



Access Control Rules

The following topics describe how to configure access control rules:

- [Introduction to Access Control Rules, on page 1](#)
- [Requirements and Prerequisites for Access Control Rules, on page 9](#)
- [Guidelines and Limitations for Access Control Rules, on page 10](#)
- [Managing Access Control Rules, on page 11](#)
- [Examples for Access Control Rules, on page 26](#)
- [History for Access Control Rules, on page 31](#)

Introduction to Access Control Rules

Within an access control policy, *access control rules* provide a granular method of handling network traffic across multiple managed devices.

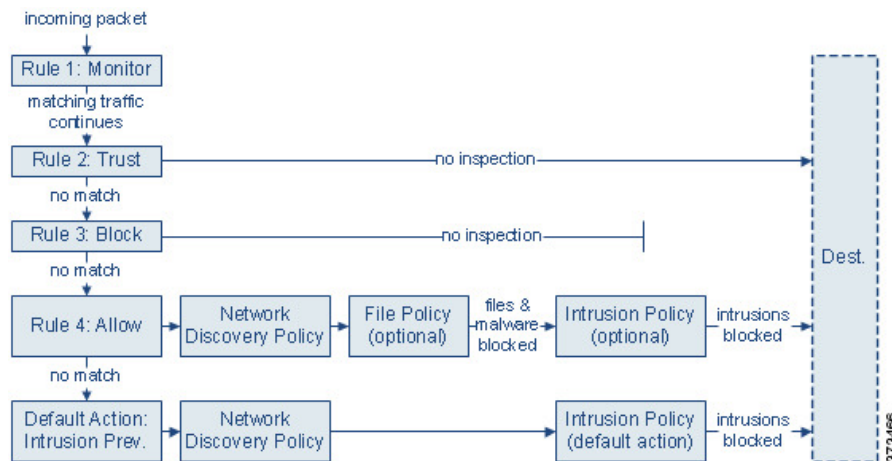


Note Security Intelligence filtering, decryption, user identification, and some decoding and preprocessing occur before access control rules evaluate network traffic.

The system matches traffic to access control rules in the order you specify. In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic.

Each rule also has an *action*, which determines whether you monitor, trust, block, or allow matching traffic. When you allow traffic, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

The following scenario summarizes the ways that traffic can be evaluated by access control rules in an inline, intrusion prevention deployment.



In this scenario, traffic is evaluated as follows:

- **Rule 1: Monitor** evaluates traffic first. Monitor rules track and log network traffic. The system continues to match traffic against additional rules to determine whether to permit or deny it. (However, see an important exception and caveat at [Access Control Rule Monitor Action, on page 6.](#))
- **Rule 2: Trust** evaluates traffic next. Matching traffic is allowed to pass to its destination without further inspection, though it is still subject to identity requirements and rate limiting. Traffic that does not match continues to the next rule.
- **Rule 3: Block** evaluates traffic third. Matching traffic is blocked without further inspection. Traffic that does not match continues to the final rule.
- **Rule 4: Allow** is the final rule. For this rule, matching traffic is allowed; however, prohibited files, malware, intrusions, and exploits within that traffic are detected and blocked. Remaining non-prohibited, non-malicious traffic is allowed to its destination, though it is still subject to identity requirements and rate limiting. You can configure Allow rules that perform only file inspection, or only intrusion inspection, or neither.
- **Default Action** handles all traffic that does not match any of the rules. In this scenario, the default action performs intrusion prevention before allowing non-malicious traffic to pass. In a different deployment, you might have a default action that trusts or blocks all traffic, without further inspection. (You cannot perform file or malware inspection on traffic handled by the default action.)

Traffic you allow, whether with an access control rule or the default action, is automatically eligible for inspection for host, application, and user data by the network discovery policy. You do not explicitly enable discovery, although you can enhance or disable it. However, allowing traffic does not automatically guarantee discovery data collection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy; additionally, application discovery is limited for encrypted sessions.

Note that access control rules handle encrypted traffic when your decryption configuration allows it to pass, or if you do not configure decryption. However, some access control rule conditions require unencrypted traffic, so encrypted traffic may match fewer rules. Also, 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.








Access Control Rule Management

The rules table of the access control policy editor allows you to add, edit, categorize, search, filter, move, enable, disable, delete, and otherwise manage access control rules in the current policy.

Properly creating and ordering access control rules is a complex task, but one that is essential to building an effective deployment. If you do not plan your policy carefully, rules can preempt other rules, require additional licenses, or contain invalid configurations. To help ensure that the system handles traffic as you expect, the access control policy interface has a robust warning and error feedback system for rules.

Use the search bar to filter the list of access control policy rules. You can deselect the **Show Only Matching Rules** option to see all rules. Matched rules are highlighted.

For each access control rule, the policy editor displays its name, a summary of its conditions, the rule action, and icons that communicate the rule's inspection options or status. These icons represent:

- **Time Range Option** ()
- **Intrusion policy** ()
- **File policy** ()
- **Logging** ()
- **Warning** ()
- **Errors** ()
- **Rule Conflict** ()

Disabled rules are dimmed and marked (disabled) after the rule name.

To create or edit a rule, use the access control rule editor.

—You can:

- Configure the rule name and select its placement in the upper portion of the editor.
- Switch to editing a different rule by selecting its row above or below the editor.
- Use the left-hand list to select the rule action, and apply intrusion policies and variable sets, file policies, and time range, and to set logging options.
- Use the options next to the rule name to select the rule action, and apply intrusion policies and variable sets, file policies, and time range, and to set logging options.
- Use the **Sources** and **Destinations and Applications** columns to add matching criteria. You can add options from the All list, or move to different tabs to more easily find the type of option you want, such as security zone or networks.
- Add comments to the rule at the bottom of the editor.



Tip

A right-click menu provides shortcuts to many rule management options, including editing, deleting, moving, enabling, and disabling.

Related Topics

[Access Control Rule Components](#), on page 4
[Best Practices for Access Control Rules](#)

Access Control Rule Components

In addition to its unique name, each access control rule has the following basic components:

State

By default, rules are enabled. If you disable a rule, the system does not use it and stops generating warnings and errors for that rule.

Position

Rules in an access control policy are numbered, starting at 1. If you are using policy inheritance, rule 1 is the first rule in the outermost policy. The system matches traffic to rules in top-down order by ascending rule number. With the exception of Monitor rules, the first rule that traffic matches is the rule that handles that traffic.

Rules can also belong to a section and a category, which are organizational only and do not affect rule position. Rule position goes across sections and categories.

