

Access Control

The following topics explain access control rules. These rules control which traffic is allowed to pass through the device, and apply advanced services to the traffic, such as intrusion inspection.

- Best Practices for Access Control, on page 1
- Access Control Overview, on page 4
- License Requirements for Access Control, on page 14
- Guidelines and Limitations for Access Control Policies, on page 15
- Configuring the Access Control Policy, on page 17
- Monitoring Access Control Policies, on page 28
- Examples for Access Control, on page 30

Best Practices for Access Control

The access control policy is your primary tool for protecting your internal networks and preventing your users from accessing undesirable external network resources, such as inappropriate web sites. Thus, we recommend that you pay special attention to this policy and fine-tune it to get the level of protection and connectivity that you need.

The following procedure provides an overview of the basic things you should do with the access control policy. This is an overview, and does not provide exhaustive steps for performing each task.

To get to the access control policy, choose **Policies** > **Access Control**.

Procedure

Step 1 Configure the default action for the policy.

The default action handles connections that do not match the specific rules in the policy. By default, this action is **Block**, so that anything you miss in the rules is blocked. Thus, you simply need to write access control rules that allow desirable traffic. This is the traditional way to configure the access control policy.

You can do the opposite, where you allow traffic by default, and write rules that drop known undesirable traffic, so that you do not need to have rules for everything you want to allow. This makes it easy for new services to be used, but opens the risk that new undesirable traffic will get through before you notice it.

Step 2 Click the **Access Policy Settings** ( ) button, and enable the **TLS Server Identity Discovery** option.

This option improves the initial application detection and URL category and reputation identification for TLS 1.3 connections. If you do not enable this option, TLS 1.3 traffic might not match your intended rules. This option can also improve the efficacy of decryption rules.

Step 3 Create as few access control rules as possible.

With traditional firewalls, you might end up with tens of thousands of rules for various combinations of IP address and port. With a next-generation firewall, you can use advanced inspection and avoid some of these detailed rules. The fewer rules you have, the faster the system can evaluate traffic, and the easier it will be for you to find and fix problems within your rule set.

Step 4 Enabling logging on your access control rules.

Statistics are collected for matching traffic only if you enable logging. Your monitoring dashboards will be inaccurate if you do not enable logging.

Step 5 Put very specific rules at the top of the policy, and ensure that specific rules are above any more general rule that would match the connections a specific rule would also match.

The policy is evaluated top-down, first match wins. Thus, if you put in a rule that blocks all traffic to a specific subnet, then follow it with a rule that allows access to a single IP address within the subnet, traffic to that address will not be allowed, because the first rule will block it.

In addition, place rules that target traffic based only on traditional criteria such as ingress/egress interface, and source/destination IP address, port, or Geolocation, ahead of rules that require deep inspection, such as those that apply to user criteria, URL filtering, or application filtering. Because these rules do not require inspection, putting them early can get you quicker access control decisions for matching connections.

For more suggestions, see Best Practices for Access Control Rule Order, on page 13.

Step 6 Pair Block and Allow rules to target subsets of traffic.

For example, it is likely that you want to allow a lot of HTTP/HTTPS traffic, yet block access to some undesirable sites such as pornography or gambling. You can accomplish this by creating the following rules and keeping them sequential within the policy (for example, rules 11 and 12).

- A URL filtering Block rule that targets undesirable URL categories, applied to the inside security zone (source) and outside security zone (destination), and any IP address, port, or Geolocation. For example, Block Botnets, Child Abuse Content, Cryptojacking, DNS Tunneling, Ebanking Fraud, Exploits, Extreme, Filter Avoidance, Gambling, Hacking, Hate Speech, High Risk Sites and Locations, Illegal Activities, Illegal Downloads, Illegal Drugs, Malicious Sites, Malware Sites, Mobile Threats, P2P Malware Node, Phishing, Pornography, Spam, Spyware and Adware.
- An application filtering Allow rule for the HTTP and HTTPS applications, applied to the inside security zone (source) and outside security zone (destination), and any IP address, port, or Geolocation. After the URL filtering rule blocks access to unwanted web resources, this rule allows all other HTTP/HTTPS access.

Step 7 Use advanced next-generation firewall features to target traffic regardless of IP address or port.

Attackers or other bad actors can frequently change IP addresses and ports to evade traditional access control traffic-matching criteria. Instead, use the following next-generation features:

- User criteria—Configure the Identity policy to obtain information about the user who is initiating the traffic. Ideally, your Active Directory server organizes users into groups, and you can create access control rules that allow or block traffic based on user group membership. For example, allow Engineers to access your development subnets, but implicitly block anyone who is not in the Engineers group. Use

groups rather than individual user names, so you do not need to continually update the rules as people are added to the network.

- Application criteria—Use application filtering criteria to allow or block types of applications. This, if a user changes ports for an HTTP connection, the system can recognize that it is HTTP, even though it isn't going to port 80. For more suggestions, see Best Practices for Application Filtering, on page 5.
- URL category and reputation criteria—Use URL filtering based on category to dynamically allow or block sites based on the type of site. Within the site type, or category, you can fine-tune your rule based on whether the site has a reputation as a good actor or a bad actor. By using category and reputation, you will not need to constantly adjust your rules as sites change URLs, which you would have to do if you tried to manually block sites by URL. For more suggestions, see Best Practices for Effective URL Filtering, on page 10.

  You can also apply URL category/reputation filtering rules to the FQDN in a DNS lookup request. The system can prevent the DNS response for blocked category/reputation, effectively blocking the user's connection attempt. For details, see Filtering DNS Requests Based on URL Category and Reputation, on page 12.

Step 8 Apply intrusion inspection to all of your Allow rules.

One of the powerful aspects of next-generation firewalls is that you can apply intrusion inspection and access control using the same device. Apply an intrusion policy to each Allow rule, so that if an attack does enter your network through a normally benign path, you can catch it and drop the attacking connection.

If your default action is Allow, you can also apply intrusion protection for traffic that matches the default action.

Step 9 Also configure the Security Intelligence policy to block unwanted IP addresses and URLs.

The Security Intelligence policy is applied before the access control policy, so it can block unwanted connections before your access control rules are even evaluated. This can provide an early block and help you reduce the complexity of your access control rules.

Step 10 Consider implementing the SSL Decryption policy.

The system cannot perform deep inspection on encrypted traffic. If you configure the SSL Decryption policy, the access control policy is applied to a decrypted version of the traffic. Thus, deep inspection can identify attacks (using the intrusion policy), and rule matching is better because application and URL filtering can be applied more effectively. Any traffic that the access control policy allows is then re-encrypted before being sent out of the device, so the end user does not lose the protection of encryption.

Step 11 Enable object group search to simplify the deployment of your rules.

Starting with release 7.2, this feature is enabled by default on new deployments, but is not automatically enabled on upgraded systems.

Enabling object group search reduces memory requirements for access control policies that include network objects. However, it is important to note that object group search might also decrease rule lookup performance and thus increase CPU utilization. You should balance the CPU impact against the reduced memory requirements for your specific access control policy. In most cases, enabling object group search provides a net operational improvement.

You can set this option using FlexConfig by issuing the **object-group-search access-control** command; use the **no** form of the command in the negate template.

Access Control Overview

The following topics explain access control policies.

Access Control Rules and the Default Action

Use the access control policy to allow or block access to network resources. The policy consists of a set of ordered rules, which are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched.

You can control access based on:

- Traditional network characteristics such as source and destination IP addresses, protocol, ports, and interfaces (in the form of security zones).
- The fully-qualified domain name (FQDN) of the source or destination (in the form of a network object). Traffic matching is based on the IP address returned for the name from a DNS lookup.
- The security group tag (SGT) assigned to the source or destination by the Cisco Identity Services Engine (ISE).
- The application that is being used. You can control access based on the specific application, or you can create rules that cover categories of applications, applications tagged with a particular characteristic, the type of application (client, server, web), or the application's risk or business relevance rating.
- The destination URL of a web request, including the generalized category of the URL. You can refine category matches based on the public reputation of the target site.
- The URL category and reputation of an FQDN in a DNS lookup request. You can block the DNS response for unwanted categories or poor reputations, effectively preventing the subsequent connection attempt.
- The user who is making the request, or the user groups to which the user belongs.

For unencrypted traffic that you allow, you can apply IPS inspection to check for threats and block traffic that appears to be an attack. You can also use file policies to check for prohibited files or malware.

Any traffic that does not match an access rule is handled by the access control **Default Action**. If you allow traffic by default, you can apply intrusion inspection to the traffic. However, you cannot perform file or malware inspection on traffic handled by the default action.

Application Filtering

You can use access control rules to filter traffic based on the application used in the connection. The system can recognize a wide variety of applications, so that you do not need to figure out how to block one web application without blocking all web applications.

For some popular applications, you can filter on different aspects of the application. For example, you could create a rule that blocks Facebook Games without blocking all of Facebook.

You can also create rules based on general application characteristics, blocking or allowing entire groups of applications by selecting risk or business relevance, type, category, or tag. **However, as you select categories in an application filter, look over the list of matching applications to ensure you are not including**

unintended applications. For a detailed explanation of the possible groupings, see Application Criteria, on page 22.

Application Control for Encrypted and Decrypted Traffic

If an application uses encryption, the system might not be able to identify the application.

The system can detect application traffic encrypted with StartTLS, including SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS ClientHello message, or the subject distinguished name value from the server certificate.

Use the application filters dialog box to determine if your application requires decryption by selecting the following Tags, then examining the list of applications.

- **SSL Protocol**—You do not need to decrypt traffic tagged as SSL Protocol. The system can recognize this traffic and apply your access control action. Access control rules for the listed applications should match to expected connections.
- **Decrypted Traffic**—The system can recognize this traffic only if you first decrypt the traffic. Configure SSL decryption rules for this traffic.

Filtering on Common Industrial Protocol (CIP) and Modbus Applications (ISA 3000)

You can enable the Common Industrial Protocol (CIP) and Modbus pre-processors on Cisco ISA 3000 devices, and filter on CIP and Modbus applications in access control rules. All CIP application names start with "CIP," such as CIP Write. There is only one application for Modbus.

