

Access Control Rules

The following topics describe how to configure access control rules.

- Access control rules, on page 1
- Requirements and prerequisites for access control rules, on page 12
- Guidelines and limitations for access control rules, on page 13
- Best practices for application control, on page 14
- Best practices for access control rules, on page 17
- Access control rule management, on page 22
- Examples for access control rules, on page 36
- History for access control rules, on page 40

Access control rules

Within an access control policy, access control rules provide a granular method that

- handles network traffic across multiple managed devices within an access control policy
- provides conditions that match traffic in a specified order, and
- determines actions to monitor, trust, block, or allow matching traffic.

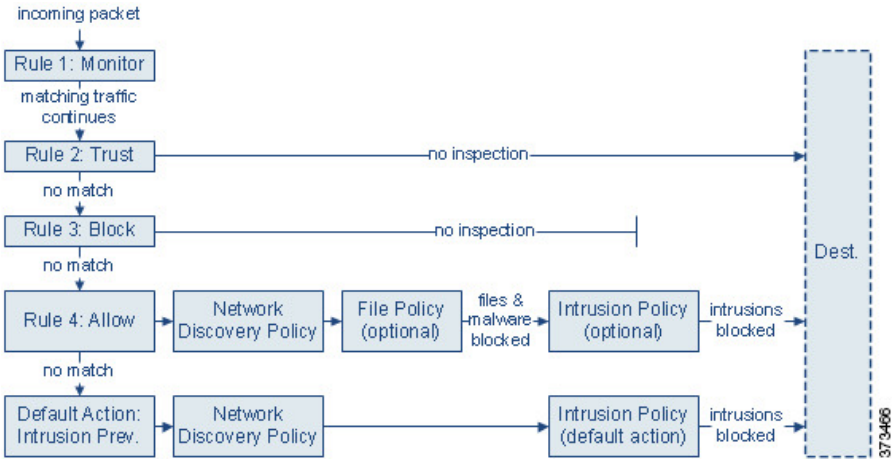
Rule evaluation and traffic handling

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.

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

This 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.

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.

The system provides specific icons to communicate the inspection options or status of each rule:

- **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, refer to Create and edit access control rules, on page 23 .

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 3
Best practices for access control rules, on page 17

Access control rule components

In addition to its unique name, each access control rule has basic components that determine how the system processes and handles network traffic.

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 set rules to operate during specified days and times

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

Logging settings govern records that the system keeps for handled traffic. You can log sessions at connection start, end, or both; log connections to databases, syslog, or SNMP trap servers. You can log connections to the database, as well as to a syslog or an SNMP trap server.

Comments

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

Related Topics

Best practices for access control rules, on page 17

Access control rule management, on page 3
Create and edit access control rules, on page 23
Access control rule actions, on page 6
Access control rule conditions, on page 25
Deep inspection using file and intrusion policies, on page 9
Access control rule comments

Access control rule order

Access control rule order is a traffic processing method that

- numbers rules in an access control policy starting at 1
- matches traffic to access control rules in top-down order by ascending rule number, and
- handles network traffic according to the first access control rule where all the rule's conditions match the traffic.

In most cases, 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.

Note 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. For more information, refer to Best practices for ordering rules, on page 18.

Access control rule actions

An access control rule action is a configuration setting that determines how the system handles and logs matching traffic, with options to monitor, trust, block, or allow traffic with or without further inspection.

Default action behavior

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 an access control rule action that forces connection logging regardless of how matching traffic is eventually handled. It does not affect traffic flow by permitting or denying traffic immediately, and allows traffic to be matched against additional rules to determine the final action.

When a connection matches a monitor rule, the next non-monitor rule that the connection matches determines traffic handling and any further inspection. If there are no additional matching rules, the Firewall Threat Defense device uses the default action.

There is an exception: 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), even if a subsequent rule would block the connection. Early packets are not evaluated against subsequent rules, so to prevent uninspected traffic, you can specify an intrusion policy for these packets in the access control policy's advanced settings. For more information, refer to 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.

Access control rule trust action

The **Trust** action is the access control rule action that allows traffic to pass without deep inspection or network discovery, still subjects trusted traffic to identity requirements and rate limiting, and bypasses additional inspection services while maintaining basic security controls.

