Best Practices: Use Cases for FTD

The following topics explain some common tasks you might want to accomplish with FTD using the FDM. These use cases assume that you completed the device configuration wizard and that you retained this initial configuration. Even if you modified the initial configuration, you should be able to use these examples to understand how to use the product.

How to Configure the Device in FDM

After you complete the setup wizard, you should have a functioning device with a few basic policies in place:

- (Except for ASA 5506-X and.) An outside and an inside interface. No other data interfaces are configured.
- (ASA 5506-X, only.) An outside interface, and an inside bridge group that includes all other data interfaces.
- Security zones for the inside and outside interfaces.
- An access rule trusting all inside to outside traffic.
- An interface NAT rule that translates all inside to outside traffic to unique ports on the IP address of the outside interface.
- A DHCP server running on the inside interface or bridge group.

The following steps provide an overview of additional features you might want to configure. Please click the help button (**?**) on a page to get detailed information about each step.

Procedure

Step 1 Choose **Device**, then click **View Configuration** in the **Smart License** group.

Click **Enable** for each of the optional licenses you want to use: Threat, Malware, URL. If you registered the device during setup, you can also enable the RA VPN license desired. Read the explanation of each license if you are unsure of whether you need it.

If you have not registered, you can do so from this page. Click **Request Register** and follow the instructions. Please register before the evaluation license expires.

For example, an enabled Threat license should look like the following:

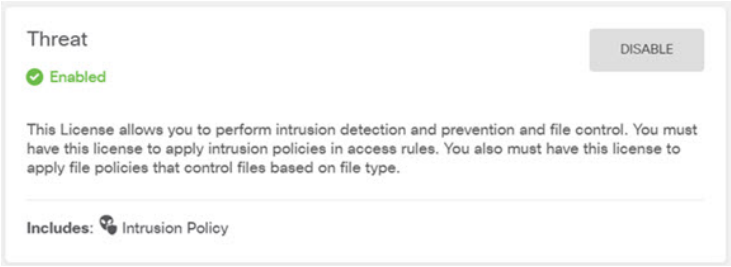

Step 2 If you wired other interfaces, choose **Device**, then click the link in the **Interfaces** summary.

- Because the ASA 5506-X and comes pre-configured with a bridge group containing all non-outside data interfaces, there is no need to configure these interfaces. If you want to break apart the bridge group, you can edit it to remove the interfaces you want to treat separately. Then you can configure those interfaces as hosting separate networks.

  For other models, you can create a bridge group for the other interfaces, or configure separate networks, or some combination of both.

Click the edit icon ( ) for each interface to define the IP address and other settings.

The following example configures an interface to be used as a “demilitarized zone” (DMZ), where you place publically-accessible assets such as your web server. Click **Save** when you are finished.

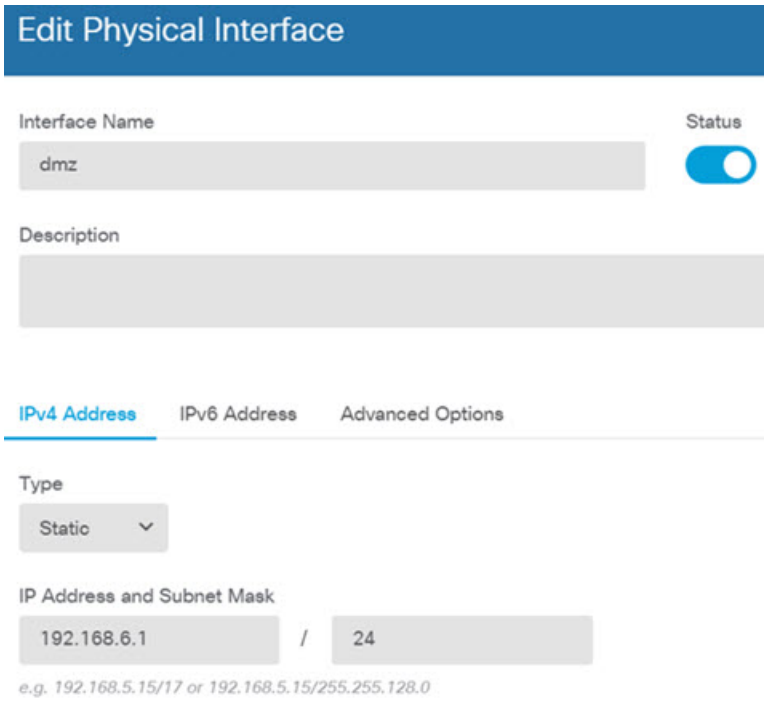

Step 3 If you configured new interfaces, choose **Objects**, then select **Security Zones** from the table of contents.

Edit or create new zones as appropriate. Each interface must belong to a zone, because you configure policies based on security zones, not interfaces. You cannot put the interfaces in zones when configuring them, so you must always edit the zone objects after creating new interfaces or changing the purpose of existing interfaces.

The following example shows how to create a new dmz-zone for the dmz interface.

Step 4 If you want internal clients to use DHCP to obtain an IP address from the device, choose **Device**, then **System Settings** > **DHCP Server**. Select the **DHCP Servers** tab.

There is already a DHCP server configured for the inside interface, but you can edit the address pool or even delete it. If you configured other inside interfaces, it is very typical to set up a DHCP server on those interfaces. Click + to configure the server and address pool for each inside interface.

You can also fine-tune the WINS and DNS list supplied to clients on the **Configuration** tab.

The following example shows how to set up a DHCP server on the inside2 interface with the address pool 192.168.4.50-192.168.4.240.

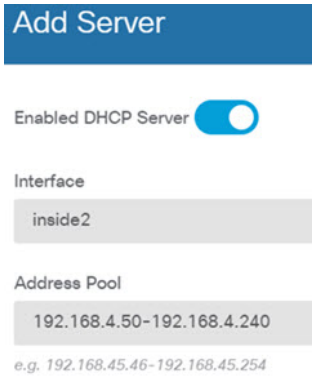

Step 5 Choose **Device**, then click **View Configuration** (or **Create First Static Route**) in the **Routing** group and configure a default route.

The default route normally points to the upstream or ISP router that resides off the outside interface. A default IPv4 route is for any-ipv4 (0.0.0.0/0), whereas a default IPv6 route is for any-ipv6 (::0/0). Create routes for each IP version you use. If you use DHCP to obtain an address for the outside interface, you might already have the default routes that you need.

The routes you define on this page are for the data interfaces only. They do not impact the management interface. Set the management gateway on **System Settings** > **Management Interface**.

The following example shows a default route for IPv4. In this example, isp-gateway is a network object that identifies the IP address of the ISP gateway (you must obtain the address from your ISP). You can create this object by clicking **Create New Network** at the bottom of the **Gateway** drop-down list.

Step 6 Choose **Policies** and configure the security policies for the network.

The device setup wizard enables traffic flow between the inside-zone and outside-zone, and interface NAT for all interfaces when going to the outside interface. Even if you configure new interfaces, if you add them to the inside-zone object, the access control rule automatically applies to them.

However, if you have multiple inside interfaces, you need an access control rule to allow traffic flow from inside-zone to inside-zone. If you add other security zones, you need rules to allow traffic to and from those zones. These would be your minimum changes.

In addition, you can configure other policies to provide additional services, and fine-tune NAT and access rules to get the results that your organization requires. You can configure the following policies:

- **Identity**—If you want to correlate network activity to individual users, or control network access based on user or user group membership, use the identity policy to determine the user associated with a given source IP address.
- **NAT** (Network Address Translation)—Use the NAT policy to convert internal IP addresses to externally routeable addresses.
- **Access Control**—Use the access control policy to determine which connections are allowed on the network. You can filter by security zone, IP address, protocol, port, application, URL, user or user group. You also apply intrusion and file (malware) policies using access control rules. Use this policy to implement URL filtering.

The following example shows how to allow traffic between the inside-zone and dmz-zone in the access control policy. In this example, no options are set on any of the other tabs except for **Logging**, where **At End of Connection** is selected.

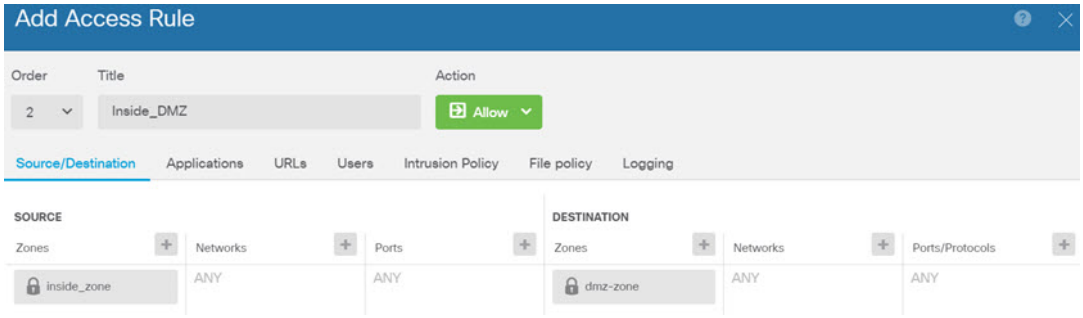

Step 7 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.

b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

How to Gain Insight Into Your Network Traffic

After completing initial device setup, you have an access control policy that allows all inside traffic access to the Internet or other upstream network, and a default action to block all other traffic. Before you create additional access control rules, you might find it beneficial to gain insight into the traffic that is actually occurring on your network.