To enable the pre-processors, you must go into expert mode in a CLI session (SSH or Console) and issue the following command to enable one or both of these Supervisory Control and Data Acquisition (SCADA) applications.

sudo /usr/local/sf/bin/enable_scada.sh {cip | modbus | both}

For example, to enable both pre-processors:

```
> expert
admin@firepower:~$ sudo /usr/local/sf/bin/enable_scada.sh both
```

Note You must issue this command after every deployment. These pre-processors are disabled during deployment.

Best Practices for Application Filtering

Please keep the following recommendations in mind when designing your application filtering access control rules.

- To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.
- Avoid combining application and URL criteria in the same rule, especially for encrypted traffic.
- If you write a rule for traffic that is tagged **Decrypted Traffic**, ensure that you have an SSL Decryption rule that will decrypt the matching traffic. These applications can be identified in decrypted connections only.

- TLS 1.3 encrpyts most handshake messages, so certificate information is not readily available. For traffic encrypted with TLS 1.3 to match access rules that use application or URL filtering effectively, the system must obtain a clear-text certificate for the server. We recommend that you enable **TLS 1.3 Certificate Visibility** in the access control settings. If you enable this option, the system checks whether a certificate for the site is stored in cache based on the IP address and Server Name Indication (SNI) in the client hello packet. If one is not available, the system uses a TLS 1.2 probe to obtain the certicate, which can then be used for application/URL category and reputation identification without decrypting the connection.
- The system can detect multiple types of Skype application traffic. To control Skype traffic, choose the Skype tag from the Application Filters list rather than selecting individual applications. This ensures that the system can detect and control all Skype traffic the same way.
- To control access to Zoho mail, select both the Zoho and Zoho Mail applications.

URL Filtering

You can use access control rules to filter traffic based on the URL used in an HTTP or HTTPS connection. Note that URL filtering for HTTP is more straight-forward than it is for HTTPS, because HTTPS is encrypted.

You can use the following techniques to implement URL filtering.

- Category and reputation-based URL filtering—With a URL filtering license, you can control access to web sites based on the URL's general classification (category) and risk level (reputation). This is by far the easiest and most effective way to block unwanted sites.
- Manual URL filtering—With any license, you can manually specify individual URLs, and groups of URLs, to achieve granular, custom control over web traffic. The main purpose of manual filtering is to create exceptions to category-based block rules, but you can use manual rules for other purposes.

The following topics provide more information on URL filtering.

Filtering URLs by Category and Reputation

With a URL filtering license, you can control access to web sites based on the category and reputation of the requested URLs:

- Category—A general classification for the URL. For example, ebay.com belongs to the Auctions category, and monster.com belongs to the Job Search category. A URL can belong to more than one category.
- Reputation—How likely the URL is to be used for purposes that might be against your organization's security policy. Reputations range from Untrusted (level 1) to Trusted (level 5). Reputation is determined by the hostname, e.g. all pages on www.cisco.com are considered to have the same reputation.

URL categories and reputations help you quickly configure URL filtering. For example, you can use access control to block untrusted URLs in the Illegal Drugs category.

For a description of the categories, see https://www.talosintelligence.com/categories.

Using category and reputation data also simplifies policy creation and administration. Sites that represent security threats, or that serve undesirable content, might appear and disappear faster than you can update and deploy new policies. As Cisco updates the URL database with new sites, changed classifications, and changed reputations, your rules automatically adjust to the new information. You do not need to edit your rules to account for new sites.

If you enable regular URL database updates, you can ensure that the system uses up-to-date information for URL filtering. You can also enable communications with Cisco Collective Security Intelligence (CSI) to obtain the latest threat intelligence for URLs with unknown category and reputation. For more information, see Configuring URL Filtering Preferences.

Note To see URL category and reputation information in events and application details, you must create at least one rule with a URL condition.

Looking Up the Category and Reputation for a URL

You can check on the category and reputation for a particular URL. You can go to the URL tab of an access control rule, or SSL decryption rule, or go to **Device** > **System Settings** > **URL Filtering Preferences**. There, you can enter the URL in the **URL to Check** field and click **Go**.

You will be taken to a web site that shows the results of the lookup. You can use this information to help you check the behavior of your category- and reputation-based URL filtering rules.

If you disagree with the categorization, you can click the **Submit a URL Category Dispute** in the Firewall Device Manager to tell us what you think.

Manual URL Filtering

You can supplement or selectively override category and reputation-based URL filtering by manually filtering individual URLs or groups of URLs. You can perform this type of URL filtering without a special license.

For example, you might use access control to block a category of web sites that are not appropriate for your organization. However, if the category contains a web site that is appropriate, and to which you want to provide access, you can create a manual Allow rule for that site and place it before the Block rule for the category.

To configure manual URL filtering, you create a URL object with the destination URL. How this URL is interpreted is based on the following rules:

- If you do not include a path (that is, there are no / characters in the URL), the match is based on the server's hostname only. If you include one or more / character, the entire URL string is used for a substring match. Then, a URL is considered a match if any of the following are true:
  - The string is at the beginning of the URL.
  - The string follows a dot.
  - The string contains a dot in the beginning.
  - The string follows the :// characters.

  For example, ign.com matches ign.com or www.ign.com, but not versign.com.

Note We recommend that you do not use manual URL filtering to block or allow individual web pages or parts of sites (that is, URL strings with / characters), as servers can be reorganized and pages moved to new paths.

- The system disregards the encryption protocol (HTTP vs HTTPS). In other words, if you block a website, both HTTP and HTTPS traffic to that website is blocked, unless you use an application condition to target a specific protocol. When creating a URL object, you do not need to specify the protocol when creating an object. For example, use example.com rather than http://example.com.
- If you plan to use a URL object to match HTTPS traffic in an access control rule, create the object using the subject common name in the public key certificate used to encrypt the traffic. Also, the system disregards subdomains within the subject common name, so do not include subdomain information. For example, use example.com rather than www.example.com.

  However, please understand that the subject common name in the certificate might be completely unrelated to a web site's domain name. For example, the subject common name in the certificate for youtube.com is *.google.com (this of course might change at any time). You will get more consistent results if you use the SSL Decryption policy to decrypt HTTPS traffic so that URL filtering rules work on decrypted traffic.



Note URL objects will not match HTTPS traffic if the browser resumes a TLS session because the certificate information is no longer available. Thus, even if you carefully configure the URL object, you might get inconsistent results for HTTPS connections.

Filtering HTTPS Traffic

Because HTTPS traffic is encrypted, performing URL filtering directly on HTTPS traffic is not as straight-forward as it is on HTTP traffic. For that reason, you should consider using SSL Decryption policies to decrypt all HTTPS traffic that you intend to filter. That way, the URL filtering access control policies work on decrypted traffic, and you get the same results you would get for regular HTTP traffic.

However, if you do intend to allow some HTTPS traffic to pass undecrypted into the access control policy, you need to understand that rules match HTTPS traffic differently than they do for HTTP traffic. To filter encrypted traffic, the system determines the requested URL based on information passed during the SSL handshake: the subject common name in the public key certificate used to encrypt the traffic. There might be little or no relationship between the web site hostname in the URL and the subject common name.

You can improve HTTPS matching for category/reputation rules if you enable DNS request filtering. The system can determine the category and reputation during the DNS resolution phase, and block the DNS reply for unwanted combinations, before the user can start the HTTPS connection attempt. For allowed DNS responses, the system will have the category/reputation information available for subsequent HTTPS connections. See DNS Request Filtering, on page 11.

HTTPS filtering, unlike HTTP filtering, disregards subdomains within the subject common name. Do not include subdomain information when manually filtering HTTPS URLs. For example, use example.com rather than www.example.com. Also, review the content of the certificates used by the site to ensure you have the right domain, the one used in the subject common name, and that this name will not conflict with your other rules (for example, the name for a site you want to block might overlap with one you want to allow). For example, the subject common name in the certificate for youtube.com is *.google.com (this of course might change at any time).



Note URL objects will not match HTTPS traffic if the browser resumes a TLS session because the certificate information is no longer available. Thus, even if you carefully configure the URL object, you might get inconsistent results for HTTPS connections.

Controlling Traffic by Encryption Protocol

The system disregards the encryption protocol (HTTP vs HTTPS) when performing URL filtering. This occurs for both manual and reputation-based URL conditions. In other words, URL filtering treats traffic to the following web sites identically:

- http://example.com
- https://example.com

To configure a rule that matches only HTTP or HTTPS traffic, but not both, either specify the TCP port in the Destination condition or add an application condition to the rule. For example, you could allow HTTPS access to a site while disallowing HTTP access by constructing two access control rules, each with an TCP port or application, and URL, condition.

The first rule allows HTTPS traffic to the website:

Action: Allow
TCP port or Application: HTTPS (TCP port 443)
URL: example.com

The second rule blocks HTTP access to the same website:

Action: Block
TCP port or Application: HTTP (TCP port 80)
URL: example.com

Comparing URL and Application Filtering

URL and application filtering have similarities. But you should use them for very distinct purposes:

- URL filtering is best used to block or allow access to an entire web server. For example, if you do not want to allow any type of gambling on your network, you can create a URL filtering rule to block the Gambling category. With this rule, users cannot get to any pages on any web server within the category.
- Application filtering is useful for blocking specific applications regardless of the hosting site, or for blocking specific features of an otherwise allowable web site. For example, you could block just the Facebook Games application without blocking all of Facebook.

Because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic, it is a good policy to create separate rules for URL and application criteria. If you do need to combine application and URL criteria in a single rule, you should place these rules after straight-forward application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule. Because URL filtering block rules are more broad than application filtering, you should place them above application-only rules.

If you do combine application and URL criteria, you might need to monitor your network more carefully to ensure that you are not allowing access to unwanted sites and applications.

Best Practices for Effective URL Filtering

Please keep the following recommendations in mind when designing your URL filtering access control rules.

