovery and simple, network-based access control only), you can improve a device's performance by following a few important guidelines when configuring its access control policy.



Note You must use an access control policy, even if it simply allows all traffic. The network discovery policy can **only** examine traffic that the access control policy allows to pass.

First, make sure your access control policy does not require complex processing and uses only simple, network-based criteria to handle network traffic. You must implement **all** of the following guidelines; misconfiguring any one of these options eliminates the performance benefit:

- Do **not** use the Security Intelligence feature. Remove any populated global Block or Do Not Block list from the policy's Security Intelligence configuration.
- Do **not** include access control rules with Monitor or Interactive Block actions. Use only Allow, Trust, and Block rules. Keep in mind that allowed traffic can be inspected by discovery; trusted and blocked traffic cannot.
- Do **not** include access control rules with application, user, URL, ISE attribute, or geolocation-based network conditions. Use only simple network-based conditions: zone, IP address, VLAN tag, and port.
- Do **not** include access control rules that perform file, malware, or intrusion inspection. In other words, do not associate a file policy or intrusion policy with any access control rule.
- In the Advanced settings for the access control policy, make sure that **Intrusion Policy used before Access Control rule is determined** is set to **No Rules Active**.
- Select **Network Discovery Only** as the policy's default action. Do **not** choose a default action for the policy that performs intrusion inspection.

In conjunction with the access control policy, you can configure and deploy the network discovery policy, which specifies the network segments, ports, and zones that the system examines for discovery data, as well as whether hosts, applications, and users are discovered on the segments, ports, and zones.

Related Topics

Inspection of Packets That Pass Before Traffic Is Identified

Intrusion Prevention Without Discovery

Disabling discovery if you don't need it (for example, in an IPS-only deployment) can improve performance. To disable discovery you must implement *all* of these changes:

- Delete *all* rules from your network discovery policy.
- Use *only* simple network-based conditions to perform access control: zone, IP address, VLAN tag, and port.

  Do *not* perform any kind of, application, user, URL, or geolocation control. Although you can disable storage of discovery data, the system still must collect and examine it to implement those features.

After you deploy, new discovery halts on target devices. The system gradually deletes information in the network map according to your timeout preferences. Or, you can purge all discovery data immediately.