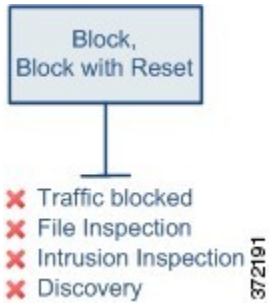


Note

- 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

A block action is an access control rule action that denies network traffic without further inspection and prevents the use of file, intrusion, or network discovery policies for blocked traffic.



Block and **Block with reset** actions immediately stop traffic from entering the network. Block with reset rules reset the connection except for web requests met with an HTTP response page. The HTTP response page is used to notify users when their web request is blocked, but it cannot be displayed if the connection is immediately reset.

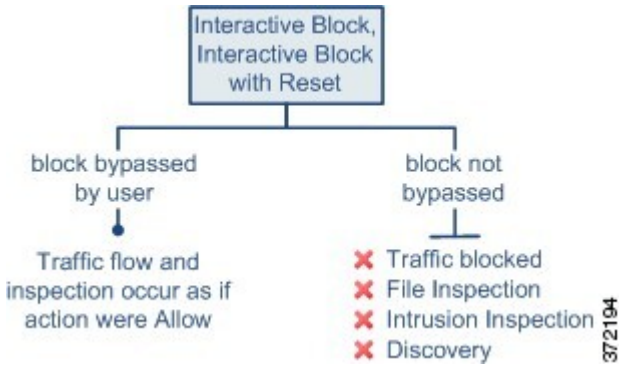
For more information, refer to HTTP response pages.

Related Topics

HTTP response pages

Access control rule interactive blocking actions

The **Interactive Block** and **Interactive Block with reset** actions are access control rule actions that 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 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, refer to HTTP response pages.

Related Topics

Decryption rule blocking actions

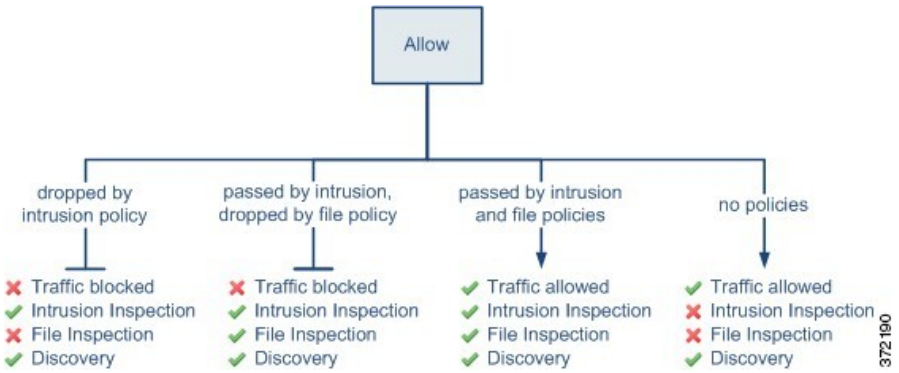
Access control rule allow action

An access control rule allow action is a rule action that allows matching traffic to pass. However, traffic 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.

This 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, permitting traffic does not guarantee that discovery inspection occurs. The system only performs discovery for connections involving IP addresses explicitly monitored by your network discovery policy. Also, application discovery is limited for encrypted sessions.

Deep inspection using file and intrusion policies

Deep inspection is a security mechanism that uses intrusion and file policies as the last line of defense before traffic is allowed to its destination.

Policy types and capabilities

Deep inspection uses two types of policies to analyze traffic:

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

Access control occurs before deep inspection; access control rules and the access control default action determine which traffic is inspected by intrusion and file policies.

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.

In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.

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.

To associate intrusion and file policies with an access control rule, see:

- Access Control Rules: Intrusion Policy Selection
- Configure an access control rule to perform malware protection

For more information about intrusion policies, refer to Network Analysis and Intrusion Policy Basics.

For more information about file policies, refer to File Policies for Network Malware Protection.

File and intrusion inspection order

File and intrusion inspection order defines the traffic processing sequence that performs file inspection before intrusion inspection for any single connection handled by an access control rule, preventing files blocked by a file policy from being inspected for intrusions while prioritizing simple file blocking by type over malware inspection.