- Use category and reputation blocking whenever possible. This ensures that new sites get blocked automatically as they are added to the categories, and that blocking based on reputation is adjusted if a site becomes more (or less) reputable.
- When using URL category matching, note that there are cases where the login page for a site is in a different category than the site itself. For example, Gmail is in the Web-based Email category, whereas the login page is in the Search Engines and Portals category. If you have different rules with different actions for the categories, you might get unintended results.
- Use URL objects to target entire web sites and to make exceptions to category blocking rules. That is, to allow specific sites that would otherwise get blocked in a category rule.
- If you want to manually block a web server (using a URL object), it is much more effective to do so in the Security Intelligence policy. The Security Intelligence policy drops connections before the access control rules are evaluated, so you get a faster, more efficient, block.
- For the most effective filtering of HTTPS connections, implement SSL decryption rules to decrypt traffic for which you are writing an access control rule. Any decrypted HTTPS connections are filtered as HTTP connections in the access control policy, so you avoid all of the limitations for HTTPS filtering.
- TLS 1.3 encrpyts most handshake messages, so certificate information is not readily available. For traffic encrypted with TLS 1.3 to match access rules that use application or URL filtering effectively, the system must obtain a clear-text certificate for the server. We recommend that you enable **TLS 1.3 Certificate Visibility** in the access control settings. If you enable this option, the system checks whether a certificate for the site is stored in cache based on the IP address and Server Name Indication (SNI) in the client hello packet. If one is not available, the system uses a TLS 1.2 probe to obtain the certicate, which can then be used for application/URL category and reputation identification without decrypting the connection.
- Place URL blocking rules before any application filtering rules, because URL filtering blocks entire web servers, whereas application filtering targets specific application usage regardless of the web server.
- If you want to block high risk sites whose category is unknown, select the Uncategorized category and adjust the reputation slider to Questionable or Untrusted.
- You can improve overall URL filtering efficacy by also enabling DNS request filtering. When you use DNS request filtering, the system determines the FQDN's URL category and reputation at DNS lookup time, so the information is available if a subsequent HTTP/HTTPS request comes for the same destination. Additionally, if you block the category/reputation, the attempted connection gets stopped at the DNS request stage, rather than the web-session establishment stage. See DNS Request Filtering, on page 11.

What the User Sees When You Block Web Sites

When you block web sites with URL filtering rules, what the user sees differs based on whether the site is encrypted.

- HTTP connections—The user sees a system default block response page instead of the normal browser page for timed out or reset connections. This page should make it clear that you blocked the connection on purpose.
- HTTPS (encrypted) connections—The user does not see the system default block response page. Instead, the user sees the browser's default page for a secure connection failure. The error message does not indicate the site was blocked due to policy. Instead, errors might indicate that there are no common

encryption algorithms. It will not be obvious from this message that you blocked the connection on purpose.

In addition, web sites might be blocked by other access control rules that are not explicitly URL filtering rules, or even by the default action. For example, if you block entire networks or geolocations, any web sites on that network or in that geographic location are also blocked. Users blocked by these rules may, or may not, get a response page as described in the limitations below.

If you implement URL filtering, consider explaining to end users what they might see when a site is intentionally blocked, and what types of site you are blocking. Otherwise, they might spend a good deal of time troubleshooting blocked connections.

Limitations of HTTP Response Pages

HTTP response pages do not always appear when the system blocks web traffic.

- The system does not display a response page when web traffic is blocked as a result of a promoted access control rule (an early-placed blocking rule with only simple network conditions).
- The system does not display a response page when web traffic is blocked before the system identifies the requested URL.
- The system does not display a response page for encrypted connections blocked by access control rules.

DNS Request Filtering

You can apply the URL category and reputation database to DNS lookup requests, even for connection attempts that are not HTTP/HTTPS.

For example, if a user tries to make an FTP connection to www.example.com, you can configure the system to look up the category and reputation of www.example.com when it sees the DNS lookup request for that fully-qualified domain name (FQDN). If your DNS/URL filtering rule for the returned category/reputation is a block rule, the system blocks the DNS reply. Thus, the user does not get an IP address for the FQDN, and their connection attempt fails.

By enabling DNS lookup request filtering, you can extend your URL filtering rules to protocols other than HTTP/HTTPS, and prevent FTP, TFTP, SCP, ICMP, and any other protocols from establishing a connection to a site that you are blocking for web access. This works so long as the user uses an FQDN name and thus requires a DNS lookup. If the user uses an IP address, there is no DNS request and DNS request blocking is not possible.

For HTTP/HTTPS traffic, doing the category/reputation lookup at DNS request time might improve system performance, as it can prevent the connection before the attempt to establish the web session. This might be especially beneficial for HTTPS, which is encrypted. By denying at the DNS request stage, the system never sees an HTTPS connection, and thus your decryption rules do not need to be evaluated, nor does the system need to perform the more difficult task of matching an encrypted session to the right access control rule.

Guidelines for DNS Request Filtering

Keep the following in mind as you configure DNS request filtering:

- DNS request filtering works on the DNS session only. If you allow the DNS reply (that is, the URL filtering rule action is Allow), then the subsequent connection that the user establishes with the returned IP address will be matched against your access control rules separately. The connection might match a

different rule, and thus be blocked or allowed for other reasons. For example, if you allow an FTP attempt to get an IP address through a DNS lookup, you might have another access control rule that prohibits FTP connections, and the connection will ultimately be blocked.

- DNS lookup requests that match access control rules that come before your URL/DNS request filtering rules will be allowed or blocked according to the matching rule. Category/reputation lookup will not be done for these connections.
- This feature requires that you implement URL filtering based on category/reputation. You must have the URL filtering license for this type of URL filtering. If you have no URL filtering rules based on category/reputation, then DNS request filtering is not relevant and you should not enable it.
- Connection events generated by DNS filtering include the following fields of special interest: DNS Query, URL Category, and URL Reputation. The DNS Query field shows the fully-qualified domain name (FQDN) for the lookup request. For DNS filtering events, the URL field will be blank.
- DNS request filtering uses the URL category and reputation database only. Any URL objects or other manual URL filtering defined in a matching access control rule are ignored. If you want to implement manual DNS name blocking, use the Security Intelligence DNS policy.

Filtering DNS Requests Based on URL Category and Reputation

The following procedure explains how to implement DNS lookup request filtering.

Before you begin

You must enable the URL license if it is not already enabled.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 If necessary, click the **Access Policy Settings** ( ) button, select the **Reputation Enforcement on DNS Traffic** option, and click **OK**.

This option enables DNS request filtering for the access control policy. The option is enabled by default.

Step 3 Evaluate existing URL filtering rules, or create new ones, to implement filtering based on URL category and reputation that will also apply to DNS requests.

URL filtering normally applies to HTTP/HTTPS traffic only, so there is no reason to restrict these rules based on application or port. However, if you do have these restrictions, ensure that the rule can also apply to DNS requests:

- On the **Source/Destination** tab, if the **Destination Ports** field has **Any**, no change is needed. If you specified ports, add **DNS over UDP** and **DNS over TCP** to the list.
- On the **Applications** tab, if the application list simply has **Any**, no change is needed. If you specified any applications or application filters, add the **DNS** application to the list or filter. The other DNS-related options are not relevant for this purpose.

For information on creating access control rules, see Configuring Access Control Rules, on page 19.

Step 4 Evaluate preceding rules to ensure that DNS requests do not match those rules.

Category and reputation determination happens only if the DNS request matches a URL filtering rule that has category and reputation specifications. Any DNS requests that match rules earlier in the access control policy than your URL filtering rule bypass DNS request filtering. Such DNS requests are handled according to the matching rule, either blocked or allowed.

Intrusion, File, and Malware Inspection

Intrusion and file policies work together as the last line of defense before traffic is allowed to its destination:

- Intrusion policies govern the system's intrusion prevention capabilities.
- File policies govern the system's file control and malware defense capabilities.

All other traffic handling occurs before network traffic is examined for intrusions, prohibited files, and malware. By associating an intrusion or 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 the traffic with an intrusion policy, a file policy, or both.

You can configure intrusion and file policies on rules that **allow** traffic only. Inspection is not performed on rules set to **trust** or **block** traffic. In addition, if the default action for the access control policy is **allow**, you can configure an intrusion policy but not a file policy.

For any single connection handled by an access control rule, file inspection occurs before intrusion inspection. That is, the system does not inspect files blocked by a file policy for intrusions. Within file inspection, simple blocking by type takes precedence over malware inspection and blocking. Until a file is detected and blocked in a session, packets from the session may be subject to intrusion inspection.



Note By default, the system disables intrusion and file inspection of encrypted payloads. This helps reduce false positives and improve performance when an encrypted connection matches an access control rule that has intrusion and file inspection configured. Inspection works with unencrypted traffic only.

Best Practices for Access Control Rule Order

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic. Consider the following recommendations:

- Specific rules should come before general rules, especially when the specific rules are exceptions to general rules.
- Any rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number) should come as early as possible. We recommend they come before any rule that requires inspection, such as those with application or URL criteria, because Layer-3/4 criteria can be evaluated quickly and without inspection. Of course, any exceptions to these rules must be placed above them.
- Whenever possible, put specific drop rules near the top of the policy. This ensures the earliest possible decision on undesirable traffic.

- Any rules that include both application and URL criteria should come after straight-forward application-only or URL-only rules, unless the application+URL rule is acting as an exception to a more general application-only or URL-only rule. Combining application and URL criteria can lead to unexpected results, especially for encrypted traffic, so we recommend that you create separate rules for URL and application filtering whenever possible.

NAT and Access Rules

Access rules always use the real IP addresses when determining an access rule match, even if you configure NAT. For example, if you configure NAT for an inside server, 10.1.1.5, so that it has a publicly routable IP address on the outside, 209.165.201.5, then the access rule to allow the outside traffic to access the inside server needs to reference the server's real IP address (10.1.1.5), and not the mapped address (209.165.201.5).