Section and Category

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize access control rules, you can create custom rule categories inside the Mandatory and Default sections.

If you are using policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default sections.

Conditions

Conditions specify the specific traffic the rule handles. Conditions can be simple or complex; their use often depends on license.

Traffic must meet all of the conditions specified in a rule. For example, if the Application condition specifies HTTP but not HTTPS, the URL category and reputation conditions will not apply to HTTPS traffic.

Applicable Time

You can specify days and times during which a rule is applicable.

Action

A rule's action determines how the system handles matching traffic. You can monitor, trust, block, or allow (with or without further inspection) matching traffic. The system does not perform deep inspection on trusted, blocked, or encrypted traffic.

Inspection

Deep inspection options govern how the system inspects and blocks malicious traffic you would otherwise allow. When you allow traffic with a rule, you can specify that the system first inspect it with intrusion or file policies to block any exploits, malware, or prohibited files before they reach your assets or exit your network.

Logging

A rule's logging settings govern the records the system keeps of the traffic it handles. You can keep a record of traffic that matches a rule. In general, you can log sessions at the beginning or end of a connection, or both. You can log connections to the database, as well as to the system log (syslog) or to an SNMP trap server. You can log sessions by application and protocol, then forward these logs to event management solutions for further analysis.

Comments

Each time you save changes to an access control rule, you can add comments.

Related Topics

- [Best Practices for Access Control Rules](#)
- [Access Control Rule Management](#), on page 3
- [Create and Edit Access Control Rules](#), on page 12
- [Access Control Rule Actions](#), on page 6
- [Access Control Rule Conditions](#), on page 14
- [Deep Inspection Using File and Intrusion Policies](#)
- [Access Control Rule Comments](#)

Access Control Rule Order

Rules in an access control policy are numbered, starting at 1. The system matches traffic to access control rules in top-down order by ascending rule number.

In most cases, the system handles network traffic according to the *first* access control rule where *all* the rule's conditions match the traffic. Except for Monitor rules, the system does not continue to evaluate traffic against additional, lower-priority rules after that traffic matches a rule.

To help you organize access control rules, every access control policy has two system-provided rule sections, Mandatory and Default. To further organize, you can create custom rule categories inside the Mandatory or Default sections. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The system assigns rule numbers across sections and categories.

If you use policy inheritance, the current policy's rules are nested between its parent policy's Mandatory and Default rule sections. Rule 1 is the first rule in the outermost policy, not the current policy, and the system assigns rule numbers across policies, sections, and categories.

Any predefined user role that allows you to modify access control policies also allows you to move and modify access control rules within and among rules categories. You can, however, create custom roles that restrict users from moving and modifying rules. Any user who is allowed to modify access control policies can add rules to custom categories and modify rules in them without restriction.

**Caution**

Failure to set up your access control rules properly can have unexpected results, including traffic being allowed that should be blocked. In general, application control rules should be lower in your access control list because it takes longer for those rules to match than rules based on IP address, for example.

Access control rules that use *specific* conditions (such as networks and IP addresses) should be ordered *before* rules that use general conditions (such as applications). If you're familiar with the Open Systems Interconnect (OSI) model, use similar numbering in concept. Rules with conditions for layers 1, 2, and 3 (physical, data link, and network) should be ordered first in your access control rules. Conditions for layers 5, 6, and 7 (session, presentation, and application) should be ordered later in your access control rules. For more information about the OSI model, see this [Wikipedia article](#).

**Tip**

Proper access control rule order reduces the resources required to process network traffic, and prevents rule preemption. Although the rules you create are unique to every organization and deployment, there are a few general guidelines to follow when ordering rules that can optimize performance while still addressing your needs.

Related Topics

[Best Practices for Ordering Rules](#)

Access Control Rule Actions

Every access control rule has an *action* that determines how the system handles and logs matching traffic. You can monitor, trust, block, or allow (with or without further inspection).

The access control policy's *default action* handles traffic that does not meet the conditions of any access control rule with an action other than Monitor.

Access Control Rule Monitor Action

The **Monitor** action is not designed to permit or deny traffic. Rather, its primary purpose is to force connection logging, regardless of how matching traffic is eventually handled.

If a connection matches a Monitor rule, the next non-Monitor rule that the connection matches should determine traffic handling and any further inspection. If there are no additional matching rules, the system should use the default action.

There is an exception, however. If a Monitor rule contains layer 7 conditions—such as an application condition—the system *allows early packets to pass* and the connection to be established (or the SSL handshake to complete). This occurs even if the connection should be blocked by a subsequent rule; this is because these early packets *are not evaluated against subsequent rules*. So that these packets do not reach their destination completely uninspected, you can specify an intrusion policy for this purpose in the access control policy's Advanced settings; see [Inspection of Packets That Pass Before Traffic Is Identified](#). After the system completes its layer 7 identification, it applies the appropriate action to the remaining session traffic.

**Caution**

As a best practice, *avoid placing layer 7 conditions on broadly-defined monitor rules high in your rule priority order*, to prevent inadvertently allowing traffic into your network. Also, if locally bound traffic matches a Monitor rule in a Layer 3 deployment, that traffic may bypass inspection. To ensure inspection of the traffic, enable **Inspect Local Router Traffic** in the advanced device settings for the managed device routing the traffic.

Access Control Rule Trust Action


The **Trust** action allows traffic to pass without deep inspection or network discovery. Trusted traffic is still subject to identity requirements and rate limiting.

**Note**

- Some protocols, such as FTP and SIP, use secondary channels, which the system opens through the process of inspection. In some cases, trusted traffic can bypass all inspection, and these secondary channels cannot be opened properly. If you run into this problem, change the trust rule to **Allow**.
- For trust rules with logging options disabled, end-of-flow events are still generated in the system. However, the events are not visible on the event pages.
- Because access control rules are evaluated after other policies, such as decryption, trusting a connection does not necessarily mean it will be fast-pathed with no inspection. For example, if a connection matches both a decryption rule that requires decryption, and a trust access control rule, the connection is decrypted and inspected as needed, prior to being allowed by the trust rule. Trust simply means no additional inspection, such as intrusion inspection, will be applied. If your intention is to allow a connection uninspected, either use the prefilter policy to fast path the connection, or ensure that no other policy applies inspection services to the connection.

Access Control Rule Blocking Actions

The **Block** and **Block with reset** actions deny traffic without further inspection of any kind.



Block with reset rules reset the connection, with the exception of web requests met with an *HTTP response page*. This is because the response page, which you configure to appear when the system blocks web requests, cannot display if the connection is immediately reset.