You can use the monitoring capabilities of the FDM to analyze network traffic. FDM reporting helps you answer the following questions:

- What is my network being used for?
- Who is using the network the most?
- Where are my users going?
- What devices are they using?
- What access control rules (policies) are being hit the most?

The initial access rule can provide some insight into traffic, including policies, destinations, and security zones. But to obtain user information, you need to configure an identity policy that requires users to authenticate (identify) themselves. To obtain information on applications used on the network, you need to make some additional adjustments.

The following procedure explains how to set up the FTD device to monitor traffic and provides an overview of the end-to-end process of configuring and monitoring policies.

Note This procedure does not provide insight into the web site categories and reputations of sites visited by users, so you cannot see meaningful information in the web categories dashboard. You must implement category-based URL filtering, and enable the URL license, to obtain category and reputation data. If you just want to obtain this information, you can add a new access control rule that allows access to an acceptable category, such as Financial Services, and make it the first rule in the access control policy. For details on implementing URL filtering, see How to Implement an Acceptable Use Policy (URL Filtering), on page 19.

Procedure

Step 1 To gain insight into user behavior, you need to configure an identity policy to ensure that the user associated with a connection is identified.

By enabling the identity policy, you can collect information about who is using the network, and what resources they are using. This information is available in the User monitoring dashboard. User information is also available for connection events shown in Event Viewer.

Users are authenticated only when they use a web browser for HTTP connections.

If a user fails to authenticate, the user is not prevented from making web connections. This just means that you do not have user identity information for the connections. If you want, you can create an access control rule to drop traffic for Failed Authentication users.

a) Click **Policies** in the main menu, then click **Identity**.

The identity policy is initially disabled. The identity policy uses your Active Directory server to authenticate users and associate them with the IP address of the workstation they are using. Subsequently, the system will identify traffic for that IP address as being the user's traffic.

b) Click **Enable Identity Policy**.

This action opens the Identity Policy Configuration dialog box.

c) Click in **Realm Server** to open the drop-down list, then select **Create New Identity Realm**.

If you already created your realm server object, simply select it and skip the steps for configuring the server.

d) Fill in the following fields, then click **OK**.

- **Name**—A name for the directory realm.
- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.
- **Directory Username**, **Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

  Note The system generates ldap-login-dn and ldap-login-password from this information. For example, Administrator@example.com is translated as cn=adminisntrator,cn=users,dc=example,dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name “users” folder.

- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, dc=example,dc=com. For information on finding the base DN, see Determining the Directory Base DN.
- **AD Primary Domain**— The fully qualified Active Directory domain name that the device should join. For example, example.com.
- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address.
- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method.
- **Encryption**—To use an encrypted connection for downloading user and group information, select the desired method, **STARTTLS** or **LDAPS**. The default is **None**, which means that user and group information is downloaded in clear text.
  - **STARTTLS** negotiates the encryption method, and uses the strongest method supported by the directory server. Use port 389. This option is not supported if you use the realm for remote access VPN.
  - **LDAPS** requires LDAP over SSL. Use port 636.
- **Trusted CA Certificate**—If you select an encryption method, upload a Certificate Authority (CA) certificate to enable a trusted connection between the system and the directory server. If you are

using a certificate to authenticate, the name of the server in the certificate must match the server Hostname / IP Address. For example, if you use 10.10.10.250 as the IP address but ad.example.com in the certificate, the connection fails.

Example:

For example, the following image shows how to create an unencrypted connection for the ad.example.com server. The primary domain is example.com, and the directory username is Administrator@ad.example.com. All user and group information is under the Distinguished Name (DN) ou=user,dc=example,dc=com.

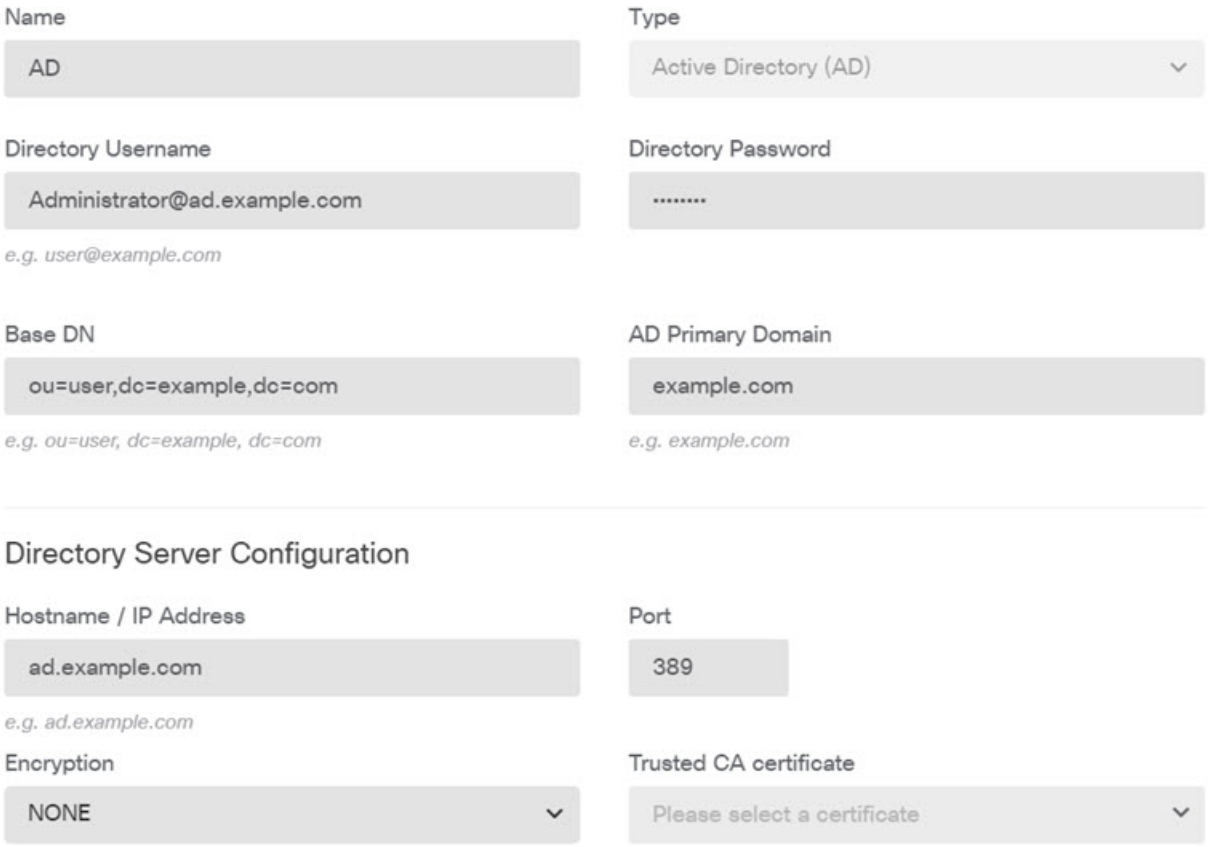

e) In the Identity Policy Configuration dialog box, **Realm Server** list, select the realm server you just created.

f) In the Identity Policy Configuration dialog box, configure the Active Authentication captive portal settings.

When an identity rule requires active authentication for a user, the user is redirected to the captive portal port on the interface through which they are connected and then they are prompted to authenticate.

- **Server Certificate**—Select the internal certificate to present to users during active authentication. You can select the predefined self-signed DefaultInternalCertificate, or you can click **Create New Internal Certificate** and upload a certificate that your browsers already trust.

  Users will have to accept the certificate if you do not upload a certificate that their browsers already trust.

- **Port**—The captive portal port. The default is 885 (TCP). If you configure a different port, it must be in the range 1025-65535.

Example:

The Identity Policy Configuration dialog box should now look like the following.