How Other Security Policies Impact Access Control

Other security policies can affect how access control rules function and match connections. As you configure your access rules, keep the following in mind:

- **SSL Decryption** policy—The SSL decryption rules are evaluated before access control. Thus, if an encrypted connection matches an SSL decryption rule that applies some type of decryption, it is the plain text (decrypted) connection that is evaluated by the access control policy. The access rules do not see the encrypted version of the connection. Additionally, any connections that match SSL decryption rules that drop traffic are never seen by the access control policy. Finally, any encrypted connection that matches a Do Not Decrypt rule is evaluated in its encrypted state.
- **Identity** policy—Connections are matched to users (and thus, user groups) only if there is a user mapping for the source IP address. Access rules that key on user or group membership can match only those connections for which user identity was successfully collected by your identity policy.
- **Security Intelligence** policy—Any connection that is dropped is never seen by the access control policy. Connections that match the do not block list are subsequently matched to access control rules and, ultimately, it is the access control rule that determines how the connection is handled (allowed or dropped).
- **VPN** (site-to-site or remote access)—VPN traffic is always evaluated against the access control policy, and connections are allowed or dropped based on the matching rule. However, the VPN tunnel itself is decrypted before the access control policy is evaluated. The access control policy evaluates the connections that are embedded within the VPN tunnel, not the tunnel itself.

License Requirements for Access Control

You do not need a special license to use the access control policy.

However, you do need the following licenses for specific features within the access control policy. For information on configuring licenses, see Enabling or Disabling Optional Licenses.

- **URL** license—To create rules that use URL categories and reputations as match criteria.
- **IPS** license—To configure an intrusion policy on an access rule or the default action. You also need this license to use a file policy (the Malware Defense license is also required).

- **Malware Defense** license—To configure a file policy on an access rule. The IPS is also required for file policies.

Guidelines and Limitations for Access Control Policies

Following are some additional limitations for access control. Please consider them when evaluating whether you are getting the expected results from your rules.

- If a URL database update includes added (new, incoming), deprecated (outgoing), or deleted categories, there is a grace period for you to make changes to any access control rules that are affected. Impacted rules are marked with informational messages, with explanations about the issues that impact the rule and links to the Cisco Talos Intelligence Group (Talos) web site for more information about the category changes. You need to update the rule so that it uses the appropriate categories available in the latest URL database.

  To accommodate the grace period, add the newly added incoming categories to the appropriate rules without removing the outgoing deprecated categories: your rules should contain both the new and old categories. The new categories will be effective when the old categories are marked for deletion. When the old categories are finally deleted, you need to edit the rules to remove the deleted categories and redeploy the configuration. You will be blocked from deploying the configuration until you fix any rules that use deleted categories. Click the **See Problem Rules** link above the table to filter on rules that need attention.

- Firewall Device Manager can download information on up to 50,000 users from the directory server. If your directory server includes more than this number of user accounts, you will not see all possible names when selecting users in an access rule or when viewing user-based dashboard information. You can write rules on only those names that were downloaded.

  The limit also applies to the names associated with groups. If a group has more than the limit of members, only the first downloaded names up to the limit can be matched against the group membership.

- If a Vulnerability Database (VDB) update removes (deprecates) applications, you must make changes to any access control rules or application filters that use the application that was deleted. You cannot deploy changes until you fix these rules. In addition, you cannot install system software updates before fixing the issue. On the Application Filters object page, or the Application tab of the rule, these applications say “(Deprecated)” after the application name.

- To use fully-qualified domain name (FQDN) network objects as source or destination criteria, you must also configure DNS for the data interfaces on **Device** > **System Settings** > **DNS Server**. The system does not use the management DNS server setting to do lookups for FQDN objects used in access control rules. For information on troubleshooting FQDN resolution, see Troubleshooting General DNS Problems.

  Note that controlling access by FQDN is a best-effort mechanism. Consider the following points:

  - Because DNS replies can be spoofed, only use fully trusted internal DNS servers.
  - Some FQDNs, especially for very popular servers, can have hundreds if not thousands of IP addresses, and these can frequently change. Because the system uses cached DNS lookup results, users might get addresses that are not yet in the cache, and their connections will not match the FQDN rule. Rules that use FQDN network objects function effectively only for names that resolve to fewer than 100 addresses.

    We recommend that you do not create network object rules for an FQDN that resolves to more than 100 addresses, as the likelihood of the address in a connection being one that has been resolved and

available in the DNS cache on the device is low. For these cases, use a URL-based rule instead of an FQDN network object rule.

- For popular FQDNs, different DNS servers can return a different set of IP addresses. Thus, if your users use a different DNS server than the one you configure, FQDN-based access control rules might not apply to all IP addresses for the site that are used by your clients, and you will not get the intended results for your rules.
- Some FQDN DNS entries have very short time to live (TTL) values. This can result in frequent recompliation of the lookup table, which can impact overall system performance.

- If you edit a rule that is actively in use, the changes do not apply to established connections that are no longer being inspected by Snort. The new rule is used to match against future connections. In addition, if Snort is actively inspecting a connection, it can apply the changed matching or action criteria to an existing connection. If you need to ensure that your changes apply to all current connections, you can log into the device CLI and use the **clear conn** command to end established connections, on the assumption that the sources for the connections will then attempt to reestablish the connection and thus be matched appropriately against the new rule.
- It takes 3 to 5 packets for the system to identify the application or URL in a connection. Thus, the correct access control rule might not be matched immediately for a given connection. However, once the application/URL is known, the connection is handled based on the matching rule. For encrypted connections, this happens after the server certificate exchange in the SSL handshake.
- The system applies the default policy action to packets that do not have a payload in a connection where an application is identified.
- Leave matching criteria empty whenever possible, especially those for security zones, network objects, and port objects. For example, the system can more efficiently match traffic for all interfaces if you simply leave the security zone criteria blank, rather than if you create zones that contain all interfaces. When you specify multiple criteria, the system must match against every combination of the contents of the criteria you specify.
- If you specify IP addresses for source or destination criteria, do not mix IPv4 and IPv6 addresses in the same rule. Create separate rules for IPv4 and IPv6 addresses.
- While operating, the Firewall Threat Defense device expands access control rules into multiple access control list entries based on the contents of any network objects used in the access rule. You can reduce the memory required to search access control rules by enabling object group search. With object group search enabled, the system does not expand network objects, but instead searches access rules for matches based on those group definitions. Object group search does not impact how your access rules are defined or how they appear in the Firewall Device Manager. It impacts only how the device interprets and processes them while matching connections to access control rules.

  Enabling object group search reduces memory requirements for access control policies that include network objects. However, it is important to note that object group search might also decrease rule lookup performance and thus increase CPU utilization. You should balance the CPU impact against the reduced memory requirements for your specific access control policy. In most cases, enabling object group search provides a net operational improvement.

  You can set this option using FlexConfig by issuing the **object-group-search access-control** command; use the **no** form of the command in the negate template.

  Starting with release 7.2, this feature is enabled by default on new deployments, but is not automatically enabled on upgraded systems.

- GRE tunnels that violate the related RFCs will be dropped. For example, if a GRE tunnel contains non-zero values in the reserved bits, contrary to the RFCs, it is dropped. If you need to allow non-compliant GRE tunnels, you need to use a remote manager and configure a prefilter rule that trusts the sessions. You cannot configure prefilter rules using the Firewall Device Manager.

Configuring the Access Control Policy

Use the access control policy to control access to network resources. The policy consists of a set of ordered rules, which are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched. If no rules match the traffic, the default action shown at the bottom of the page is applied.

To configure the access control policy, select **Policies** > **Access Control**.

The access control table lists all rules in order. For each rule:

- Click the > button next to the rule number in the left-most column to open the rule diagram. The diagram can help you visualize how the rule controls traffic. Click the button again to close the diagram.
- Most cells allow inline editing. For example, you can click the action to select a different one, or click a source network object to add or change the source criteria.
- To move a rule, hover over the rule until you get the move icon (), then click, drag, and drop the rule to the new location. You can also move a rule by editing it and selecting the new location in the **Order** list. It is critical that you put the rules in the order that you want them processed. Specific rules should be near the top, especially for rules that define exceptions to more general rules
- The right-most column contains the action buttons for a rule; mouse over the cell to see the buttons. You can edit () or delete () a rule.
- Click the **Access Control Settings** () button to configure settings that apply to the access control policy, rather than to specific rules within the policy.
- Click the **Toggle Hit Counts** icon () above the table to add or remove the Hit Counts column in the table. The Hit Count column appears to the right of the Name column with the total hit count for the rule and the date and time of the last hit. The hit count information is fetched at the time you click the toggle button. Click the **refresh** icon () to get the latest information.
- If any rules have problems, for example, because of removed or changed URL categories, click the **See Problem Rules** link next to the search box to filter the table to show only those rules. Please edit and correct (or delete) these rules, so that they will provide the service that you require.

The following topics explain how to configure the policy.

Configuring the Default Action

If a connection does not match a specific access rule, it is handled by the default action for the access control policy.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Click anywhere in the **Default Action** field.

Step 3 Select the action to apply to matching traffic.

- **Trust**—Allow traffic without further inspection of any kind.
- **Allow**—Allow the traffic subject to the intrusion policy.
- **Block**—Drop the traffic unconditionally. The traffic is not inspected.

Step 4 If the action is **Allow**, select an intrusion policy.

For an explanation of the policy options, see Intrusion Policy Settings, on page 26.

Step 5 (Optional.) Configure logging for the default action.

You must enable logging for traffic that matches the default action to be included in dashboard data or Event Viewer. See Logging Settings, on page 27.

Step 6 Click **OK**.

Configuring Access Control Policy Settings

You can configure settings that apply to the access control policy, rather than to specific rules within the policy.

Procedure

Step 1 Select **Policy** > **Access Control**.

Step 2 Click the **Access Policy Settings** ( ) button.