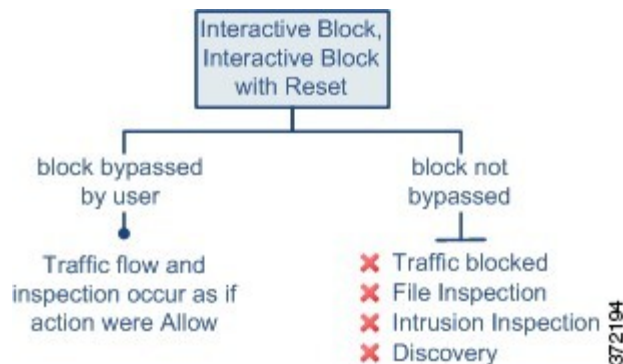
For more information, see [Configure HTTP Response Pages](#).

Related Topics

[Configure HTTP Response Pages](#)

Access Control Rule Interactive Blocking Actions

The **Interactive Block** and **Interactive Block with reset** actions give web users a choice to continue to their intended destinations.



If a user bypasses the block, the rule mimics an allow rule. Therefore, you can associate interactive block rules with file and intrusion policies, and matching traffic is also eligible for network discovery.

If a user does not (or cannot) bypass the block, the rule mimics a block rule. Matching traffic is denied without further inspection.

Note that if you enable interactive blocking, you cannot reset *all* blocked connections. This is because the response page cannot display if the connection is immediately reset. Use the **Interactive Block with reset** action to (non-interactively) block-with-reset all non-web traffic, while still enabling interactive blocking for web requests.

For more information, see [Configure HTTP Response Pages](#).

Related Topics

[Rule-based decryption rule blocking actions](#)

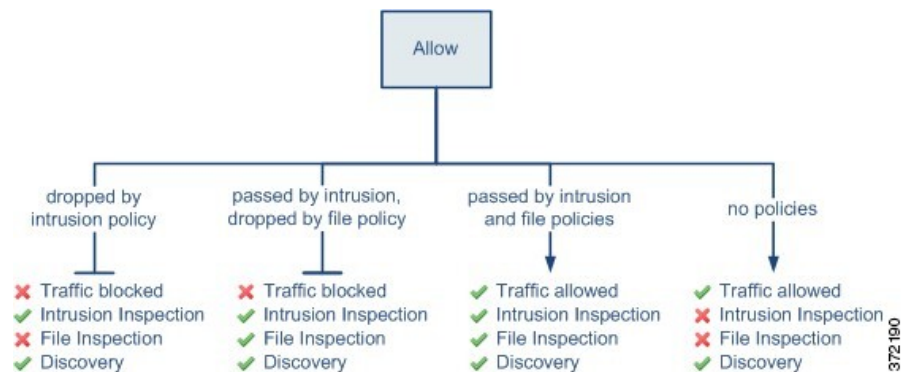
Access Control Rule Allow Action

The **Allow** action allows matching traffic to pass, though it is still subject to identity requirements and rate limiting.

Optionally, you can use deep inspection to further inspect and block unencrypted or decrypted traffic before it reaches its destination:

- You can use an intrusion policy to analyze network traffic according to intrusion detection and prevention configurations, and drop offending packets depending on the configuration.
- You can perform file control using a file policy. File control allows you to detect and block your users from uploading (sending) or downloading (receiving) files of specific types over specific application protocols.
- You can perform network-based advanced malware protection (AMP), also using a file policy. malware defense can inspect files for malware, and block detected malware depending on the configuration.

The following diagram illustrates the types of inspection performed on traffic that meets the conditions of an Allow rule (or a user-bypassed Interactive Block rule. Notice that file inspection occurs before intrusion inspection; blocked files are not inspected for intrusion-related exploits.



For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with an access control rule. You can, however, configure one without the other. Without a file policy, traffic flow is determined by the intrusion policy; without an intrusion policy, traffic flow is determined by the file policy.

Regardless of whether the traffic is inspected or dropped by an intrusion or file policy, the system can inspect it using network discovery. However, allowing traffic does not automatically guarantee discovery inspection. The system performs discovery only for connections involving IP addresses that are explicitly monitored by your network discovery policy; additionally, application discovery is limited for encrypted sessions.

Requirements and Prerequisites for Access Control Rules

Model Support

Any

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin
- You can define custom user roles to differentiate between the intrusion configuration in access control policy and rules and the rest of the access control policy and rules. Using these permissions, you can separate the responsibilities of your network administration team and your intrusion administration teams. The existing pre-defined user roles that included the Modify Access Control Policy permission support all sub-permissions; you need to create your own custom roles if you want to apply granular permissions. The granular permissions are:
 - **Policies > Security policies > Access Control** and choose the **Access Control Policy > Modify Access Control Policy > Modify Threat Configuration** allows the selection of intrusion policy, variable set, and file policy in a rule, the configuration of the advanced options for Network Analysis and Intrusion Policies, the configuration of the Security Intelligence policy for the access control policy, and intrusion actions in the policy default action. If a user has this option only, the user can modify no other part of the policy or rule.
 - **Modify Remaining Access Control Policy Configuration** controls the ability to edit all other aspects of the policy.

Guidelines and Limitations for Access Control Rules

- There is a limit of 1000 rules that can be shown at a time on the current page. Thus, if you have a very large number of rules, such as 3000 rules within a single category, actions like selecting all rules in a category and deleting them does not delete all of the rules. You might need to select/delete rules again to delete all the rules that you are targeting.
- If you edit an access control rule that is actively in use, the changes do not apply to established connections at deploy-time. The updated rule is used to match against future connections. However, if the system is actively inspecting a connection (for example, with an intrusion policy), it *will* apply changed matching or action criteria to existing connections.

For Firewall Threat Defense, you can ensure that your changes apply to all current connections by using the Firewall Threat Defense **clear conn** CLI command to end established connections. Note that you should only do this if it is acceptable to end those 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.

- VLAN tags in access rules only apply to inline sets; they cannot be used in access rules applied to firewall interfaces.
- To use fully-qualified domain name (FQDN) network objects as source or destination criteria, you must also configure DNS for the data interfaces in the platform settings policy. The system does not use the management DNS server setting to do lookups for FQDN objects used in access control rules.

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

- Whenever possible, use Security Intelligence or URL filtering instead of FQDN rules.
- 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 recompilation of the lookup table, which can impact overall system performance.
 - If more than 8 FQDNs resolve to the same IP address, the system cannot reliably match traffic to the rules for those FQDN. At most 8 FQDNs per IP address can be handled.
- The maximum objects per match criteria per access control rule is 200. For example, you can have up to 200 network objects in a single access control rule.