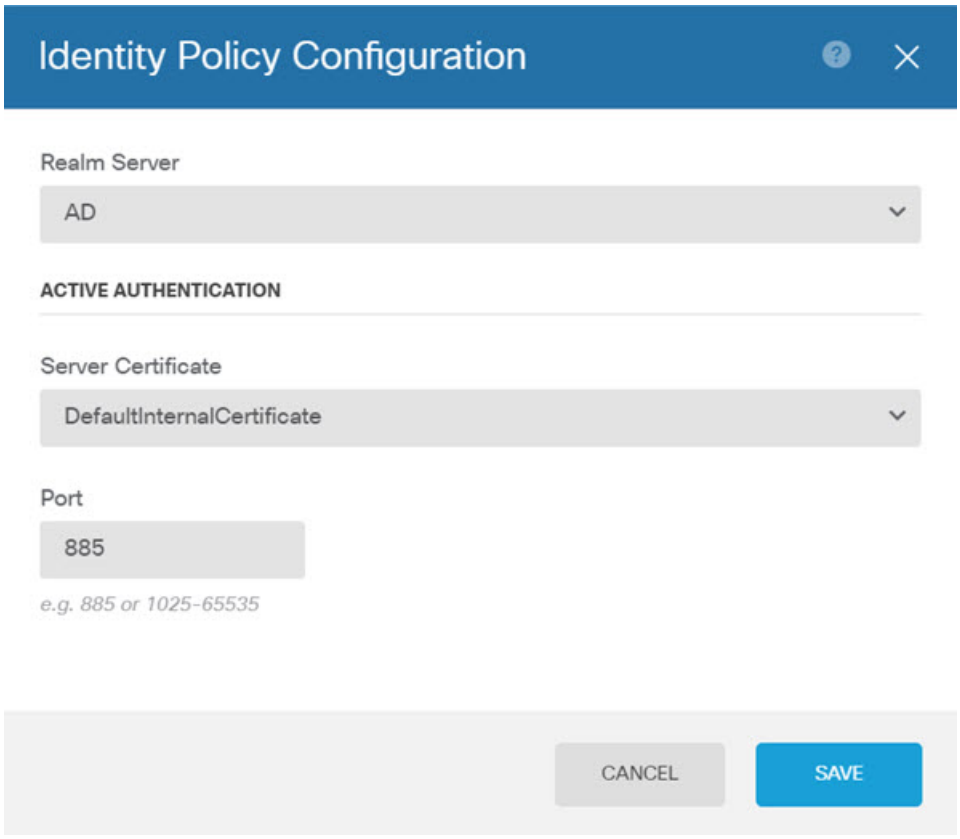

g) Click **Save**.

Now, create a rule to require active authentication.

h) Click the **Create Identity Rule** button, or the + button.

i) Fill in the identity rule properties.

Assuming you want to require everyone to authenticate, you could use the following settings:

- **Name**—Anything you choose, for example, Require_Authentication.
- **User Authentication**—**Active** should already be selected; keep it.
- **Type**—Select **HTTP Negotiate**. This allows the browser and directory server to negotiate the strongest authentication protocol, in order, NTLM, then HTTP basic.

  Note For the HTTP Basic, HTTP Response Page, and NTLM authentication methods, the user is redirected to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate. If you cannot, or do not want to, update the DNS server, select one of the other authentication methods.

- **Source/Destination**—Leave all fields to default to Any.

You can constrain the policy as you see fit to a more limited set of traffic. However, active authentication will only be attempted for HTTP traffic, so it does not matter that non-HTTP traffic matches the source/destination criteria. For more details about identity policy properties, see Configure Identity Rules.

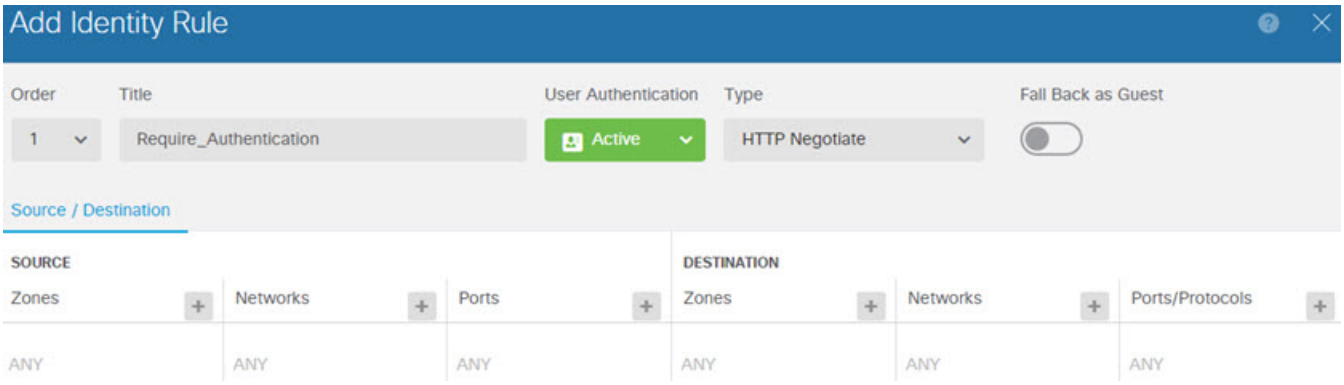

j) Click **OK** to add the rule.

If you look in the upper right of the window, you can see that the **Deploy** icon button now has a dot, which indicates that there are undeployed changes. Making changes in the user interface is not sufficient for getting the changes configured on the device, you must deploy changes. Thus, you can make a set of related changes before you deploy them, so that you do not face the potential problems of having a partially-configured set of changes running on the device. You will deploy changes later in this procedure.

Step 2 Change the action on the Inside_Outside_Rule access control rule to **Allow**.

The Inside_Outside_Rule access rule is created as a trust rule. However, trusted traffic is not inspected, so the system cannot learn about some of the characteristics of trusted traffic, such as application, when the traffic matching criteria does not include application or other conditions besides zone, IP address, and port. If you change the rule to allow rather than trust traffic, the system fully inspects the traffic.

Note (ASA 5506-X .) Also consider changing the Inside_Inside_Rule from Trust to Allow. This rule covers traffic going between the inside interfaces.

a) Click **Access Control** on the **Policies** page.
b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon ( ) to open the rule.
c) Select **Allow** for the **Action**.

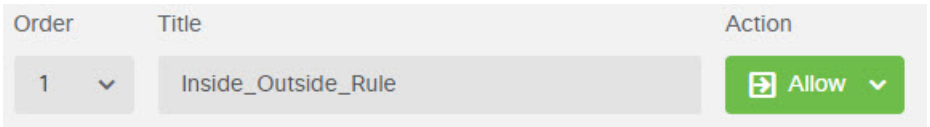

d) Click **OK** to save the change.

Step 3 Enable logging on the access control policy default action.

Dashboards contain information about connections only if the connection matches an access control rule that enables connection logging. The Inside_Outside_Rule enables logging, but the default action has logging disabled. Thus, dashboards show information for the Inside_Outside_Rule only, and do not reflect connections that do not match any rules.

a) Click anywhere in the default action at the bottom of the access control policy page.

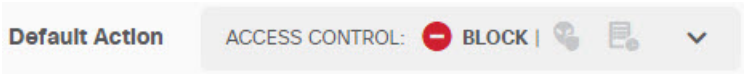

b) Select **Select Log Action** > **At Beginning and End of Connection**.
c) Click **OK**.

Step 4 Set an update schedule for the vulnerability database (VDB).

Cisco regularly releases updates to the VDB, which includes the application detectors that can identify the application used in a connection. You should update the VDB on a regular basis. You can either manually download updates, or you can set up a regular schedule. The following procedure shows how to set up a schedule. By default, VDB updates are disabled, so you need to take action to get VDB updates.

a) Click **Device**.
b) Click **View Configuration** in the Updates group.

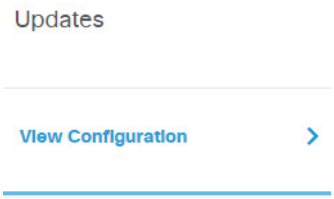

c) Click **Configure** in the VDB group.

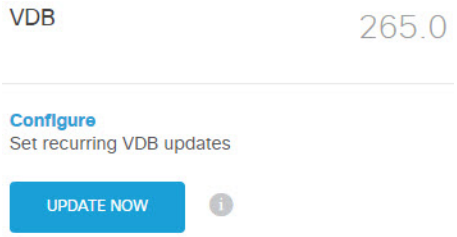

d) Define the update schedule.

Choose a time and frequency that will not be disruptive to your network. Also, please understand that the system will do an automatic deployment after downloading the update. This is necessary to activate the new detectors. Thus, any configuration changes that you have made and saved but have not yet deployed will also be deployed.

For example, the following schedule updates the VDB once a week on Sunday at 12:00 AM (using the 24-hour clock notation).

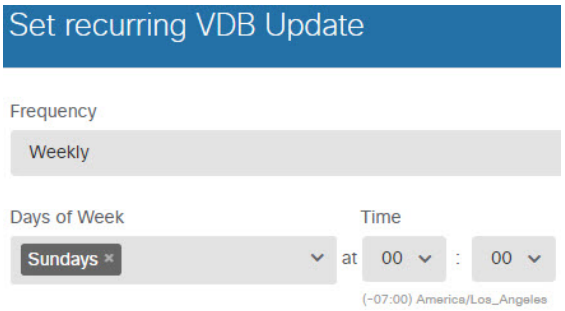

e) Click **Save**.

Step 5 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.

b) Click the **Deploy Now** button and wait for deployment to finish.