Step 3 Configure the settings:

- **TLS Server Identity Discovery**—TLS 1.3 encrpyts most handshake messages, so certificate information is not readily available. For traffic encrypted with TLS 1.3 to match access rules that use application or URL filtering, the system must have a clear-text certificate for the server. If you enable this option, the system checks whether a certificate for the site is stored in cache based on the IP address and Server Name Indication (SNI) in the client hello packet. If one is not available, the system uses a TLS 1.2 probe to obtain the certicate, which can then be used for application/URL category and reputation identification. We recommend that you enable this option to ensure encrypted connections are matched to the right access control rule. The setting obtains the certificate only; the connection remains encrypted. Enabling this option is sufficient to obtain TLS 1.3 certificates; you do not need to create a corresponding SSL decryption rule. However, the cached certificates are also used for more effective decryption rule processing in addition to access control processing.
- **Reputation Enforcement on DNS Traffic**—Enable this option to apply your URL filtering category and reputation rules to DNS lookup requests. If the fully-qualified domain name (FQDN) in the lookup request has a category and reputation that you are blocking, the system blocks the DNS reply. Because

the user does not receive a DNS resolution, the user cannot complete the connection. Use this option to apply URL category and reputation filtering to non-web traffic. For more information, see DNS Request Filtering, on page 11.

Step 4 Click **OK**.

Configuring Access Control Rules

Use access control rules to control access to network resources. Rules in the access control policy are evaluated from top to bottom. The rule applied to traffic is the first one where all the traffic criteria are matched.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Do any of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

To delete a rule you no longer need, click the delete icon ( ) for the rule.

Step 3 In **Order**, select where you want to insert the rule in the ordered list of rules.

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic.

The default is to add the rule to the end of the list. If you want to change a rule's location later, edit this option.

Step 4 In **Title**, enter a name for the rule.

The name cannot contain spaces. You can use alphanumeric characters and these special characters: + . _ -

Step 5 Select the action to apply to matching traffic.

- **Trust**—Allow traffic without further inspection of any kind.
- **Allow**—Allow the traffic subject to the intrusion and other inspection settings in the policy.
- **Block**—Drop the traffic unconditionally. The traffic is not inspected.

Step 6 Define the traffic matching criteria using any combination of the following tabs:

- **Source/Destination**—The security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, the security group tags (SGT) assigned to the address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, SGT, protocol, and port. See Source/Destination Criteria, on page 20.
- **Application**—The application, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any application. See Application Criteria, on page 22.
- **URL**—The URL or URL category of a web or DNS lookup request. The default is any URL. See URL Criteria, on page 24.

- **Users**—The identity source, user or user group. Your identity policies determine whether user and group information is available for traffic matching. You must configure identity policies to use this criteria. See User Criteria, on page 25.

To modify a condition, you click the **+** button within that condition, select the desired object or element, and click **OK** in the popup dialog box. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

When adding conditions to access control rules, consider the following tips:

- You can configure multiple conditions per rule. Traffic must match all the conditions in the rule for the rule to apply to traffic. For example, you can use a single rule to perform URL filtering for specific hosts or networks.
- For each condition in a rule, you can add up to 50 criteria. Traffic that matches any of a condition's criteria satisfies the condition. For example, you can use a single rule to apply application control for up to 50 applications or application filters. Thus, there is an OR relationship among the items in a single condition, but an AND relationship between condition types (for example, between source/destination and application).
- Some features require that you enable the appropriate license.

Step 7 (Optional.) For policies that use the Allow action, you can configure further inspection on unencrypted traffic. Click one of the following links:

- **Intrusion Policy**—Select **Intrusion Policy** > **On** and select the intrusion inspection policy to inspect traffic for intrusions and exploits. See Intrusion Policy Settings, on page 26.
- **File Policy**—Select the file policy to inspect traffic for files that contain malware and for files that should be blocked. See File Policy Settings, on page 26.

Step 8 (Optional.) Configure logging for the rule.

By default, connection events are not generated for traffic that matches a rule, although file events are generated by default if you select a file policy. You can change this behavior. You must enable logging for traffic that matches the policy to be included in dashboard data or Event Viewer. See Logging Settings, on page 27.

Intrusion events are always generated for intrusion rules set to drop or alert regardless of the logging configuration on the matching access rule.

Step 9 Click **OK**.

Source/Destination Criteria

The Source/Destination criteria of an access rule define the security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, the security group tags (SGT) assigned to the address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, SGT, protocol, and port.

To modify a condition, you click the **+** button within that condition, select the desired object or element, and click **OK**. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

You can use the following criteria to identify the source and destination to match in the rule.

Source Zones, Destination Zones

The security zone objects that define the interfaces through which the traffic passes. You can define one, both, or neither criteria: any criteria not specified applies to traffic on any interface.

- To match traffic leaving the device from an interface in the zone, add that zone to the **Destination Zones**.
- To match traffic entering the device from an interface in the zone, add that zone to the **Source Zones**.
- If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones and egress through one of the destination zones.

Use this criteria when the rule should apply based on where the traffic enters or exits the device. For example, if you want to ensure that all traffic going to inside hosts gets intrusion inspection, you would select your inside zone as the **Destination Zones** while leaving the source zone empty. To implement intrusion filtering in the rule, the rule action must be **Allow**, and you must select an intrusion policy in the rule.



Note You cannot mix passive and routed security zones in a single rule. In addition, you can specify passive security zones as source zones only, you cannot specify them as destination zones.

Source Networks, Destination Networks

The network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, configure the **Source Networks**.
- To match traffic to an IP address or geographical location, configure the **Destination Networks**.
- If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control. You can use objects that define the address using the fully-qualified domain name (FQDN); the address is determined through a DNS lookup.
- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.



Note To ensure that you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Source Ports, Destination Ports/Protocols

The port objects that define the protocols used in the traffic. For TCP/UDP, this can include ports. For ICMP, it can include codes and types.

- To match traffic from a protocol or port, configure the **Source Ports**. Source ports can be TCP/UDP only.
- To match traffic to a protocol or port, configure the **Destination Ports/Protocols**. If you add only destination ports to a condition, you can add ports that use different transport protocols. ICMP and other non-TCP/UDP specifications are allowed in destination ports only; they are not allowed in source ports.
- To match traffic both originating from specific TCP/UDP ports and destined for specific TCP/UDP ports, configure both. If you add both source and destination ports to a condition, you can only add ports that share a single transport protocol, TCP or UDP. For example, you could target traffic from port TCP/80 to port TCP/8080.

Source SGT Groups, Destination SGT Groups

The Security Group Tag (SGT) group objects that identify the SGTs assigned to the traffic, as downloaded from the Identity Services Engine (ISE). You can use these objects only if you define an ISE identity source; otherwise, this section will not appear. For detailed information on how to use SGTs for access control, see How to Control Network Access Using TrustSec Security Group Tags, on page 31.

- To match traffic whose source has one of the SGT defined in the group, configure the **Source SGT Groups**.
- To match traffic to a destination that has one of the SGT defined in the group, configure the **Destination SGT Groups**.
- If you add both source and destination SGT conditions to a rule, matching traffic must originate from a source with one of the specified tags and be destined for one of the destination tags.

Application Criteria

The Application criteria of an access rule defines the application used in an IP connection, or a filter that defines applications by type, category, tag, risk, or business relevance. The default is any application.

Although you can specify individual applications in the rule, application filters simplify policy creation and administration. 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.

In addition, 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.

You can specify applications and filters directly in the rule, or create application filter objects that define those characteristics. The specifications are equivalent, although using objects can make it easier to stay within the 50-items-per-criteria system limit if you are creating a complex rule.

To modify the application and filters list, you click the + button within the condition, select the desired applications or application filter objects, which are listed on separate tabs, and click **OK** in the popup dialog box. On either tab, you can click **Advanced Filter** to select filter criteria or to help you search for specific applications. Click the **x** for an application, filter, or object to remove it from the policy. Click the **Save As Filter** link to save the combined criteria that is not already an object as a new application filter object.

Note If a selected application was removed by a VDB update, "(Deprecated)" appears after the application name. You must remove these applications from the filter, or subsequent deployments and system software upgrades will be blocked.

You can use the following **Advanced Filter** criteria to identify the application or filter to match in the rule. These are the same elements used in application filter objects.

Note Multiple selections within a single filter criteria have an OR relationship. For example, Risk is High OR Very High. The relationship between filters is AND, so Risk is High OR Very High, AND Business Relevance is Low OR Very Low. As you select filters, the list of applications in the display updates to show only those that meet the criteria. You can use these filters to help you find applications that you want to add individually, or to verify that you are selecting the desired filters to add to the rule.

Risks

The likelihood that the application is used for purposes that might be against your organization's security policy, from very low to very high.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally, from very low to very high.

Types

The type of application:

- **Application Protocol**—Application protocols such as HTTP and SSH, which represent communications between hosts.
- **Client Protocol**—Clients such as web browsers and email clients, which represent software running on the host.
- **Web Application**—Web applications such as MPEG video and Facebook, which represent the content or requested URL for HTTP traffic.

Categories

A general classification for the application that describes its most essential function.

Tags

Additional information about the application, similar to category.

For encrypted traffic, the system can identify and filter traffic using only the applications tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic. Also, the system assigns the **decrypted traffic** tag to applications that the system can detect in decrypted traffic only, not encrypted or unencrypted.

Applications List (bottom of the display)

This list updates as you select filters from the options above the list, so you can see the applications that currently match the filter. Use this list to verify that your filter is targeting the desired applications when

you intend to add filter criteria to the rule. If your intention is to add specific applications, select them from this list.

URL Criteria

The URL criteria of an access rule defines the URL used in a web request, or the category to which the requested URL belongs. For category matches, you can also specify the relative reputation of sites to allow or block. The default is to allow all URLs.

If you enable DNS lookup request filtering, the category and reputation settings also apply to the fully-qualified domain name (FQDN) in the lookup request. Only the category and reputation settings apply for DNS request filtering. Manual URL filtering is ignored.

URL categories and reputations allow you to quickly create URL conditions for access control rules. For example, you could block all Gambling sites, or untrusted Social Networking sites. 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.

To modify the URL list, you click the **+** button within the condition and select the desired categories or URLs using one of the following techniques. Click the **x** for a category or object to remove it from the policy.

URL Tab

Click **+**, select URL objects or groups, and click **OK**. You can click **Create New URL** if the object you require does not exist.



Note Before configuring URL objects to target specific sites, carefully read the information on manual URL filtering.