Managing Access Control Rules

The following topics explain how to manage access control rules.

Adding an Access Control Rule Category

You can divide an access control policy's Mandatory and Default rule sections into custom categories. After you create a category, you cannot move it, although you can delete it, rename it, and move rules into, out of, within, and around it. The system assigns rule numbers across sections and categories.

Procedure

-
- Step 1** In the access control policy editor, click **Add Category**.

Tip

If your policy already contains rules, you can click a blank area in the row for an existing rule to set the position of the new category before you add it. You can also right-click an existing rule and select **Insert new category**.

Step 2 Enter a **Name**.

Step 3 From the **Insert** drop-down list, choose where you want to add the category:

- To insert a category below all existing categories in a section, choose **into Mandatory** or **into Default**.
- To insert a category above an existing category, choose **above category**, then choose a category.
- To insert a category above or below an access control rule, choose **above rule** or **below rule**, then enter an existing rule number.

Step 4 Click **Apply**.

Step 5 Click **Save** to save the policy.

Create and Edit Access Control Rules

Use access control rules to apply actions to specific traffic classes. Rules allow you to selectively allow desirable traffic and drop unwanted traffic.

Before you begin

If you have enabled the Policy Analyzer & Optimizer tool, your edits are evaluated as you make them. If anomalies are detected, you are notified during the edit and prompted to make changes when you click **Apply** to save the rule. You can view the anomalies, such as redundant and shadowed rules, and choose to edit or delete the rule. You can also proceed with saving the rule as is, to deal with the issue later.

Procedure

Step 1 In the access control policy editor, you have the following options:

- To add a new rule, click **Add Rule**.
- To edit an existing rule, click **Edit** (✎).
- To start with a copy of an existing rule, choose one of the following commands from the **More** (⋮) menu.
 - **Copy Rule**, to copy the rule to the clipboard, so you can use the **Paste Above/Below** commands to place it anywhere in the same policy.
 - **Clone Rule**, to create the copy immediately below the duplicated rule.
- To edit multiple rules, use the checkboxes to select multiple rules, then select **Edit** or another action from the **Select Action** list next to the search box.
- To do inline editing, where you change the configuration of an object in a rule condition, right-click the value and choose **Edit**. You can also use the right-click menu to remove a item, add it to the filter, or copy the text or value.

If **View** (👁) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.

Step 2 If this is a new rule, enter a **Name**.

Step 3 Configure the rule components.

If you are bulk-editing multiple rules, only a subset of options are available.

- Position—Specify the rule position; see [Access Control Rule Order, on page 5](#).
- Action—Choose a rule **Action**; see [Access Control Rule Actions, on page 6](#).
- Deep Inspection—(Optional.) For Allow and Interactive Block rules, select options for **Intrusion Policy**, **Variable Set**, and **File Policy**. You can apply intrusion and file policies independently; you do not need to configure both.
- Time Range—(Optional.) For Firewall Threat Defense devices, choose the days and times when the rule is applicable. If you do not choose an option, the rule is always active. For details, see [Creating Time Range Objects](#).
- Logging—Click **Logging** to specify options for connection logging and SNMP traps. The same logging settings can be configured in the **Default Action Configuration** dialog box (navigation path: Edit access control policy. In the **Default Action** area, click the **Default logging and inspection** icon (Cog (⚙️))).

See *Best Practices for Connection Logging* in the [Cisco Secure Firewall Management Center Administration Guide](#) for more information.

Note

You can integrate Splunk or any SIEM syslog server to directly receive connection events from the Firewall Management Center or from its managed devices. If you have configured Firewall Management Center as the source of connection events in Splunk configuration, under **Send Connection Events to**, choose **Firewall Management Center**. If you have configured Firewall Threat Defense as the connection event source, select at least one destination. For the Splunk configuration procedure, see the [Cisco Secure Firewall Management Center Administration Guide](#).

- Application Logging—Click **Application Logging** to specify the protocols for application logging. See *About Application-Aware Event Logging* in the [Cisco Secure Firewall Management Center Administration Guide](#) for more information.

Note

The **Application Logging** option is available when the feature is enabled in the access control policy's advanced settings. See *Enable Application-Aware Event Logging* in the [Cisco Secure Firewall Management Center Administration Guide](#).

- Conditions—Select the objects you want to add or either source or destination, then click either **Add to Sources** or **Add to Destinations and Applications** to add matching conditions for connections. You can click a tab to restrict the list of available objects, for example, to Networks, Security Zones, Applications, and so forth. However, the sources and destination column always show all selected objects regardless of the tab you are on. See [Access Control Rule Conditions, on page 14](#) for more information.
- Comments—Open the comments list at the bottom of the dialog box, enter your comment, and click **Post** to add a comment.

Step 4 Click **Add** or **Apply** to save the rule. Click **Apply and Add New Rule** to leave the dialog box open so you can create a new rule.

Step 5 Click **Save** to save the policy.

What to do next

If you will deploy time-based rules, specify the time zone of the device to which the policy is assigned. See [Time Zone](#).

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

Related Topics

[Best Practices for Access Control Rules](#)

Access Control Rule Conditions

Rule conditions define the characteristics of the connections you want to target with each rule. Use the conditions precisely to fine-tune the rule to apply to all, and only, the traffic that should be handled by the rule. The following topics explain the match conditions that you can use.

Security/Tunnel Zone Rule Conditions

You can use security zones and tunnel zones to select traffic for a rule.

Security zones segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices. Tunnel zones allow you to identify tunneled traffic, such as GRE, that should be handled as a tunnel rather than apply access control rules to the encapsulated connections within the tunnel.

You can use security zones to control traffic by its source and destination interfaces. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones for it to match the rule. Just as all interfaces in a security zone must be of the same type (all inline, passive, switched, or routed), all zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

When using tunnel zones, ensure that you have matching rules in the prefilter policy to associate tunneled traffic with the zone. Then, you can select the tunnel zone as a source zone in a rule; tunnel zones cannot be destinations. If you do not have prefilter rules to rezone the tunnels into the tunnel zone, an access control rule for the tunnel will never apply to any connections. You can specify destination security zones to target tunnels that leave the device through specific interfaces.

Security Zone Considerations

Consider the following when deciding on security zone criteria:

- Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.
- Access control rules generate ACL entries (ACEs) in the device configuration to provide early processing and drops whenever possible. If you specify security zones in rules, ACEs are created for each interface in the zone, which can greatly increase the size of the ACL. Excessively large ACLs generated from access control rules can impact system performance.