The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

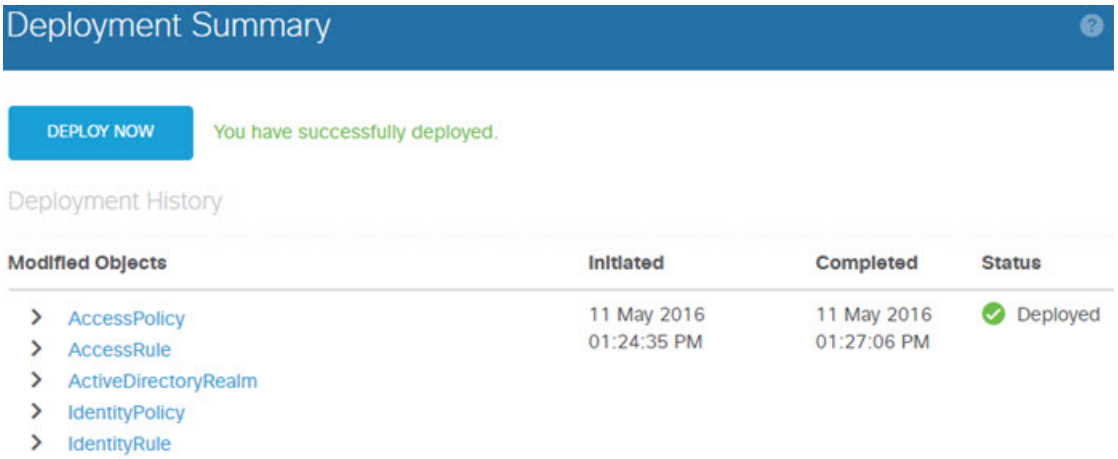

What to do next

At this point, the monitoring dashboards and events should start showing information about users and applications. You can evaluate this information for undesirable patterns and develop new access rules to constrain unacceptable use.

If you want to start collecting information about intrusions and malware, you need to enable intrusion and file policies on one or more access rule. You also need to enable the licenses for those features.

If you want to start collecting information about URL categories, you must implement URL filtering.

How to Block Threats

You can implement next generation Intrusion Prevention System (IPS) filtering by adding intrusion policies to your access control rules. Intrusion policies analyze network traffic, comparing the traffic contents against known threats. If a connection matches a threat you are monitoring, the system drops the connection, thus preventing the attack.

All other traffic handling occurs before network traffic is examined for intrusions. By associating an intrusion policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect the traffic with an intrusion policy.

You can configure intrusion policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, you can configure an intrusion policy as part of the default action if the default action is **allow**.

The intrusion policies are designed by the Cisco Talos Intelligence Group (Talos), who set the intrusion and preprocessor rule states and advanced settings.

Procedure

Step 1 If you have not already done so, enable the Threat license.

You must enable the Threat license to use intrusion policies. If you are currently using the evaluation license, you are enabling an evaluation version of the license. If you have registered the device, you must purchase the required license and add it to your Smart Software Manager account on Cisco.com.

a) Click **Device**.
b) Click **View Configuration** in the Smart License group.

c) Click **Enable** in the **Threat** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.

Step 2 Select an intrusion policy for one or more access rules.

Determine which rules cover traffic that should be scanned for threats. For this example, we will add intrusion inspection to the Inside_Outside_Rule. For ASA 5506-X models, you might also want to add it to the Inside_Inside_Rule.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon ( ) to open the rule.

c) If you have not already done so, select **Allow** for the **Action**.

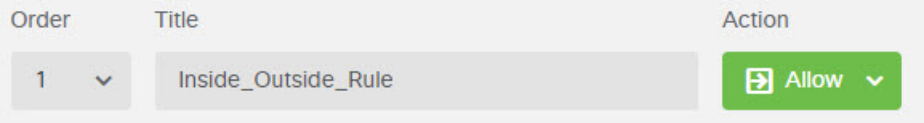

d) Click the **Intrusion Policy** tab.

e) Click the **Intrusion Policy** toggle to enable it, then select the intrusion policy.

The **Balanced Security and Connectivity** policy is appropriate for most networks. It provides a good intrusion defense without being overly aggressive, which has the potential of dropping traffic that you might not want to be dropped. If you determine that too much traffic is getting dropped, you can ease up on intrusion inspection by selecting the **Connectivity over Security** policy.

If you need to be aggressive about security, try the **Security over Connectivity** policy. The **Maximum Detection** policy offers even more emphasis on network infrastructure security with the potential for even greater operational impact.

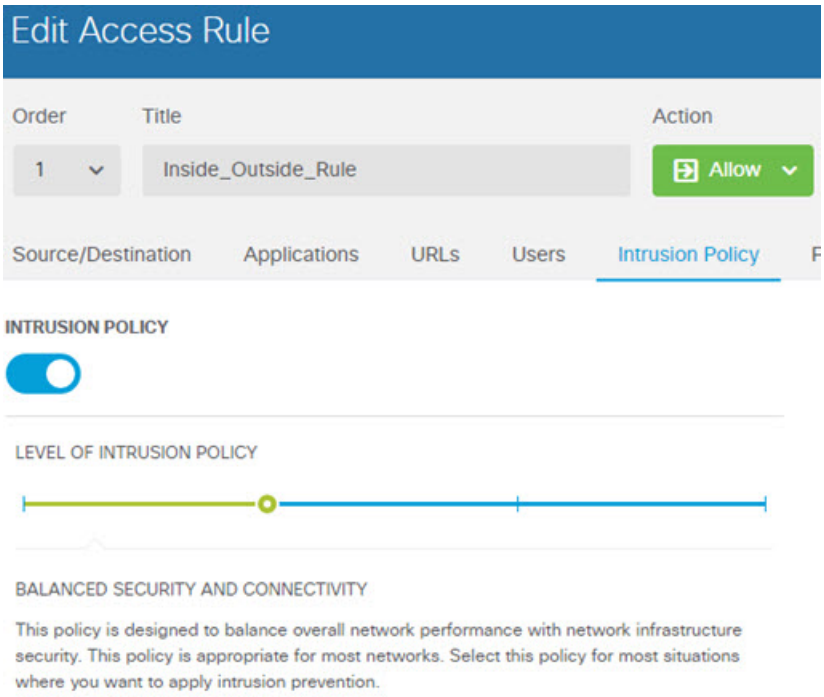

f) Click **OK** to save the change.

Step 3 Set an update schedule for the intrusion rule database.

Cisco regularly releases updates to the intrusion rule database, which is used by intrusion policies to determine whether connections should be dropped. You should update the rule database on a regular basis. You can

either manually download updates, or you can set up a regular schedule. The following procedure shows how to set up a schedule. By default, database updates are disabled, so you need to take action to get updated rules.

a) Click **Device**.
b) Click **View Configuration** in the Updates group.

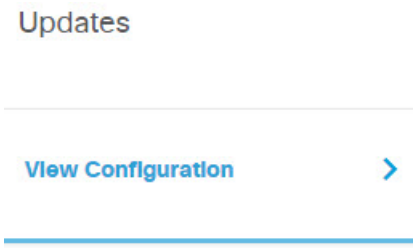

c) Click **Configure** in the Rule group.

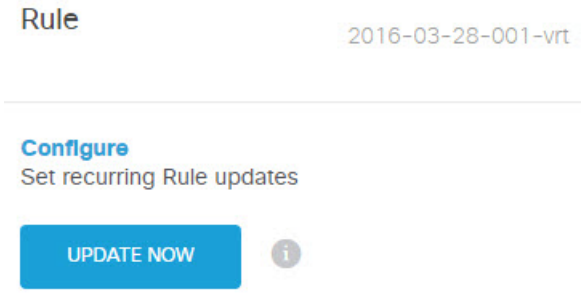

d) Define the update schedule.

Choose a time and frequency that will not be disruptive to your network. Also, please understand that the system will do an automatic deployment after downloading the update. This is necessary to activate the new rules. Thus, any configuration changes that you have made and saved but have not yet deployed will also be deployed.

For example, the following schedule updates the rule database once a week on Monday at 12:00 AM (using the 24-hour clock notation).

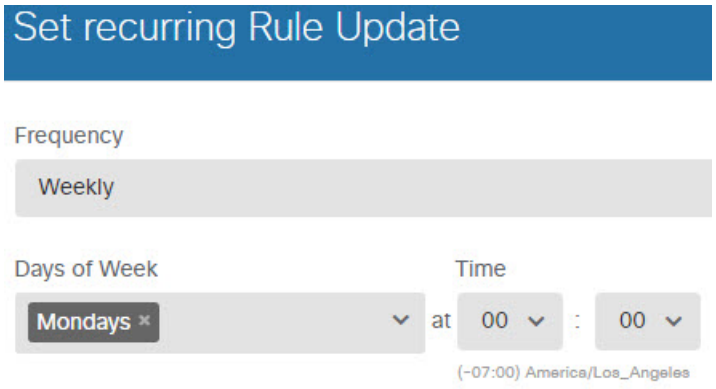

e) Click **Save**.

Step 4 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.

b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

At this point, the monitoring dashboards and events should start showing information about attackers, targets, and threats, if any intrusions are identified. You can evaluate this information to determine if your network needs more security precautions, or if you need to reduce the level of intrusion policy you are using.

How to Block Malware

Users are continually at risk of obtaining malicious software, or *malware*, from Internet sites or other communication methods, such as e-mail. Even trusted web sites can be hijacked to serve malware to unsuspecting users. Web pages can contain objects coming from different sources. These objects can include images, executables, Javascript, advertisements, and so forth. Compromised web sites often incorporate objects hosted on external sources. Real security means looking at each object individually, not just the initial request.