Categories Tab

Click **+**, select the desired categories, and click **OK**.

For a description of the categories, see https://www.talosintelligence.com/categories.

The default is to apply the rule to all URLs in each selected category regardless of reputation. To limit the rule based on reputation, click the down arrow for each category, deselect the **Any** checkbox, and then use the **Reputation** slider to choose the reputation level. The left of the reputation slider indicates sites that will be allowed, the right side are sites that will be blocked. How reputation is used depends on the rule action:

- If the rule blocks or monitors web access, selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to block or monitor **Questionable sites** (level 2), it also automatically blocks or monitors **Untrusted** (level 1) sites.
- If the rule allows web access, selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to allow **Favorable sites** (level 4), it also automatically allows **Trusted** (level 5) sites.

Select the **Include Sites with Unknown Reputation** option to have URLs with unknown reputation included in the reputation match. New sites typically are unrated, and there can be other reasons a site's reputation is unknown or cannot be determined.

Check the Category for a URL

You can check on the category and reputation for a particular URL. Enter the URL in the **URL to Check** box and click **Go**. You will be taken to an external website to see the results. If you disagree with a categorization, click the **Submit a URL Category Dispute** link and let us know.

User Criteria

The User criteria of an access rule defines the user or user group for an IP connection. You must configure identity policies and the associated directory server to include user or user group criteria in an access rule.

Your identity policies determine whether user identity is collected for a particular connection. If identity is established, the IP address of the host is associated with the identified user. Thus, traffic whose source IP address is mapped to a user is considered to be from that user. IP packets themselves do not include user identity information, so this IP-address-to-user mapping is the best approximation available.

Because you can add a maximum of 50 users or groups to a rule, selecting groups usually makes more sense than selecting individual users. For example, you could create a rule allowing the Engineering group access to a development network, and create a subsequent rule that denies all other access to the network. Then, to make the rule apply to new engineers, you only need to add the engineer to the Engineering group in the directory server.

You an also select identity sources to apply to all users within that source. Thus, if you support multiple Active Directory domains, you can provide differential access to resources based on the domain.

To modify the users list, you click the **+** button within the condition and select the desired identities using one of the following techniques. Click the **x** for an identity to remove it from the policy.

- **Identity Sources**—Select an identity source, such as an AD realm or the local user database, to apply the rule to all users obtained from the selected sources. If the realm you need does not yet exist, click **Create New Identity Realm** and create it now.
- **Groups**—Select the desired user groups. Groups are available only if you configure groups in the directory server. If you select a group, the rule applies to any member of the group, including subgroups. If you want to treat a sub-group differently, you need to create a separate access rule for the sub-group and place it above the rule for the parent group in the access control policy.
- **Users**—Select individual users. The user name is prefixed with the identity source, such as Realm\username.

  There are some built-in users under the Special-Identities-Realm:

  - **Failed Authentication**—The user was prompted to authenticate, but failed to enter a valid username/password pair within the maximum number of allowed attempts. Failure to authenticate does not itself prevent the user from accessing the network, but you can write an access rule to limit network access for these users.
  - **Guest**—Guest users are like Failed Authentication users, except that your identity rule is configured to call these users Guest. Guest users were prompted to authenticate and failed to do so within the maximum number of attempts.
  - **No Authentication Required**—The user was not prompted to authentication, because the user's connections matched identity rules that specified no authentication.

- **Unknown**—There is no user mapping for the IP address, and there is no record of failed authentication yet. Typically, this means that no HTTP traffic has yet been seen from that address.

Intrusion Policy Settings

Cisco delivers several intrusion policies with the firewall system. The several intrusion policies delivered by Cisco Cisco Talos Intelligence Group (Talos) are designed by the Cisco.Talos set the intrusion and preprocessor rule states and advanced settings. For access control rules that allow traffic, you can select an intrusion policy to inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic.

When running Snort 2, these are the only policies available and you cannot modify them. However, you can change the action to take for a given rule, as described in Changing Intrusion Rule Actions (Snort 2).

When running Snort 3, you can select one of these policies, or you can create your own intrusion policies.

To enable intrusion inspection, select **Intrusion Policy** > **On** and select the desired policy. Click the info icon for a policy in the drop-down list to see a description for each policy.

For more information on the pre-defined policies, see System-Defined Network Analysis and Intrusion Policies.

File Policy Settings

Use file policies to detect malicious software, or *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 Secure Malware Analytics Cloud to retrieve dispositions for possible malware detected in network traffic, and to obtain local malware analysis and file pre-classification updates. The management interface must have a path to the Internet to reach the Secure Malware Analytics Cloud and perform malware lookups. When the device detects an eligible file, it uses the file's SHA-256 hash value to query the Secure Malware Analytics Cloud for the file's disposition. The possible dispositions are:

- Malware—The Secure Malware Analytics Cloud categorized the file as malware. An archive file (e.g. a zip file) is marked as malware if any file within it is malware.
- Clean—The Secure Malware Analytics Cloud categorized the file as clean, containing no malware. An archive file is marked as clean if all files within it are clean.
- Unknown—The Secure Malware Analytics Cloud has not assigned a disposition to the file yet. An archive file is marked as unknown if any file within it is unknown.
- Unavailable—The system could not query the Secure Malware Analytics Cloud to determine the file's disposition. You may see a small percentage of events with this disposition; this is expected behavior. If you see a number of "unavailable" events in succession, ensure that the Internet connection for the management address is functioning correctly.

Available File Policies

You can select one of the following file policies:

- **None**—Do not evaluate transmitted files for malware and do no file-specific blocking. Select this option for rules where file transmissions are trusted or where they are unlikely (or impossible), or for rules where you are confident your application or URL filtering adequately protects your network.

- **Block Malware All**—Query the Secure Malware Analytics Cloud to determine if files traversing your network contain malware, then block files that represent threats.
- **Cloud Lookup All**—Query the Secure Malware Analytics Cloudto obtain and log the disposition of files traversing your network while still allowing their transmission.
- **(Custom File Policy)**—You can create your own file policies using the Firewall Threat Defense API filepolicies resource, and the other FileAndMalwarePolicies resources (such as filetypes, filetypecategories, ampcloudconfig, ampservers, and ampcloudconnections). After you create the policies, and deploy changes, you can select your policies when editing an access control rule in the Firewall Device Manager. The policy description appears below the policy when you select it.

Logging Settings

The logging settings for an access rule determine whether connection events are issued for traffic that matches the rule. You must enable logging to see events related to the rule in the Event Viewer. You must also enable logging for matching traffic to be reflected in the various dashboards you can use to monitor the system.

You should log connections according to the security and compliance needs of your organization. If your goal is to limit the number of events you generate and improve performance, only enable logging for the connections critical to your analysis. However, if you want a broad view of your network traffic for profiling purposes, you can enable logging for additional connections.



Caution Logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance and overwhelm the database with multiple similar events. Before you enable logging for a Block rule, consider whether the rule is for an Internet-facing interface or other interface vulnerable to DoS attack.

You can configure the following logging actions.

Select Log Action

You can select one of the following actions:

- **Log at Beginning and End of Connection**—Issue events at the start and end of a connection. Because end-of-connection events contain everything that start-of-connection events contain, plus all of the information that could be gleaned during the connection, Cisco recommends that you do not select this option for traffic that you are allowing. Logging both events can impact system performance. However, this is the only option allowed for blocked traffic.
- **Log at End of Connection**—Select this option if you want to enable connection logging at the end of the connection, which is recommended for allowed or trusted traffic.
- **No Logging at Connection**—Select this option to disable logging for the rule. This is the default.



Note When an intrusion policy invoked by an access control rule detects an intrusion and generates an intrusion event, the system automatically logs the end of the connection where the intrusion occurred, regardless of the logging configuration of the rule. For connections where an intrusion was blocked, the action for the connection in the connection log is **Block**, with a reason of **Intrusion Block**, even though to perform intrusion inspection you must use an Allow rule.

File Events

Select **Log Files** if you want to enable logging of prohibited files or malware events. You must select a file policy in the rule to configure this option. The option is enabled by default if you select a file policy for the rule. Cisco recommends you leave this option enabled.

When the system detects a prohibited file, it automatically logs one of the following types of event:

- *File events*, which represent detected or blocked files, including malware files.
- *Malware events*, which represent detected or blocked malware files only.
- *Retrospective malware events*, which are generated when the malware disposition for a previously detected file changes.

For connections where a file was blocked, the action for the connection in the connection log is **Block** even though to perform file and malware inspection you must use an Allow rule. The connection's Reason is either **File Monitor** (a file type or malware was detected), or **Malware Block** or **File Block** (a file was blocked).

Send Connection Events To

If you want to send a copy of the events to an external syslog server, select the server object that defines the syslog server. If the required object does not already exist, click **Create New Syslog Server** and create it. (To disable logging to a syslog server, select **Any** from the server list.)

Because event storage on the device is limited, sending events to an external syslog server can provide more long term storage and enhance your event analysis.

This setting applies to connection events only. To send intrusion events to syslog, configure the server in the intrusion policy settings. To send file/malware events to syslog, configure the server on **Device** > **System Settings** > **Logging Settings**.

Monitoring Access Control Policies

The following topics explain how you can monitor the access control policy.

Monitoring Access Control Statistics in the Dashboards

Most of the data on the **Monitoring** dashboards are directly related to your access control policy. See Monitoring Traffic and System Dashboards.

- **Monitoring** > **Access And SI Rules** shows the most-hit access rules and Security Intelligence rule-equivalents and related statistics.
- You can find general statistics on the **Network Overview**, **Destinations**, and **Zones**, dashboards.
- You can find URL filtering results on the **URL Categories** and **Destinations** dashboards. You must have at least one URL filtering policy to see any information on the **URL Categories** dashboard.
- You can find application filtering results on the **Applications** and **Web Applications** dashboards.
- You can find user-based statistics on the **Users** dashboard. You must implement identity policies to collect user information.