For a connection matching an Allow or Interactive Block rule, the system determines traffic flow based on the following configurations:

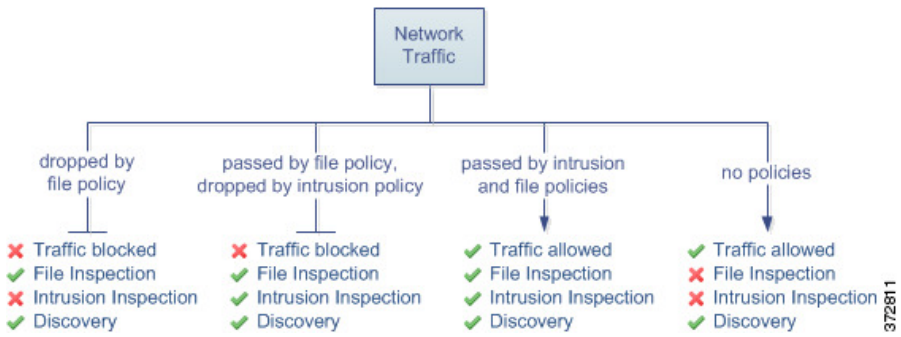
- Without a file policy, traffic flow is determined by the intrusion policy.
- Without an intrusion policy, traffic flow is determined by the file policy.
- Without either, allowed traffic is inspected by network discovery only.

In your access control policy, you can associate multiple Allow and Interactive Block rules with different intrusion and file policies to match inspection profiles to various types of traffic.

Note The system does not perform any kind of inspection on trusted traffic. Although configuring an Allow rule with neither an intrusion nor file policy passes traffic like a Trust rule, Allow rules let you perform discovery on matching traffic.

The diagram below illustrates the types of inspection you can perform on traffic that meets the conditions of either an Allow or user-bypassed Interactive Block access control rule. For simplicity, the diagram displays traffic flow for situations where both (or neither) an intrusion and a file policy are associated with a single access control rule.

Figure 1: Traffic inspection flowchart



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.

For example, consider a scenario where you normally want to allow certain network traffic as defined in an access control rule. However, as a precaution, you want to block the download of executable files, examine downloaded PDFs for malware and block any instances you find, and perform intrusion inspection on the traffic.

You create an access control policy with a rule that matches the characteristics of the traffic you want to provisionally allow, and associate it with both an intrusion policy and a file policy. The file policy blocks the download of all executables, and also inspects and blocks PDFs containing malware:

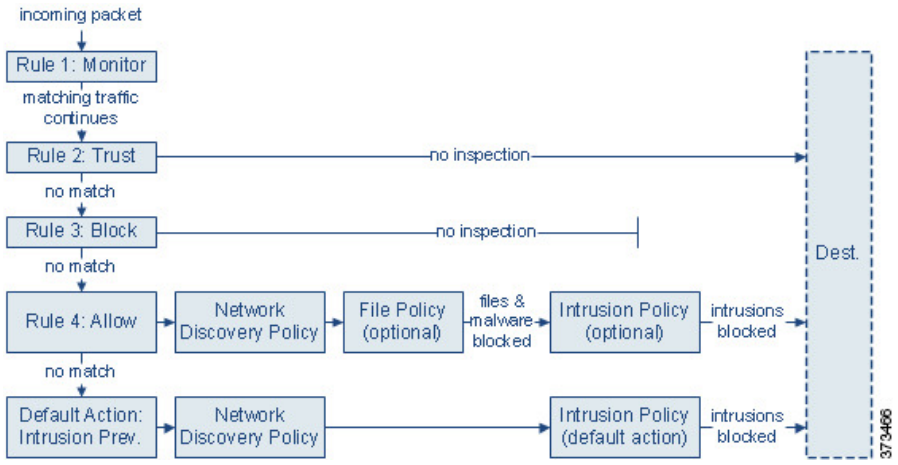
- First, the system blocks the download of all executables, based on simple type matching specified in the file policy. Because they are immediately blocked, these files are subject to neither malware nor intrusion inspection.
- Next, the system performs malware cloud lookups for PDFs downloaded to a host on your network. Any PDFs with a malware disposition are blocked, and are not subject to intrusion inspection.
- Finally, the system uses the intrusion policy associated with the access control rule to inspect any remaining traffic, including files not blocked by the file policy.

Note Until a file is detected and blocked in a session, packets from the session might be subject to intrusion inspection.

How access control traffic handling works with intrusion and file policies

Summary

This diagram shows the flow of traffic in an inline intrusion prevention and malware defense deployment, as governed by an access control policy that contains four different types of access control rules and a default action.