Use file policies to detect malware using malware defense. You can also use file policies to perform file control, which allows control over all files of a specific type regardless of whether the files contain malware.

Malware defense uses the AMP Cloud to retrieve dispositions for possible malware detected in network traffic. The management interface must have a path to the Internet to reach the AMP Cloud and perform malware lookups. When the device detects an eligible file, it uses the file's SHA-256 hash value to query the AMP Cloud for the file's disposition. The possible disposition can be **clean**, **malware**, or **unknown** (no clear verdict). If the AMP Cloud is unreachable, the disposition is **unknown**.

By associating a file policy with an access control rule, you are telling the system that before it passes traffic that matches the access control rule's conditions, you first want to inspect any files in the connection.

You can configure file policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic.

Procedure

Step 1 If you have not already done so, enable the Malware and Threat licenses.

You must enable the Malware to use file policies in addition to the Threat license, which is required for intrusion policies. If you are currently using the evaluation license, you are enabling an evaluation version of the licenses. If you have registered the device, you must purchase the required licenses and add them to your Smart Software Manager account on Cisco.com.

a) Click **Device**.
b) Click **View Configuration** in the Smart License group.

c) Click **Enable** in the **Malware** group, and if not already enabled, the **Threat** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.

Step 2 Select a file policy for one or more access rules.

Determine which rules cover traffic that should be scanned for malware. For this example, we will add file inspection to the Inside_Outside_Rule. For ASA 5506-X models, you might also want to add it to the Inside_Inside_Rule.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Hover over the **Actions** cell on the right side of the Inside_Outside_Rule row to expose the edit and delete icons, and click the edit icon ( ) to open the rule.

c) If you have not already done so, select **Allow** for the **Action**.

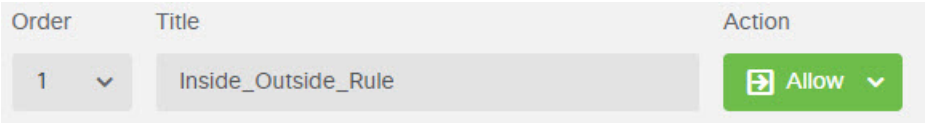

d) Click the **File Policy** tab.

e) Click the file policy you want to use.

Your main choice is between **Block Malware All**, which drops any files that are considered malware, or **Cloud Lookup All**, which queries the AMP Cloud to determine the file's disposition, but does no blocking. If you want to first see how files are being evaluated, use cloud lookup. You can switch to the blocking policy later if you are satisfied with how files are being evaluated.

There are other policies available that block malware. These policies are coupled with file control, blocking the upload of Microsoft Office, or Office and PDF, documents. That is, these policies prevent users from sending these file types to other networks in addition to blocking malware. You can select these policies if they fit your needs.

For this example, select **Block Malware All**.

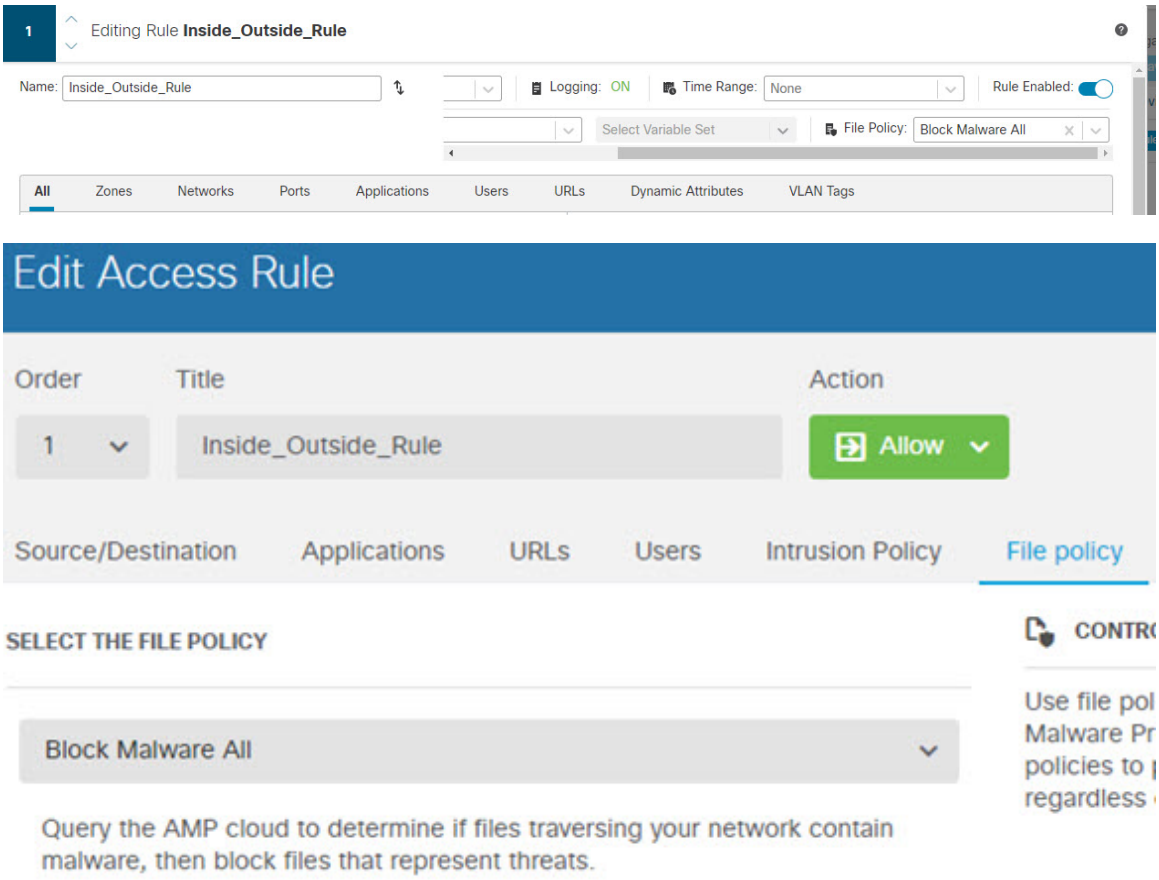

f) Click the **Logging** tab and verify that **Log Files** under File Events is selected.

By default, file logging is enabled whenever you select a file policy. You must enable file logging to get file and malware information in events and dashboards.

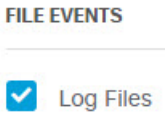

g) Click **OK** to save the change.

Step 3 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.

b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

At this point, the monitoring dashboards and events should start showing information about file types and file and malware events, if any files or malware are transmitted. You can evaluate this information to determine if your network needs more security precautions related to file transmissions.

How to Implement an Acceptable Use Policy (URL Filtering)

You might have an acceptable use policy for your network. Acceptable use policies differentiate between network activity that is appropriate in your organization and activity that is considered inappropriate. These policies are typically focused on Internet usage, and are geared towards maintaining productivity, avoiding legal liabilities (for example, maintaining a non-hostile workplace), and in general controlling web traffic.

You can use URL filtering to define an acceptable use policy with access policies. You can filter on broad categories, such as Gambling, so that you do not need to identify every individual web site that should be blocked. For category matches, you can also specify the relative reputation of sites to allow or block. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked.

Using category and reputation data also simplifies policy creation and administration. It grants you assurance that the system will control web traffic as expected. Finally, because Cisco's threat intelligence is continually updated with new URLs, as well as new categories and risks for existing URLs, you can ensure that the system uses up-to-date information to filter requested URLs. Malicious sites that represent security threats such as malware, spam, botnets, and phishing may appear and disappear faster than you can update and deploy new policies.

The following procedure explains how to implement an acceptable use policy using URL filtering. For purposes of this example, we will block sites of any reputation in several categories, risky Social Networking sites, and an unclassified site, badsite.example.com.

Procedure

Step 1 If you have not already done so, enable the **URL** license.

You must enable the URL license to use URL category and reputation information, or to see the information in dashboards and events. If you are currently using the evaluation license, you are enabling an evaluation version of the license. If you have registered the device, you must purchase the required license and add it to your Smart Software Manager account on Cisco.com.

a) Click **Device**.
b) Click **View Configuration** in the Smart License group.

c) Click **Enable** in the **URL License** group.

The system registers the license with your account, or activates the evaluation license, as appropriate. The group should indicate that the license is enabled, and the button changes to a Disable button.

Step 2 Create a URL filtering access control rule.

You might want to first see the categories for sites your users are visiting before making a blocking rule. If that is the case, you can create a rule with the Allow action for an acceptable category, such as Financial Services. Because all web connections must be inspected to determine if the URL belongs to this category, you would get category information even for non-Financial Services sites.

But there are probably URL categories that you already know you want to block. A blocking policy also forces inspection, so you get category information on connections to unblocked categories, not just the blocked categories.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Click + to add a new rule.

c) Configure the order, title, and action.

- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use the same Source/Destination as the Inside_Outside_Rule created during initial device configuration. You might have created other rules as well. To maximize access control efficiency, it is best to have specific rules early, to ensure the quickest decision on whether a connection is allowed or dropped. For the purposes of this example, select **1** as the rule order.
- **Title**—Give the rule a meaningful name, such as Block_Web_Sites.
- **Action**—Select **Block**.