Network Rule Conditions

Network rule conditions are 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, add the criteria to the Sources list.
- To match traffic to an IP address or geographical location, add the criteria to the Destinations list.
- 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.

Whenever possible, combine multiple network objects into a single object group. The system automatically creates an object group (during deployment) when you select more than one object (for source or destination separately). Selecting existing groups can avoid object group duplication and reduce the potential impact on CPU usage when there are a large number of duplicate objects.

You can use objects that define the address using the fully-qualified domain name (FQDN); the address is determined through a DNS lookup. However, FQDN objects are not supported in the following sections in access control policies: Original Client networks, SGT/ISE attributes, Network Analysis And Intrusion policy, Security Intelligence, Threat Detection, Elephant Flow Settings.

- **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).

Original Client in Network Conditions (Filtering Proxied Traffic)

For some rules, you can handle proxied traffic based on the originating client. Use a source network condition to specify proxy servers, then add an original client constraint to specify original client IP addresses. The system uses a packet's X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header field to determine original client IP.

Traffic matches the rule if the proxy's IP address matches the rule's source network constraint, **and** the original client's IP address matches the rule's original client constraint. For example, to allow traffic from a specific original client address, but only if it uses a specific proxy, create three access control rules:

Access Control Rule 1: Blocks proxied traffic from a specific IP address (209.165.201.1)

Source Networks: 209.165.201.1

Original Client Networks: none/any

Action: Block

Access Control Rule 2: Allows proxied traffic from the same IP address, but only if the proxy server for that traffic is one you choose (209.165.200.225 or 209.165.200.238)

Source Networks: 209.165.200.225 and 209.165.200.238

Original Client Networks: 209.165.201.1

Action: Allow

Access Control Rule 3: Blocks proxied traffic from the same IP address if it uses any other proxy server.

Source Networks: any

Original Client Networks: 209.165.201.1

Action: Block

VLAN tags rule conditions



Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
 - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
 - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

User rule conditions

Matches traffic based the user who initiates the connection, or the group to which the user belongs. For example, you could configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

You can configure user rule conditions for users in Microsoft Active Directory realms only.

In addition to configuring users and groups for configured realms, you can set policies for the following Special Identities users:

- Failed Authentication: User that failed authentication with the captive portal.
- Guest: Users configured as guest users in the captive portal.

- No Authentication Required: Users that match an identity **No Authentication Required** rule action.
- Unknown: Users that cannot be identified; for example, users that are not downloaded by a configured realm.

For access control rules only, you must first associate an identity policy with the access control policy as discussed in [Associating other policies with access control](#).

Application rule conditions

When the system analyzes IP traffic, it can identify and classify the commonly used applications on your network. This discovery-based *application awareness* is the basis for *application control*—the ability to control application traffic.

System-provided *application filters* help you perform application control by organizing applications according to basic characteristics: type, risk, business relevance, category, and tags. You can create reusable user-defined filters based on combinations of the system-provided filters, or on custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled for an application, the system automatically enables all system-provided detectors for the application; if none exist, the system enables the most recently modified user-defined detector for the application. For more information about application detectors, see [Application Detector Fundamentals](#).

You can use both application filters and individually specified applications to ensure complete coverage. However, understand the following note before you order your access control rules.

Benefits of application filters

Application filters help you quickly configure application control. For example, you can easily use system-provided filters to 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 system blocks the session.

Using application filters simplifies policy creation and administration. It assures you that the system controls application traffic as expected. Because Cisco frequently updates and adds application detectors via system and vulnerability database (VDB) updates, you can ensure that the system uses up-to-date detectors to monitor application traffic. You can also create your own detectors and assign characteristics to the applications they detect, automatically adding them to existing filters.

Application characteristics

The system characterizes each application that it detects using the criteria described in the following table. Use these characteristics as application filters.

Table 1: Application Characteristics

Characteristic	Description	Example
Type	Application protocols represent communications between hosts. Clients represent software running on a host. Web applications represent the content or requested URL for HTTP traffic.	HTTP and SSH are application protocols. Web browsers and email clients are clients. MPEG video and Facebook are web applications.
Risk	The likelihood that the application is being used for purposes that might be against your organization's security policy.	Peer-to-peer applications tend to have a very high risk.

Characteristic	Description	Example
Business Relevance	The likelihood that the application is being used within the context of your organization's business operations, as opposed to recreationally.	Gaming applications tend to have a very low business relevance.
Category	A general classification for the application that describes its most essential function. Each application belongs to at least one category.	Facebook is in the social networking category.
Tag	Additional information about the application. Applications can have any number of tags, including none.	Video streaming web applications often are tagged <code>high bandwidth</code> and <code>displays ads</code> .

Related Topics

[Best Practices for Configuring Application Control](#)

Configuring Application Conditions and Filters

To build an application condition or filter, choose the applications whose traffic you want to control from a list of available applications. Optionally (and recommended), constrain the available applications using filters. You can use filters and individually specified applications in the same condition.

Before you begin

- Adaptive profiling must be enabled (its default state) for access control rules to perform application control.

Procedure


Step 1 Invoke the rule or configuration editor:

- Access control, decryption, QoS rule condition—In the rule editor, click **Applications**.
- Identity rule condition—In the rule editor, click **Realms & Settings** and enable active authentication; see [Create an Identity Rule](#).
- Application filter—On the Application Filters page of the object manager, add or edit an application filter. Provide a unique **Name** for the filter.
- Intelligent Application Bypass (IAB)—In the access control policy editor, click **Advanced**, edit IAB settings, then click **Bypassable Applications and Filters**.

Step 2 Find and choose the applications you want to add from the **Available Applications** list.

To constrain the applications displayed in **Available Applications**, choose one or more **Application Filters** or search for individual applications.

Tip

Click **Information**  next to an application to display summary information and internet search links. **Unlock** marks applications that the system can identify only in decrypted traffic.

When you choose filters, singly or in combination, the Available Applications list updates to display only the applications that meet your criteria. You can choose system-provided filters in combination, but not user-defined filters.

- Multiple filters for the same characteristic (risk, business relevance, and so on)—Application traffic must match only one of the filters. For example, if you choose both the medium and high-risk filters, the Available Applications list displays all medium and high-risk applications.
- Filters for different application characteristics—Application traffic must match both filter types. For example, if you choose both the high-risk and low business relevance filters, the Available Applications list displays only applications that meet both criteria.

Step 3 Click **Add Application** or **Add to Rule**, or drag and drop.

Tip

Before you add more filters and applications, click **Clear Filters** to clear your current choices.

Step 4 (Access control rules only.) If destination ports are not specified on the Ports tab, select the **Port** for the applications.