Workflow

In the scenario above, the first three access control rules in the policy—Monitor, Trust, and Block—cannot inspect matching traffic. Monitor rules track and log but do not inspect network traffic, so 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.) Trust and Block rules handle matching traffic without further inspection of any kind, while traffic that does not match continues to the next access control rule.

The fourth and final rule in the policy, an Allow rule, invokes various other policies to inspect and handle matching traffic, in the following order:

1. Discovery: Network Discovery Policy—First, the network discovery policy inspects traffic for discovery data. Discovery is passive analysis and does not affect the flow of traffic. Although you do not explicitly enable discovery, 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.
2. malware defense and File Control: File Policy—After traffic is inspected by discovery, the system can inspect it for prohibited files and malware. malware defense detects and optionally blocks malware in many types of files, including PDFs, Microsoft Office documents, and others. If your organization wants to block not only the transmission of malware files, but all files of a specific type (regardless of whether the files contain malware), file control allows you to monitor network traffic for transmissions of specific file types, then either block or allow the file.
3. Intrusion Prevention: Intrusion Policy—After file inspection, the system can inspect traffic for intrusions and exploits. An intrusion policy examines decoded packets for attacks based on patterns, and can block or alter malicious traffic. Intrusion policies are paired with variable sets, which allow you to use named values to accurately reflect your network environment.
4. Destination—Traffic that passes all the checks described above passes to its destination.

An Interactive Block rule (not shown in the diagram) has the same inspection options as an Allow rule. This is so you can inspect traffic for malicious content when a user bypasses a blocked website by clicking through a warning page.

Traffic that does not match any access control rules in the policy with an action other than Monitor is handled by the default action. In this scenario, the default action is an Intrusion Prevention action, which allows traffic to its final destination as long as it is passed by the intrusion policy you specify. In a different deployment, you might have a default action that trusts or blocks all traffic without further inspection. Note that the system can inspect traffic allowed by the default action for discovery data and intrusions, but not prohibited files or malware. You cannot associate a file policy with the access control default action.

Note Sometimes, when a connection is analyzed by an access control policy, the system must process the first few packets in that connection, allowing them to pass, before it can decide which access control rule (if any) will handle the traffic. However, so these packets do not reach their destination uninspected, you can specify an intrusion policy (in the advanced settings for the access control policy) to inspect these packets and generate intrusion events.

Requirements and prerequisites for access control rules

Ensure that your deployment meets these requirements before you configure access control rules.

Model support

Any

Supported domains

Any

User roles

- Admin
- Access Admin
- Network Admin

Guidelines and limitations for access control rules

Consider these limitations when you configure rules in the access control policy:

- 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 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.
- Block with reset can be CPU-intensive for rules that drop a high rate of TCP traffic. When the system drops TCP traffic with this action, it also generates TCP reset packets, which can affect device performance during high-volume drops. For access control rules that are expected to drop large amounts of TCP traffic, use **Block** instead of **Block with reset**, unless TCP reset behavior is required.

  If TCP reset behavior is required for denied inbound TCP traffic, you can also evaluate using a regular **Block** rule with the **service resetinbound** command through FlexConfig for the applicable interface.

Best practices for application control

These topics discuss our recommended best practices for controlling applications with access control rules.

Recommendations for application control

Ensure that adaptive profiling is enabled

If adaptive profiling is not enabled (its default state), access control rules cannot perform application control.

Enable application detectors automatically

If no detector is enabled for an application you want to detect, 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.

Configure your policy to examine the packets that must pass before an application is identified

The system cannot perform application control before *both* of these conditions occur:

- A monitored connection is established between a client and server
- The system identifies the application in the session

This identification should occur in 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted.

If early traffic matches all other criteria but application identification is incomplete, the system allows the packet to pass and the connection to be established (or the SSL handshake to complete). After the system completes its identification, the system applies the appropriate action to the remaining session traffic.

To ensure that your system examines these initial packets, select an intrusion policy in the **Intrusion Policy used before Access Control rule is determined** options in the access control policy advanced settings.

Understand limitations to identifying applications

A server must adhere to the protocol requirements of an application for the system to be able to recognize it. For example, if you have a server that sends a keep-alive packet rather than an ACK when an ACK is expected, that application might not be identified, and the connection will not match the application-based rule. Instead, it will be handled by another matching rule or the default action. This might mean that connections you want to allow can be denied instead. If you run into this problem, and you cannot fix the server to follow the protocol standards, you need to write a non-application-based rule to cover traffic for that server, for example, by matching the IP address and port number.