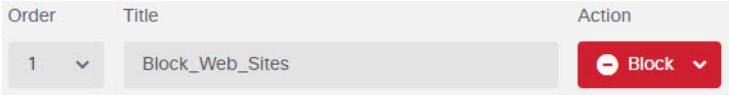

d) On the **Source/Destination** tab, click + for **Source** > **Zones**, select **inside_zone**, then click **OK** in the zones dialog box.

Adding any of the criteria works the same way. Clicking + opens a little dialog box, where you click the items you want to add. You can click multiple items, and clicking a selected item de-selects it; the check marks indicate the selected items. But nothing is added to the policy until you click the **OK** button; simply selecting the items is not sufficient.

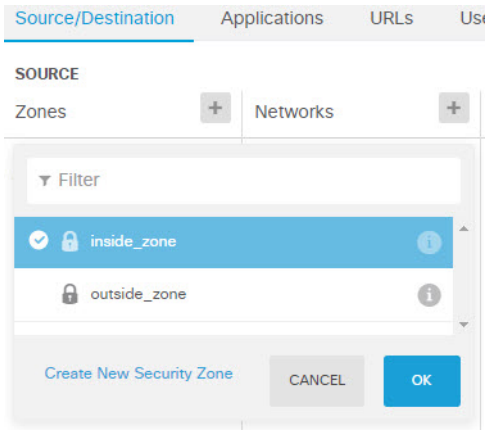

e) Using the same technique, select **outside_zone** for **Destination** > **Zones**.

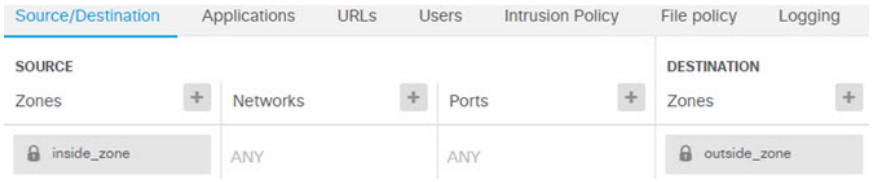

f) Click the **URLs** tab.

g) Click the + for **Categories**, and select the categories you want to fully or partially block.

For purposes of this example, select Adult and Pornography, Bot Nets, Confirmed SPAM Sources, and Social Network. There are additional categories that you would most likely want to block.

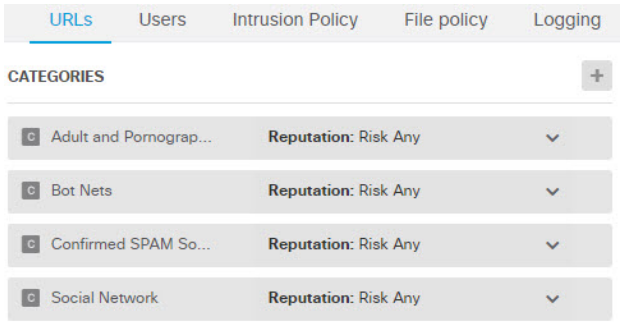

h) To implement reputation-sensitive blocking for the Social Network category, click **Reputation: Risk Any** for that category, deselect **Any**, then move the slider to **Benign sites with security risks**. Click away from the slider to close it.

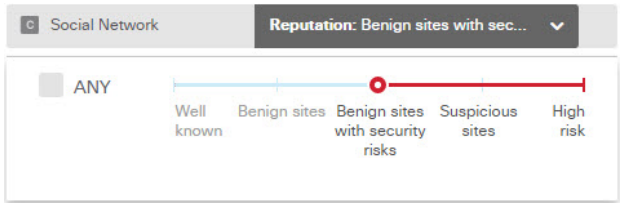

The left of the reputation slider indicates sites that will be allowed, the right side are sites that will be blocked. In this case, only Social Networking sites with reputations in the Suspicious Sites and High Risk

ranges will be blocked. Thus, your users should be able to get to commonly-used Social Networking sites, where there are fewer risks.

Using reputation, you can selectively block sites within a category you otherwise want to allow.

i) Click the + next to the **URLS** list to the left of the categories list.
j) At the bottom of the popup dialog box, click the **Create New URL** link.
k) Enter **badsite.example.com** for both the name and URL, then click **OK** to create the object.

You can name the object the same as the URL or give the object a different name. For the URL, do not include the protocol portion of the URL, just add the server name.

l) Select the new object, then click **OK**.

Adding new objects while editing policies simply adds the object to the list. The new object is not automatically selected.

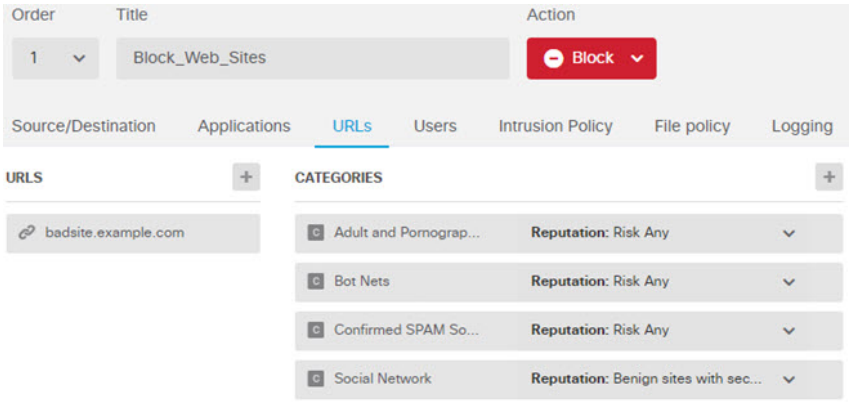

m) Click the **Logging** tab and select **Select Log Action** > **At Beginning and End of Connection**.

You must enable logging to get category and reputation information into the web category dashboard and connection events.

n) Click **OK** to save the rule.

Step 3 (Optional.) Set preferences for URL filtering.

When you enable the URL license, the system automatically enables updates to the web category database. The system checks for updates every 30 minutes, although the data is typically updated once per day. You can turn off these updates if for some reason you do not want them.

You can also elect to send URLs that are not categorized to Cisco for analysis. Thus, if the installed URL database does not have a categorization for a site, the Cisco Cloud might have one. The cloud returns the category and reputation, and your category-based rules can then be applied correctly to the URL request. Selecting this option is important for lower-end systems, which install a smaller URL database due to memory limitations.

a) Click **Device**.
b) Click **System Settings** > **Traffic Settings** > **URL Filtering Preferences**.
c) Select **Query Cisco CSI for Unknown URLs**.
d) Click **Save**.

Step 4 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.

b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

At this point, the monitoring dashboards and events should start showing information about web categories and reputations, and which connections were dropped. You can evaluate this information to determine if your URL filtering is dropping just those sites that are objectionable, or if you need to ease up on the reputation setting for certain categories.

Consider informing users beforehand that you will be blocking access to web sites based on their categorization and reputation.

How to Control Application Usage

The Web has become the ubiquitous platform for application delivery in the enterprise, whether that is browser based application platforms, or rich media applications that use web protocols as the transport in and out of enterprise networks.

FTD inspects connections to determine the application being used. This makes it possible to write access control rules targeted at applications, rather than just targeting specific TCP/UDP ports. Thus, you can selectively block or allow web-based applications even though they use the same port.

Although you can select specific applications to allow or block, you can also write rules based on type, category, tag, risk, or business relevance. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

Cisco frequently updates and adds additional application detectors via system and vulnerability database (VDB) updates. Thus, a rule blocking high risk applications can automatically apply to new applications without you having to update the rule manually.

In this use case, we will block any application that belongs to the **anonymizer/proxy** category.

Before you begin

This use case assumes that you completed the use case How to Gain Insight Into Your Network Traffic, on page 6. That use case explains how to collect application usage information, which you can analyze in the Applications dashboard. Understanding what applications are actually being used can help you design effective application-based rules. The use case also explains how to schedule VDB updates, which will not be repeated here. Ensure that you update the VDB regularly so that applications can be correctly identified.

Procedure

Step 1 Create the application-based access control rule.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Click + to add a new rule.

c) Configure the order, title, and action.

- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use the same Source/Destination as the Inside_Outside_Rule created during initial device configuration. You might have created other rules as well. To maximize access control efficiency, it is best to have specific rules early, to ensure the quickest decision on whether a connection is allowed or dropped. For the purposes of this example, select **1** as the rule order.
- **Title**—Give the rule a meaningful name, such as Block_Anonymizers.
- **Action**—Select **Block**.

d) On the **Source/Destination** tab, click + for **Source** > **Zones**, select **inside_zone**, then click **OK** in the zones dialog box.