The **Port** specification applies to all the applications in the list; it does not apply to any filters. You cannot specify different Port options for each application. If you already specified destination ports on the Ports tab, that selection is reflected in the application list, and you cannot make the following selection.

If you have not already specified destination ports for the rule, choose one of the following:

- **Application Default**—The system looks at the default port only when evaluating traffic for an application match. Click the info button (i) to see what is considered the default for each application.

When editing a rule, if the rule already has source ports, and you add an application that has an incompatible port and select Application Default port, you are warned. If you elect to use the application default port, the incompatible source port is removed.

- **Any**—The system does not restrict the application identification based on the port used in the connection. If there are source ports specified on the Port tab, you might need to select this option to be able to save the rule.

Regardless of the option you select, if you subsequently configure any destination value on the Ports tab, the ports specified on the Ports tab override your selection here, and change the application port to Any. If you want to limit the rule to one, multiple, or a range of ports that are not the application default port, use the Ports tab.

This option does not apply to custom detectors.

Note

If the access control policy is assigned to devices running 7.x versions, the rule is handled differently if you have included applications that have different default ports. When you select to apply the rule to the default ports only, then for 7.7 systems, the rule is split into separate rules, one per application. For systems running versions 7.3-7.6, the rule is kept as a single rule, but the ports page is updated with each of the default ports. If you are managing devices running down-level FTD versions, and you want to use default port matching, we recommend that you limit the application's list in each rule to applications that have the same default port.

Step 5 Save or continue editing the rule or configuration.

What to do next

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

Port, Protocol, and ICMP Code Rule Conditions

Port conditions match traffic based on the source and destination ports. Depending on the rule type, “port” can represent any of the following:

- TCP and UDP—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- ICMP—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- Protocol—You can control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic. Note that application filtering is not available in prefilter rules.

Application filtering is also recommended for applications, like FTP, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although you can configure port conditions to match non-TCP traffic, there are some restrictions:

- Access control rules—For Classic devices, you can match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, you can add only network-based conditions: zone, IP address, port, and VLAN tag. Also, the system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- Decryption rules—These rules support TCP port conditions only.

- IMCP echo—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

URL Rule Conditions

Use URL conditions to control the websites that users on your network can access.

For complete information, see [About URL Filtering with Category and Reputation](#).

Dynamic Attributes Rule Conditions

Dynamic attributes include the following:

- (Source or destination.) Dynamic objects (such as from the Dynamic Attributes Connector)

The dynamic attributes connector enables you to collect data (such as networks and IP addresses) from cloud providers and send it to the Secure Firewall Management Center so it can be used in access control rules.

For more information about the dynamic attributes connector, see [About the Dynamic Attributes Connector](#).

- (Source only.) SGT objects contain tags either manually defined or defined in ISE. For more information, see [Security Group Tag](#).
- (Source only.) Location IP objects, defined by Cisco ISE
- (Source only.) Device type objects, defined by Cisco ISE (also referred to as endpoint profile objects)

Dynamic attributes can be used as source criteria and destination criteria in access control rules. Use the following guidelines:

- Objects of different types are ANDd together
- Objects of a similar type are ORd together

For example, if you choose source destination criteria SGT 1, SGT 2, and device type 1; the rule is matched if device type 1 is detected on either SGT 1 or SGT 2. As another example, if you select both a security group tag, and a dynamic object that lists IP addresses, the rule matches if traffic with the tag originates from (or is destined to) one of those IP addresses.

Time and Day Rule Conditions

You can specify a continuous time range or a recurring time period.

For example, a rule can apply only during weekday working hours, or every weekend, or during a holiday shutdown period.

Time-based rules are applied based on the local time of the device that processes the traffic.

Time-based rules are supported only on Firewall Threat Defense devices. If you assign a policy with a time-based rule to a different type of device, the time restriction associated with the rule is ignored on that device. You will see warnings in this case.

Enabling and Disabling Access Control Rules

When you create an access control rule, it is enabled by default. If you disable a rule, the system does not use it to evaluate network traffic and stops generating warnings and errors for that rule. When viewing the list of rules in an access control policy, disabled rules are grayed out, although you can still modify them.

You can also enable or disable an access control rule using the rule editor.

Procedure

-
- Step 1** In the access control policy editor, right-click the rule and choose a rule state.
- If **View** (👁) appears next to a rule instead, the rule belongs to an ancestor policy, or you do not have permission to modify the rule.
- Step 2** Click **Save**.
-

What to do next

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

Copying Access Control Rules from One Access Control Policy to Another

You can copy access control rules from one access control policy to another. You can copy the rules either to the **Default** section or the **Mandatory** section of the access control policy.

All the settings of the copied rules, except the comments, are retained in the pasted version.

Procedure

-
- Step 1** Do one of the following:
- To copy a single rule, right-click the rule and select **Copy to Different Policy**.
 - To copy multiple rules, select their checkboxes, then select **Copy to Different Policy** from the **Select Bulk Action** menu.
- Step 2** Select the destination access control policy from the **Access Policy** drop-down list.
- Step 3** From the **Place Rules** drop-down list, choose where you want to position the copied rules. You can place them at the top or bottom of either the Mandatory or Default sections.
- Step 4** Click **Copy**.
-

What to do next

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

Moving Access Control Rules to a Prefilter Policy

You can move access control rules from an access control policy to the associated non-default prefilter policy.

You must first apply a user-defined prefilter policy to the access control policy. The access control rules cannot be moved to the default prefilter policy because the default prefilter policy cannot have rules.

Before you begin

Note the following conditions before you proceed:

- When moving an access control rule to a prefilter policy the layer 7 (L7) parameters in the access control rule cannot be moved. The L7 parameters are dropped during the operation.
- The comments in the access control rule configuration are lost after moving the rule. However, a new comment is added in the moved rule mentioning the source access control policy.
- You cannot move access control rules with **Monitor** set as the **Action** parameter.
- The **Action** parameter in the access control rule is changed to a suitable action in the prefilter rule when moved. To know what each action in the access control rule maps to, see the following table:

Action in the access control rule	Action in the prefilter rule
Allow	Analyze
Block	Block
Block with reset	Block
Interactive Block	Block
Interactive Block with reset	Block
Trust	Fastpath

- Similarly, based on the action configured in the access control rule, the logging configuration is set to an appropriate setting after the rule is moved, as mentioned in the following table.