- You can find intrusion policy statistics on the **Attackers** and **Targets** dashboards. You must apply an intrusion policy to at least one access control rule to see any information on these dashboards.
- You can find file policy and malware filtering statistics on the **File Logs** and **Malware** dashboards. You must apply a file policy to at least one access control rule to see any information on these dashboards.
- **Monitoring** > **Events** also shows events for connections and data related to the access control rules.

Examining Rule Hit Counts

You can view the hit count for each access control rule. The hit count indicates how often connections matched the rule. You can use this information to identify your most active rules and the rules that are less active.

The count persists through reboots and upgrades.

You can also see rule hit count information in the device CLI using the **show rule hits** command.

Procedure

Step 1 Select **Policies** > **Access Control**.

Step 2 Click the **Toggle Hit Counts** icon ( ).

The Hit Count column appears to the right of the Name column with the total hit count for the rule and the date and time of the last hit. The hit count information is fetched at the time you click the toggle button.

You can do the following with the hit count information:

- To the left of the button, you will see information on when the hit count was last updated. Click the **refresh** icon ( ) to get the latest numbers.
- To open a detailed view of the hit count for a given rule, click the hit count number in the table to open the Hit Count dialog box. The hit count information includes the number of hits and the date and time of the last connection that matched the rule. Click the **Reset** link to reset the counter to zero.

  If you want to reset the hit count for all rules at once, open an SSH session to the device and issue the **clear rule hits** command.
- Click the **Toggle Hit Counts** icon ( ) again to remove the hit count column from the table.

Monitoring Syslog Messages for Access Control

In addition to seeing events in the Event Viewer, you can configure access control rules, intrusion policies, file/malware policies, and Security Intelligence policies to send events to a syslog server. The events use the following message IDs:

- 430001—Intrusion event.
- 430002—Connection event logged at the beginning of a connection.
- 430003—Connection event logged at the end of a connection.

- 430004—File events.
- 430005—Malware events.

Monitoring Access Control Policies in the CLI

You can also open the CLI console or log into the device CLI and use the following commands to get more detailed information about access control policies and statistics.

- **show access-control-config** displays summary information about the access control rules along with per-rule hit counts.
- **show access-list** displays the access control lists (ACLs) that were generated from the access control rules. The ACLs provide an initial filter and attempt to provide quick decisions whenever possible, so that connections that should be dropped do not need to be inspected (and thus consume resources unnecessarily). This information includes hit counts.
- **show rule hits** displays consolidated hit counts that are more accurate than the counts shown with **show access-control-config** and **show access-list**. If you want to reset the hit count, use the **clear rule hits** command.
- **show snort statistics** displays information about the Snort inspection engine, which is the main inspector. Snort implements application filtering, URL filtering, intrusion protection, and file and malware filtering.
- **show conn** displays information about the connections currently established through the interfaces.
- **show traffic** displays statistics about traffic flowing through each interface.
- **show ipv6 traffic** displays statistics about IPv6 traffic flowing through the device.

Examples for Access Control

The use case chapter includes several examples of implementing access control rules. Please see the following examples:

- How to Gain Insight Into Your Network Traffic. This example shows some basic ideas for collecting overall connection and user information.
- How to Block Threats. This example shows how to apply intrusion policies.
- How to Block Malware. This example shows how to apply file policies.
- How to Implement an Acceptable Use Policy (URL Filtering). This example shows how to perform URL filtering.
- How to Control Application Usage. This example shows how to perform application filtering.
- How to Add a Subnet. This example shows how to integrate a new subnet into your overall network, including the access rules needed to allow traffic flow.
- How to Passively Monitor the Traffic on a Network

Following are additional examples.

How to Control Network Access Using TrustSec Security Group Tags

If you use Cisco Identity Services Engine (ISE) to define and use security group tags (SGT) for classifying traffic in a Cisco TrustSec network, you can write access control rules that use SGT as matching criteria. Thus, you can block or allow access based on security group membership rather than IP addresses directly.

About Security Group Tags (SGT)

In Cisco Identity Services Engine (ISE), you can create security group tags (SGT) and assign host or network IP addresses to each tag. You can also assign SGTs to user accounts, and the SGT is assigned to the user's traffic. If the switches and routers in the network are configured to do so, these tags are then assigned to packets as they enter the network controlled by ISE, the Cisco TrustSec cloud.

When you configure an ISE identity source in the Firewall Device Manager, the Firewall Threat Defense system automatically downloads the list of SGTs from ISE. You can then use SGT as a traffic matching condition in access control rules.

For example, you could create a Production Users tag, and associate the 192.168.7.0/24 network to the tag. This would be appropriate if you use that network for user end points, such as laptops, Wi-Fi clients, and similar devices. You could create a separate tag for Production Servers, and assign the IP addresses of the relevant servers or subnet to the tag. Then, in the Firewall Threat Defense, you could allow or block access from the user network to the production servers based on the tag. If you later alter the host or network addresses associated with the tag in ISE, you do not need to change the access control rule defined for the Firewall Threat Defense device.

When the Firewall Threat Defense evaluates SGT as a traffic matching criteria for an access control rule, it uses the following priority:

1. The source SGT tag defined in the packet, if any. For the SGT tag to be in the packet, the switches and routers in the network must be configured to add them. See the ISE documentation for information on how to implement this method.
2. The SGT assigned to the user session, as downloaded from the ISE session directory. You need to enable the option to listen to session directory information for this kind of SGT matching, but this option is on by default when you first create the ISE identity source. The SGT can be matched to source or destination. Although not required, you would also normally set up a passive authentication identity rule, using the ISE identity source along with an AD realm, to collect user identity information.
3. The SGT-to-IP address mapping downloaded using SXP. If the IP address is within the range for an SGT, then the traffic matches the access control rule that uses the SGT. The SGT can be matched to source or destination.

   ISE uses the Security-group eXchange Protocol (SXP) to propagate the IP-to-SGT mapping database to network devices. When you configure the Firewall Threat Defense device to use an ISE server, you must turn on the option to listen to the SXP topic from ISE. Thus, the Firewall Threat Defense device learns about the security group tags and mappings directly from ISE, and is notified whenever ISE publishes updated security group tags and mappings. This ensures that the list of security group tags and mappings stay up-to-date on the device, so that the Firewall Threat Defense can effectively enforce the policy defined in ISE.

Configure security groups and SXP publishing in ISE

This task enables static SGT-to-IP address mappings for Secure Firewall Threat Defense devices, allowing TrustSec policies and access control based on source and destination security group tags (SGT).

There is a lot of configuration that you must do in Cisco Identity Services Engine (ISE) to create the TrustSec policy and security group tags (SGT). Please look at the ISE documentation for more complete information on implementing TrustSec.

The following procedure picks out the highlights of the core settings you must configure in ISE for the Firewall Threat Defense device to be able to download and apply static SGT-to-IP address mappings, which can then be used for source and destination SGT matching in access control rules. See the ISE documentation for detailed information.

The figures in this procedure are based on ISE 2.4. The exact paths to these features might change in subsequent releases, but the concepts and requirements will be the same. Although ISE 2.4 or later is recommended, and preferably 2.6 or later, the configuration should work starting with ISE 2.2 patch 1.

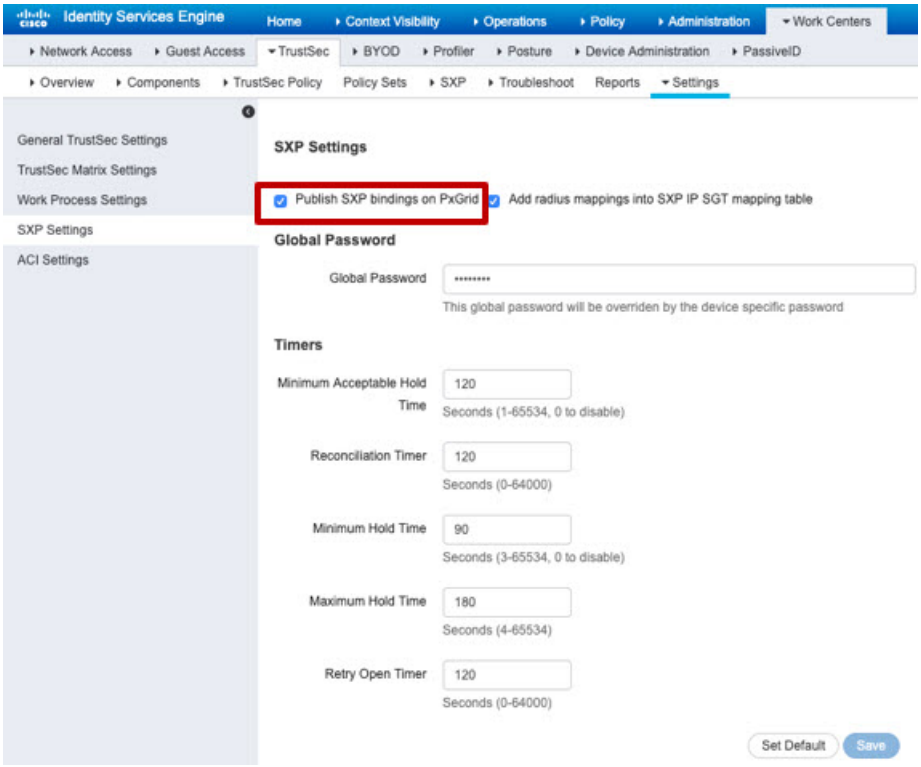
Before you begin

You must have the ISE Plus license to publish SGT-to-IP address static mappings and to get user session-to-SGT mappings so that the Firewall Threat Defense device can receive them.