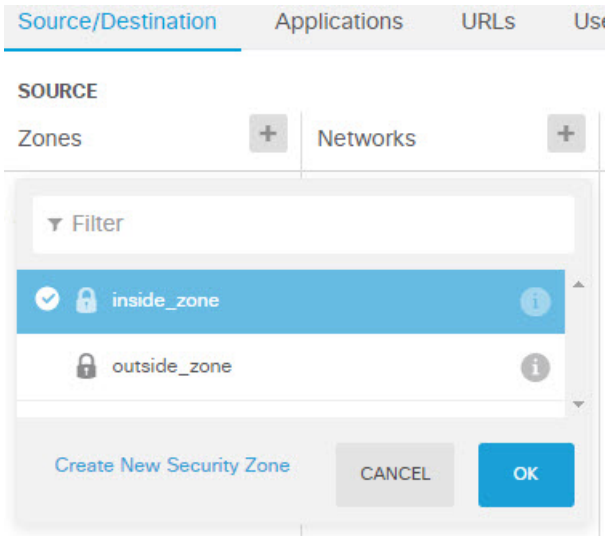

e) Using the same technique, select **outside_zone** for **Destination** > **Zones**.

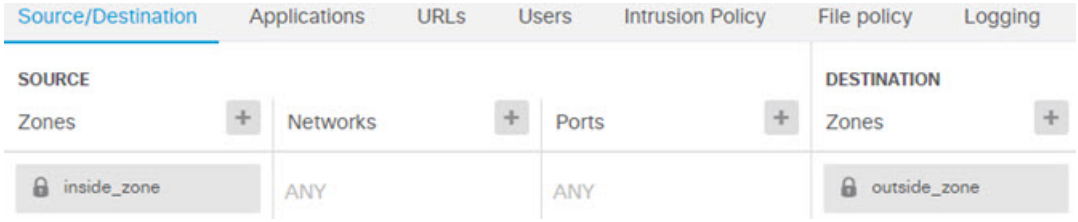

f) Click the **Applications** tab.

g) Click the + for **Applications**, and then click the **Advanced Filter** link at the bottom of the popup dialog box.

Although you can create application filter objects beforehand and select them on the Application Filters list here, you can also specify criteria directly in the access control rule, and optionally save the criteria as a filter object. Unless you are writing a rule for a single application, it is easier to use the Advanced Filter dialog box to find applications and construct appropriate criteria.

As you select criteria, the Applications list at the bottom of the dialog box updates to show exactly which applications match the criteria. The rule you are writing applies to these applications.

Look at this list carefully. For example, you might be tempted to block all very high risk applications. However, as of this writing, TFPT is classified as very high risk. Most organizations would not want to block this application. Take the time to experiment with various filter criteria to see which applications match your selections. Keep in mind that these lists can change with every VDB update.

For purposes of this example, select anonymizers/proxies from the Categories list.

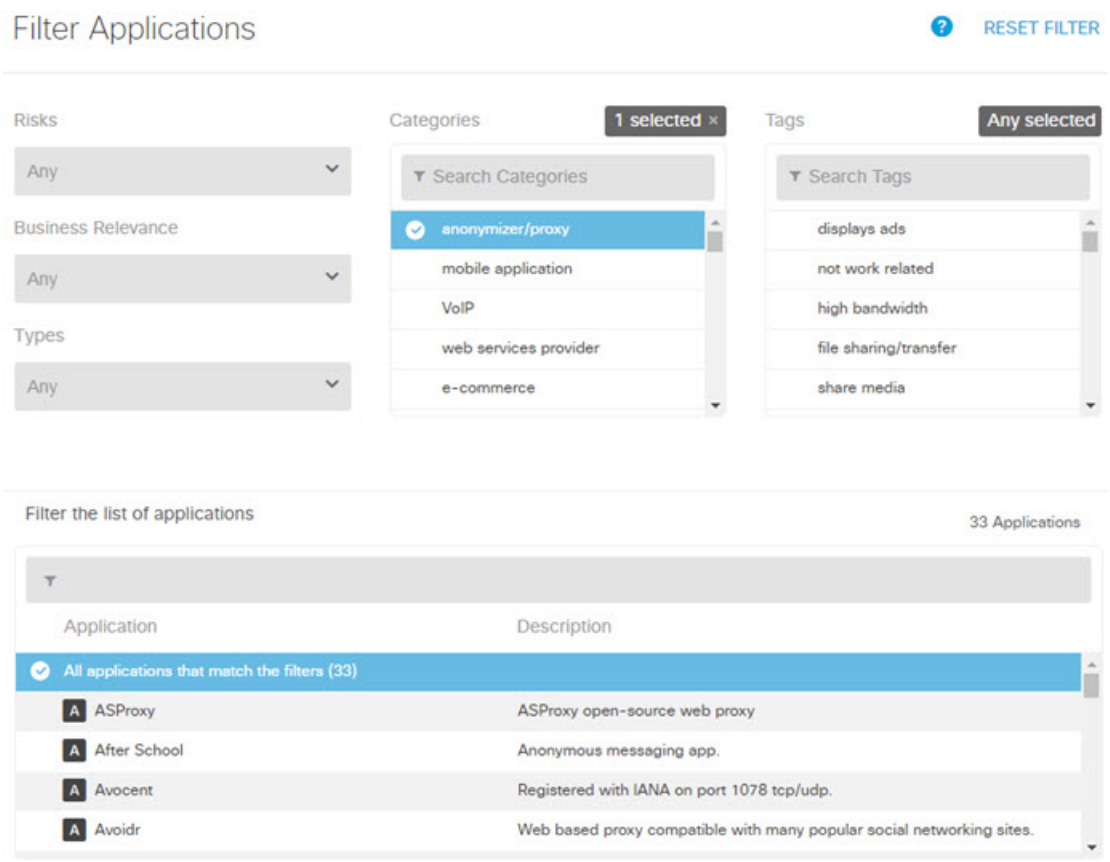

h) Click **Add** in the Advanced Filters dialog box.

The filter is added and shown on the Applications tab.

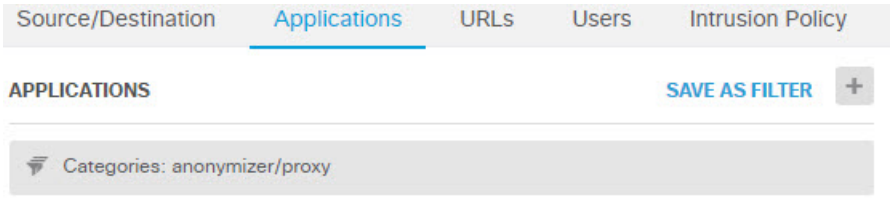

i) Click the **Logging** tab and select **Select Log Action** > **At Beginning and End of Connection**.

You must enable logging to get information about any connections blocked by this rule.

j) Click **OK** to save the rule.

Step 2 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.

b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

Step 3 Click **Monitoring** and evaluate the results.

You might now see dropped connections on the Applications widget on the **Network Overview** dashboard. Use the **All/Denied/Allowed** drop-down options to focus just on dropped applications.

The **Applications** dashboard shows these results as well. If someone tries to use these applications, you should be able to correlate the application with the user attempting the connection, assuming that you enable identity policies and require authentication.

How to Add a Subnet

If you have an available interface on your device, you can wire it to a switch (or another router) to provide services to another subnet.

There are many potential reasons you would add a subnet. For this use case, we will address the following typical scenario.

- The subnet is an inside network using the private network 192.168.2.0/24.
- The interface for the network has the static address 192.168.2.1. In this example, the physical interface is devoted to the network. Another option is to use an already-wired interface and create a subinterface for the new network.
- The device will provide addresses to workstations on the network using DHCP, using 192.168.2.2-192.168.2.254 as the address pool.
- Network access to other inside networks, and to the outside network, will be allowed. Traffic going to the outside network will use NAT to obtain a public address.

Note This example assumes the unused interface is not part of a bridge group. If it is currently a bridge group member, you must first remove it from the bridge group before following this procedure.

Before you begin

Physically connect the network cable to the interface and to the switch for the new subnet.

Procedure

Step 1 Configure the interface.

a) Click **Device**, then click the link in the **Interfaces** summary.
b) Hover over the **Actions** cell on the right side of the row for the interface you wired, and click the edit icon ().
c) Configure the basic interface properties.

- **Name**—A unique name for the interface. For this example, **inside_2**.
- **Status**—Click the status toggle to enable the interface.
- **IPv4 Address** tab—Select **Static** for **Type**, then enter **192.168.2.1/24**.

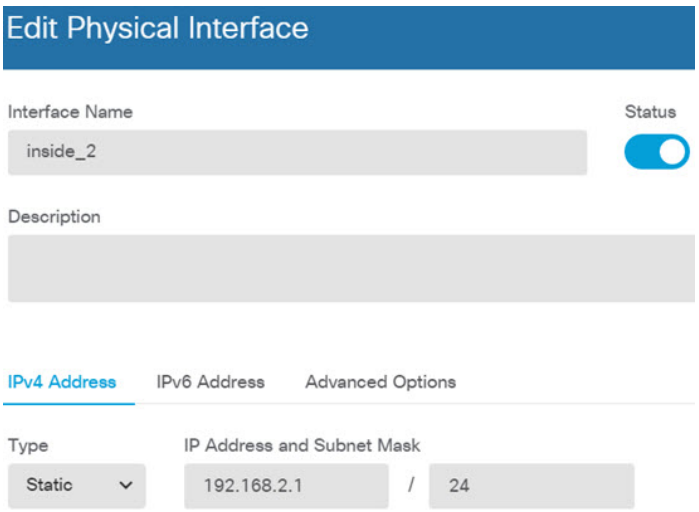

d) Click **Save**.