Action in the access control rule	Enabled Logging configurations in the prefilter rule
Allow	None of the check boxes are checked.
Block	<ul style="list-style-type: none">• Log at Beginning of Connection• Event Viewer• Syslog Server• SNMP Trap

Action in the access control rule	Enabled Logging configurations in the prefilter rule
Block with reset	<ul style="list-style-type: none"> • Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Interactive Block	<ul style="list-style-type: none"> • Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Interactive Block with reset	<ul style="list-style-type: none"> • Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Trust	<ul style="list-style-type: none"> • Log at Beginning of Connection • Log at End of Connection • Event Viewer • Syslog Server • SNMP Trap

- While moving rules from the source policy, if another user modifies those rules, you will see get a message. You may continue with the process after refreshing the page.

Procedure

Step 1

Do one of the following:

- To move a single rule, right-click the rule and select **Move to Prefilter Policy**.
- To move multiple rules, select their checkboxes, then select **Move to Prefilter Policy** from the **Select Bulk Action** menu.

Step 2

From the **Place Rules** drop-down list, choose where you want to position the moved rules:

- To position as the last set of rules, choose **At the bottom**.
- To position as the first set of rules, choose **At the top**.

Step 3 Click **Move**.

What to do next

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

Positioning an Access Control Rule

You can move an existing rule within an access control policy, or insert new rules in a desired location. When you add or move a rule to a category, the system places it last in the category.

Before you begin

Review rule order guidelines in [Best Practices for Access Control Rules](#).

Procedure

- Step 1** Do one of the following:
- New rule—Insert a new rule by hovering over the line between the existing rules, and clicking **Add Rule**. The location is selected in the **Insert** box in the Add Rule dialog box; you can select a different rule to adjust the location. You can also select **Add Rule Above** or **Add Rule Below** from the right-click menu.
 - Existing rules when viewing the rule table—Click and drag the rule to the new position.
 - Existing rules when viewing the rule table—Right-click a single rule and select **Reposition Rule**. To move multiple rules as a group, select their checkboxes, then select **Reposition Rules** from the **Select Bulk Action** menu.
 - Existing rule when editing the rule—Click the **Reposition Rule** icon next to the rule name.
- Step 2** Choose where you want to move or insert the rule:
- Choose **into Mandatory** or **into Default**.
 - Choose a **into Category**, then choose the category.
 - Choose **above rule** or **below rule**, then select the rule.
- Step 3** Click **Move** or **Confirm**, and save the rule if you are editing it.
- Step 4** Click **Save** to save the policy.
-

What to do next

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

Adding Comments to an Access Control Rule

When you create or edit an access control rule, you can add a comment. For example, you might summarize the overall configuration for the benefit of other users, or note when you change a rule and the reason for the change. You can display a list of all comments for a rule along with the user who added each comment and the date the comment was added.

When you save a rule, all comments made since the last save become read-only.

To search access control rule comments, use the "Search Rules" bar on the rule listing page.

Procedure

-
- Step 1** In the access control rule editor, click **Comments**.
 - Step 2** Enter your comment and click **Add Comment**. You can edit or delete this comment until you save the rule.
 - Step 3** Save the rule.
-

Examples for Access Control Rules

The following topics provide examples of access control rules.

How to Control Access Using Security Zones

Consider a deployment where you want hosts to have unrestricted access to the internet, but you nevertheless want to protect them by inspecting incoming traffic for intrusions and malware.

First, create two security zones: Internal and External. Then, assign interface pairs on one or more devices to those zones, with one interface in each pair in the Internal zone and one in the External zone. Hosts connected to the network on the Internal side represent your protected assets.



Note You are not required to group all internal (or external) interfaces into a single zone. Choose the grouping that makes sense for your deployment and security policies.

Then, configure an access control rule with a destination zone condition set to Internal. This simple rule matches traffic that leaves the device from any interface in the Internal zone. To inspect matching traffic for intrusions and malware, choose a rule action of **Allow**, then associate the rule with an intrusion and a file policy.

How to Control Application Usage

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

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

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

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

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

Procedure

Step 1 Choose **Policies > Security policies > Access Control** and edit the access control policy.

Step 2 Click **Add Rule** and configure the rule for application control.

- a) Give the rule a meaningful name, such as Block_Anonymizers.
- b) Select **Block** for **Action**.

Name: Action:  Block

- c) Assuming that you have configured zones, and want this rule to apply to traffic going from the inside to outside, select the **Zones** tab, and choose your inside zone as the source zone, and the outside zone as the destination zone.
- d) Click the **Applications** tab, select the applications to match, and click **Add Application**.

As you select criteria, such as category and risk level, the list to the right of the criteria updates to show exactly which applications match the criteria. The rule you are writing applies to these applications.

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

For purposes of this example, select anonymizers/proxies from the Categories list and add it to Destinations and Applications. The match criteria should now look like the following graphic.

Selected Sources: 1		Selected Destinations and Applications: 2	
Collapse All	Remove All	Collapse All	Remove All
<div>ZONE</div> <div>▼ 1 Object</div> <div>inside-zone</div>		<div>ZONE</div> <div>▼ 1 Object</div> <div>outside-zone</div>	
		<div>APP</div> <div>▼ 1 Object</div> <div>Categories: anonymizer...</div>	

- e) Click **Logging** next to the rule action, and enable logging at the start of the connection. You can select a syslog server if you use one.

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

Note

To optimize performance, log either the beginning or the end of any connection, but not both. See [Cisco Secure Firewall Management Center Administration Guide](#) for more information.

- Step 3** Move the rule so that it comes after any rules that use protocol and port criteria only, but that would not allow traffic that should be blocked by the application rule.

Matching applications requires Snort inspection. Because Snort inspection is not needed by rules that use protocol and port only, you can improve system performance by grouping these simple rules at the top of the access control policy as much as possible.

- Step 4** Deploy the changes.

You can use the application rule hit counts and analysis dashboards to see how this rule is performing and how often users try these applications.

How to Block Threats

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

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

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

Besides inspecting traffic that you allow for potential intrusions, you can use the Security Intelligence policy to preemptively block all traffic to or from known bad IP addresses, or to known bad URLs.

This example adds an intrusion policy that allows the internal 192.168.1.0/24 network to go outside, and assumes you already have block rules to selectively eliminate unwanted connections, while also adding a Security Intelligence policy to do pre-emptive blocking.

Before you begin

You must apply the IPS license to any managed device that uses this rule.

This example assumes you have already created security zones for inside and outside interfaces, and the network object for the inside network.

Procedure

Step 1 Create the access control rule that applies the intrusion policy.

- While editing the access control policy, click **Add Rule**.
- Give the rule a meaningful name, such as `Inside_Outside`, and ensure the rule action is **Allow**.