Procedure

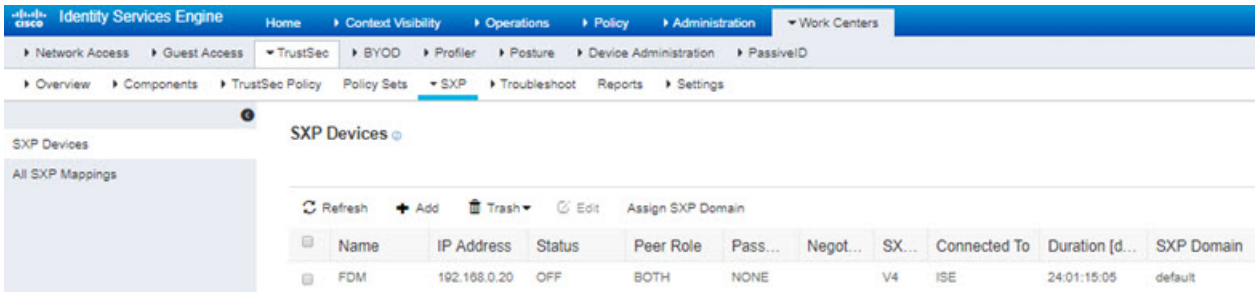
Step 1 Choose **Work Centers** > **TrustSec** > **Settings** > **SXP Settings**, and select the **Publish SXP Bindings on PxGrid** option.

This option makes ISE send the SGT mappings out using SXP. You must select this option for the Firewall Threat Defense device to "hear" anything from listing to the SXP topic. This option must be selected for the Firewall Threat Defense device to get static SGT-to-IP address mapping information. It is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

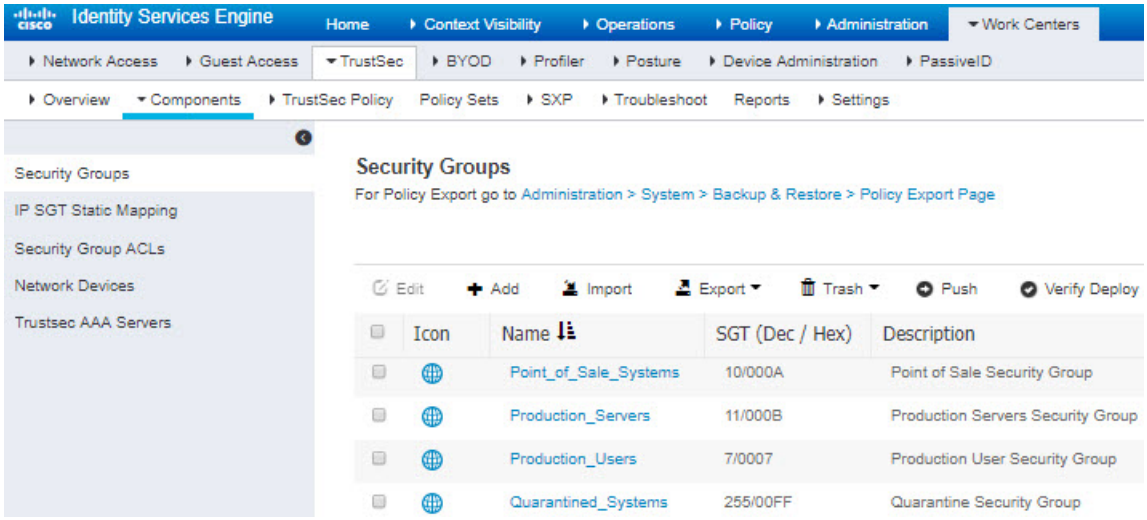


Step 2 Choose **Work Centers** > **TrustSec** > **SXP** > **SXP Devices**, and add a device.

This does not have to be a real device, you can even use the management IP address of the Firewall Threat Defense device. The table simply needs at least one device to induce ISE to publish the static SGT-to-IP address mappings. This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.

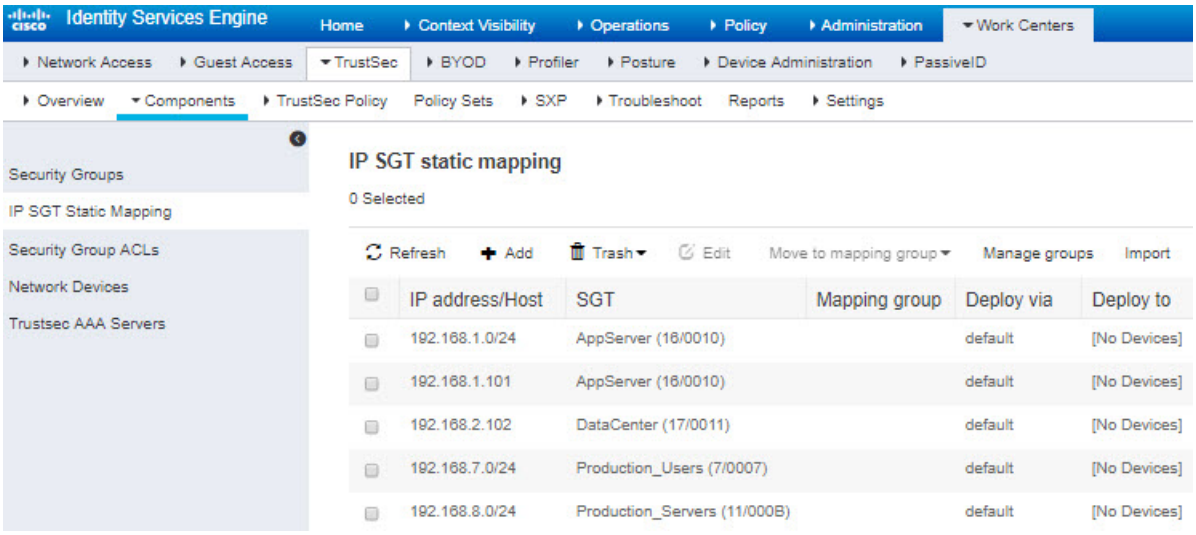


Step 3 Choose **Work Centers** > **TrustSec** > **Components** > **Security Groups** and verify there are security group tags defined. Create new ones as necessary.



Step 4 Choose **Work Centers** > **TrustSec** > **Components** > **IP SGT Static Mapping** and map host and network IP addresses to the security group tags.

This step is not necessary if you simply want to use SGT tags defined in the packets, or SGTs that are assigned to a user session.



You have configured ISE with the necessary security groups, SXP publishing settings, and static SGT-to-IP mappings required for Secure Firewall Threat Defense devices to receive and apply TrustSec policies.

Configure Access Control Based on Security Group Tag (SGT)

To configure access control rules that use security group tags (SGT) as matching criteria, you must first configure the device to obtain the SGT mappings from an ISE server.

The following procedure explains the end-to-end process based on the assumption that you want to get all mappings that are defined in ISE, including SGT-to-IP address mappings published through SXP. Alternatively:

- If you want to use SGT information in the packets only, and not use mappings downloaded from ISE, simply create SGT group dynamic objects and use them as source SGT criteria in access control rules. Note that in this case, you can use SGT tags as a source condition only; these tags will never match destination criteria.
- If you want to use SGT in packets plus user session SGT mappings only, you do not need to turn on the option to subscribe to the SXP topic in the ISE identity source, nor do you need to configure ISE to publish SXP mappings. You can use this information for both source and destination matching conditions.

Before you begin

The assumption here is that you have already configured Cisco TrustSec in your network, and you are simply adding the Firewall Threat Defense device as a policy enforcement point. If you have not deployed Cisco TrustSec, please start with ISE and configure your network, then return to this procedure. Explaining Cisco TrustSec is outside the scope of this document.

Procedure

Step 1 Ensure that SGTs are defined, that ISE is configured correctly to publish the SXP topic, and that any needed static mappings are in place.

See Configure security groups and SXP publishing in ISE, on page 31.

Step 2 Update the Identity Services Engine object to listen to the SXP topic.

You can use ISE to obtain user session SGT mappings, static SGT-to-IP address mappings through SXP, or both. By default, when you configure the ISE identity source, it obtains user session mappings only; you must turn on the option to listen to the SXP topic from ISE.

a) Choose **Objects** > **Identity Sources**.
b) Edit the ISE object. If you have not configured one yet, click **+** > **Identity Services Engine** and see Configure Identity Services Engine.
c) Under **Subscribe To**, select **SXP Topic**.

Ensure that **Session Directory Topic** is also selected if you are using passive authentication or want user-to-SGT mappings.



d) Click **OK**.

Step 3 Deploy your changes and wait for the system to download tags and mappings from ISE.

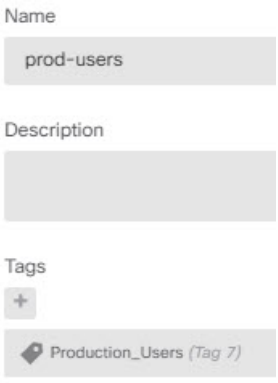
After you configure the ISE identity source and deploy changes, the system retrieves security group tag (SGT) information from the ISE server. The download will not happen until after you deploy changes.

Step 4 Create the SGT group objects that are required for your access control rules.

You cannot use the information retrieved from ISE directly in an access control rule. Instead, you need to create SGT groups, which refer to the downloaded SGT information. Your SGT groups can refer to more than one SGT, so you can apply policy based on relevant collections of tags if that is appropriate.

The number and contents of the objects depend on the access control rules that you intend to write. Repeat the following process to create all the objects that you need.

a) Choose **Objects** > **SGT Groups**.
b) Click **+** to add a new object, or edit an existing object.
c) For new objects, enter a name and optionally a description.
d) Under **Tags**, click **+** and select all the tags that should be included in the group.



e) Click **OK**.

Step 5 Create access control rules that use the SGT group objects.

As an example, the following rule allows traffic from production users to production servers. The rule depends completely on SGT; it is not limited by source/destination interface or any other criteria. Thus, the rule will dynamically apply to traffic as it comes from different interfaces, and as you change security group membership in ISE. If the packet does not explicitly contain a source SGT, source/destination matching will be based on the packet IP addresses compared to the SGT-to-IP address mappings obtained from user session information or from SXP-published mappings.

a) Choose **Policies** > **Access Control**.
b) Click **+** to create a new rule, or edit an existing rule.
c) Enter a rule name and select **Allow** as the action.
d) On the **Source/Destination** tab, click **+** under **Source** > **SGT Groups** and select the object you created for production users.
e) On the **Source/Destination** tab, click **+** under **Destination** > **SGT Groups**, and select the object you created for production servers.
f) Configure other options as needed. For example, you can enable logging and apply an intrusion policy.
g) Click **OK**.

Step 6 Deploy the configuration.