The interface list shows the updated interface status and the configured IP address.

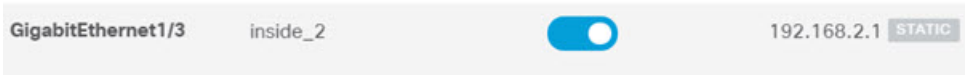

Step 2 Configure the DHCP server for the interface.

a) Click **Device**.
b) Click **System Settings** > **DHCP Server**.
c) Click the **DHCP Servers** tab.

The table lists any existing DHCP servers. If you are using the default configuration, the list includes one for the inside interface.

d) Click + above the table.
e) Configure the server properties.

- **Enable DHCP Server**—Click this toggle to enable the server.
- **Interface**—Select the interface on which you are providing DHCP services. In this example, select inside_2.
- **Address Pool**—The addresses the server can supply to devices on the network. Enter 192.168.2.2-192.168.2.254. Make sure you do not include the network address (.0), the interface address (.1), or the broadcast address (.255). Also, if you need static addresses for any devices on the network, exclude those addresses from the pool. The pool must be a single continuous series of addresses, so choose static addresses from the beginning or ending of the range.

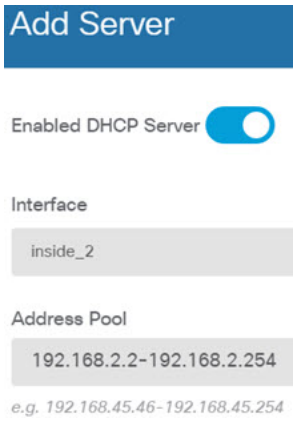

f) Click **Add**.

#	INTERFACE	ENABLED DHCP SERVER	ADDRESS POOL
1	**inside**	Enabled	192.168.1.5-192.168.1.254
2	**inside_2**	Enabled	192.168.2.2-192.168.2.254

Step 3 Add the interface to the inside security zone.

To write policies on an interface, the interface must belong to a security zone. You write policies for the security zones. Thus, as you add and remove interfaces in the zones, you automatically change the policies applied to the interface.

a) Click **Objects** in the main menu.
b) Select **Security Zones** from the objects table of contents.
c) Hover over the **Actions** cell on the right side of the row for the **inside_zone** object, and click the edit icon ( ).
d) Click + under **Interfaces**, select the inside_2 interface, and click **OK** in the interfaces list.

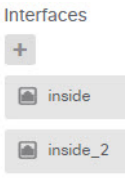

e) Click **Save**.

Step 4 Create an access control rule that allows traffic between the inside networks.

Traffic is not automatically allowed between any interfaces. You must create access control rules to allow the traffic that you want. The only exception is if you allow traffic in the access control rule's default action. For the purposes of this example, we will assume you retained the block default action that the device setup wizard configures. Thus, you need to create a rule that will allow traffic between the inside interfaces. If you have already created a rule like this, skip this step.

a) Click **Policies** in the main menu.

Ensure that the **Access Control** policy is displayed.

b) Click + to add a new rule.

c) Configure the order, title, and action.

- **Order**—The default is to add new rules to the end of the access control policy. However, you must place this rule ahead of (above) any rule that would match the same Source/Destination and other criteria, or the rule will never be matched (a connection matches one rule only, and that is the first rule it matches in the table). For this rule, we will use unique Source/Destination criteria, so adding the rule to the end of the list is acceptable.
- **Title**—Give the rule a meaningful name, such as Allow_Inside_Inside.
- **Action**—Select **Allow**.

d) On the **Source/Destination** tab, click + for **Source** > **Zones**, select **inside_zone**, then click **OK** in the zones dialog box.

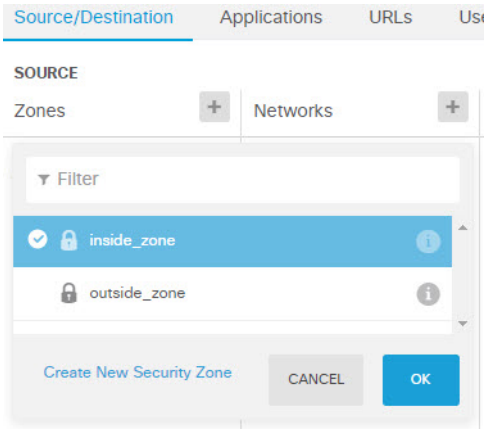

e) Using the same technique, select **inside_zone** for **Destination** > **Zones**.

A security zone must contain at least two interfaces to select the same zone for source and destination.

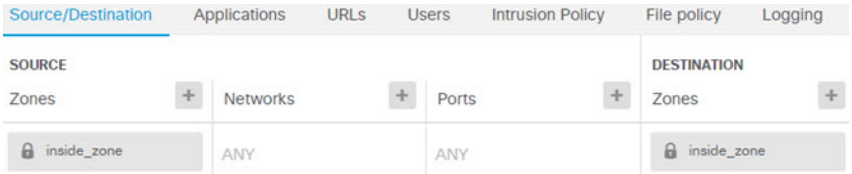

f) (Optional.) Configure intrusion and malware inspection.

Although the inside interfaces are in a trusted zone, it is typical for users to connect laptops to the network. Thus, a user might unknowingly bring a threat inside your network from an outside network or a Wi-Fi hot spot. Thus, you might want to scan for intrusions and malware in traffic that goes between your inside networks.

Consider doing the following.

- Click the **Intrusion Policy** tab, enable the intrusion policy, and use the slider to select the Balanced Security and Connectivity policy.
- Click the **File Policy** tab, then select the Block Malware All policy.

g) Click the **Logging** tab and select **Select Log Action** > **At Beginning and End of Connection**.

You must enable logging to get information about any connections that match this rule. Logging adds statistics to the dashboard as well as showing events in the event viewer.

h) Click **OK** to save the rule.

Step 5 Verify that required policies are defined for the new subnet.

By adding the interface to the inside_zone security zone, any existing policies for inside_zone automatically apply to the new subnet. However, take the time to inspect your policies and ensure that no additional policies are needed.

If you completed the initial device configuration, the following policies should already apply.

- **Access Control**—The Inside_Outside_Rule should allow all traffic between the new subnet and the outside network. If you followed the previous use cases, the policy also provides intrusion and malware inspection. You must have a rule that allows some traffic between the new network and the outside network, or users cannot access the Internet or other external networks.
- **NAT**—The InsideOutsideNATrule applies to any interface going to the outside interface, and applies interface PAT. If you kept this rule, traffic from the new network going to the outside will have the IP address translated to a unique port on the outside interface's IP address. If you do not have a rule that applies to all interfaces, or the inside_zone interfaces, when going to the outside interface, you might need to create one now.
- **Identity**—There is no default identity policy. However, if you followed previous use cases, you might have an identity policy that already requires authentication for the new network. If you do not have an identity policy that applies, create one now if you want to have user-based information for the new network.

Step 6 Commit your changes.

a) Click the **Deploy Changes** icon in the upper right of the web page.

b) Click the **Deploy Now** button.

Wait for deployment to finish. The deployment summary should indicate that you have successfully deployed your changes, and the task status for the job should be Deployed.

What to do next

Verify that workstations on the new subnet are getting IP addresses using DHCP, and that they can reach other inside networks and the outside network. Use the monitoring dashboards and the event viewer to evaluate network usage.

More Examples

In addition to the examples in the Use Case chapter, there are example configurations in some of the chapters that explain specific services. You might find the following examples of interest.

Network Address Translation (NAT)

NAT for IPv4 addresses

- Providing Access to an Inside Web Server (Static Auto NAT)
- Single Address for FTP, HTTP, and SMTP (Static Auto NAT-with-Port-Translation)
- Different Translation Depending on the Destination (Dynamic Manual PAT)
- Different Translation Depending on the Destination Address and Port (Dynamic Manual PAT)
- DNS Reply Modification, DNS Server on Outside
- DNS Reply Modification, DNS Server on Host Network
- Exempting Site-to-Site VPN Traffic from NAT

NAT for IPv6 addresses

- NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet
- NAT64/46 Example: Inside IPv6 Network with Outside IPv4 Internet and DNS Translation
- NAT66 Example, Static Translation between Networks
- NAT66 Example, Simple IPv6 Interface PAT
- DNS 64 Reply Modification

Remote Access Virtual Private Network (RA VPN)

- How to Provide Internet Access on the Outside Interface for Remote Access VPN Users (Hair Pinning)
- How to Use a Directory Server on an Outside Network with Remote Access VPN
- How to Customize the AnyConnect Client Icon and Logo

Site-to-Site Virtual Private Network (VPN)

- Exempting Site-to-Site VPN Traffic from NAT
- How to Provide Internet Access on the Outside Interface for External Site-to-Site VPN Users (Hair Pinning)