The screenshot shows a configuration form with two fields. The 'Name' field is a text input containing the text 'Inside_Outside'. The 'Action' field is a dropdown menu with a blue arrow icon and the text 'Allow'.

- For **Intrusion Policy**, select **Balanced Security and Connectivity**. You can either accept the default variable set or select your own if you want to customize it.

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

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

If you create your own custom policy, you can select that one instead.

A discussion of variable sets is beyond the scope of this example. Please read the chapters on intrusion policy for detailed information about variable sets and custom policies.

The screenshot shows a dropdown menu for 'Intrusion Policy'. The selected option is 'Balanced Security and C...' with a small 'x' icon to its right. Below it, 'Default-Set' is visible as another option.

- Select the **Zones** tab, and add your inside security zone to the source criteria, and outside zone to the destination criteria.
- Select the **Networks** tab, and add the network object that defines your inside network to the source criteria.

The match criteria should look similar to the following:

The screenshot shows a configuration interface for match criteria. It is divided into two main sections: 'Selected Sources: 2' and 'Selected Destinations and Applications: 1'. Each section has a 'Collapse All' and 'Remove All' link. Under 'Selected Sources', there are two categories: 'ZONE' with one object 'inside-zone' and 'NET' with one object 'Inside-Network'. Under 'Selected Destinations and Applications', there is one category: 'ZONE' with one object 'outside-zone'.

- Click **Logging** and enable logging at the beginning or end of the connection, or both, as desired.
- Click **Apply** to save the rule, and then **Save** to save the updated policy.
- Move the rule to the appropriate location in the access control policy.

Step 2 Configure the Security Intelligence policy to preemptively drop connections with known bad hosts and sites.

By using Security Intelligence to block connections with hosts or sites that are known to be threats, you save your system the time needed to do deep packet inspection to identify threats in each connection. Security Intelligence provides an early block of undesirable traffic, leaving more system time to handle the traffic you really care about.

- a) While editing the access control policy, click the **Security Intelligence** link in the packet path.

The link includes two policies: the DNS policy at the top, and the Security Intelligence (Network and URL) at the bottom. In this example, we are configuring the Network and URL lists. By default, these lists already include the global block and do not block lists, but these lists are empty by default until you add items to them.

- b) With **Networks** selected and the **Any** security zone selected, scroll down in the list until you get to the global lists, and the first Security Intelligence category (probably Attackers). Click Attackers, then scroll to the end of the categories (probably Tor_exit_node), and Shift+Click to select all of the categories. Click **Add To Block List**.
- c) Select the **URL** tab, and **Any** security zone, and use Shift+Click to select URL versions of the same categories. Click **Add To Block List**.
- d) Click **Save** to save the policy.
- e) As necessary, you can add network and URL objects to the block or do not block lists.

The **Do Not Block** lists are not real "allow" lists. They are exception lists. If an address or URL in the exception list also appears in the blocked list, the connection for the address or URL is allowed to pass on to the access control policy. This way, you can block a feed, but if you later find that a desirable address or site is being blocked, you can use the exception list to override that block without needing to remove the feed entirely. Keep in mind that these connections are subsequently evaluated by access control, and if configured, an intrusion policy. Thus, if any connections do contain threats, they can be identified and blocked during intrusion inspection.

Use the events and dashboards to determine what traffic is actually being dropped by the policy, and whether you need to add addresses or URLs to the **Do Not Block** lists.

Step 3 Deploy your changes.

History for Access Control Rules

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Default Port Specification for Application-based Access Control Rules.	10.0	10.0	<p>When adding applications to access control rules, you can specify that the rule be limited to the default ports for the application. This is different from manually specifying ports on the Ports tab. You can alternatively specify that the application be identified on Any port, which is the behavior of the system prior to this enhancement.</p> <p>If you specify the default port and also make a specification on the Ports tab, the Ports tab overrides the application default port. You can use the Ports tab to limit the rule to one, multiple, or a range of ports.</p> <p>On upgrade, existing application-based rules continue to be applied to Any port. If you want to limit these rules to the default port, you must edit the rules, clear any specifications on the Ports tab, select Application Default port on the Applications tab, then save the rule. Deploy changes after editing all the rules you want to change.</p> <p>We added the ability to select Application Default or Any port when adding an application to an access control rule, so long as nothing is configured on the Ports tab. This option does not apply to application filters of any kind, nor does it apply to any other policy or object except access control rules.</p>
Maximum objects per match criteria per access control rule is 200.	7.3	Any	<p>Previously, you could have a maximum of 50 objects per match criteria in a single access control rule. For example, you could have up to 50 network objects in a single access control rule. The limit is now 200 objects per match criteria in a single rule.</p> <p>We updated the access control policy to allow the increased object limit.</p>
Search for access control rule comments	6.7	Any	<p>The Search Rules bar now offers the option to search for comments.</p> <p>New/modified pages: Access control rules page, Search Rules text entry field.</p> <p>Supported platforms: Firewall Management Center</p>
Copy or move rules between access control and prefilter policies	6.7	Any	<p>You can copy access control rules from one access control policy to another. You can also move access control rules from an access control policy to the associated prefilter policy.</p> <p>New/modified pages: Access control policy page; the right-click menu for the selected rules provides additional options to copy and move.</p> <p>Supported platforms: Firewall Management Center</p>

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Bulk edit of certain settings in access control rules	6.6	Any	<p>In the list of rules in a policy, shift-click or control-click to select multiple rules, then right-click and choose an option. Example bulk operations: You can enable or disable the rules, select a rule action, and edit most inspection and logging settings.</p> <p>New/modified pages: Access control rules page.</p> <p>Supported platforms: Firewall Management Center</p>
Enhanced searching on configured rules	6.6	Any	<p>Enhanced searching on configured rules.</p> <p>New/modified pages: Access control rules page.</p> <p>Supported platforms: Firewall Management Center</p>
Time ranges for rule application	6.6	Any	<p>Ability to specify an absolute or recurring time or time range for a rule to be applied. The rule is applied based on the time zone of the device that processes the traffic.</p> <p>New/modified pages:</p> <ul style="list-style-type: none"> • A new option on the access control Add Rule page. • A related new option on the Devices > Platform Settings > Threat Defense page to specify time zones for managed devices. <p>Supported platforms: Firewall Threat Defense devices only</p>
View object details from access control rule pages	pre-6.6	Any	<p>To see information about an object in a list of rules or from the rule configuration dialog, right-click the object.</p> <p>New/modified pages: Policies > Access Control > Access control, and Add Rule page.</p> <p>Supported platforms: Firewall Management Center</p>