Create separate rules for URL and application filtering

Create separate rules for URL and application filtering whenever possible, because combining application and URL criteria can lead to unexpected results, especially for encrypted traffic.

Rules that include both application and URL criteria should come after 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.

Place URL rules before application and other rules

For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of these criteria:

- They include application conditions.
- The traffic to be inspected is encrypted.

Handling application traffic packets without payloads

When performing access control, the system applies the default policy action to packets that do not have a payload in a connection where an application is identified.

Handling referred application traffic

To handle traffic referred by a web server, such as advertisement traffic, match the referred application rather than the referring application.

Controlling application traffic that uses multiple protocols

Some applications use multiple protocols. To control their traffic, make sure your access control policy covers all relevant options. These applications typically include the application name in each aspect that you can control. Also see Application-specific requirements for filtering, on page 16.

Controlling evasive application traffic

See Application-specific requirements for filtering, on page 16.

Best practice for choosing between application matching and port matching

In traditional firewalls, you can match traffic based on OSI layers 3 (protocol) and 4 (transport), such as IP or TCP/80. All traffic on a given port (or a whole protocol) is then either allowed or blocked based on that rule's action.

Application criteria, on the other hand, are OSI layer 7. Different applications can use the same TCP/UDP port. By using application criteria, you can selectively allow or block different applications on the same port without allowing or blocking all applications on that port.

Whether you use port-based or application-based criteria in your rule can impact rule performance. Because TCP/UDP ports are quickly identifiable in packets, the system can match the correct rule on the first packet. With application-layer criteria, it can take 3-5 packets to identify the specific application (if you do not also specify ports).

Consider these best practices:

- If you want to handle all of the traffic on a given TCP/UDP port in the same way for the specified interfaces and networks, use port-based matching. For example, to handle all SSH traffic the same way, select the SSH port (TCP/22) on the Ports tab.
- If you want to narrowly match a specific application that uses the same port as other applications, select the application on the Applications tab. This is how you can handle web applications, which all use TCP/80 or 443, so that you can selectively block or allow certain web applications without blocking or allowing all web applications.
- If you want to control application use by a user group, select the user group on the Users tab and the application on the Application tab. For example, you could block the Gaming applications category for members of the Contractor user group.
- For rules that go from a less secure network (such as the internet) to a more secure network (such as an internal protected network), use the Ports tab whenever possible. For example, you can allow/block ICMP traffic from the internet to an internal network.
- If you do have a mix of port-based and application-based rules, place the port-based rules higher in the rules list, so that connections can be matched to those rules first. Protocol and port can be more quickly identified than applications.

Application-specific requirements for filtering

Office 365 admin portal

If the access policy has logging enabled at the beginning as well as at the end of a connection, the first packet will be detected as Office 365 and the end of connection will be detected as Office 365 Admin Portal. This should not affect blocking.

Skype

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.

GoToMeeting

In order to fully detect GoToMeeting, your rule must include all of these applications:

- GoToMeeting
- Citrix Online
- Citrix GoToMeeting Platform
- LogMeIn
- STUN

Zoho mail

To control Zoho mail, choose *both* **Zoho** and **Zoho mail** from the Available Application list.

Evasive application detection considerations

For evasive applications such as Bittorrent, Tor, Psiphon, and Ultrasurf, only the highest-confidence scenarios are detected by default. If you need to take action on this traffic (such as block or implement QoS), it may be necessary to configure more aggressive detection with better effectiveness. To do this, contact TAC to review your configurations as these changes may result in false positives.

WeChat

It is not possible to selectively block WeChat Media if you allow WeChat.

RDP file transfer

If allowing the RDP (Remote Desktop Protocol) application does not allow file transfers, ensure that the rule for RDP includes both the TCP and UDP port 3389. RDP file transfer uses UDP.

Best practices for access control rules

Configuring and ordering access control rules properly is essential to protect your network. Consider these best practices to achieve optimal results:

- Always evaluate device compatibility of access control rules before deployment.
- Ensure resources, such as physical memory and processors, can handle expanded rule criteria.

Note Ensure the system evaluates all rules together and creates a set of criteria for devices when you deploy configuration changes. You cannot deploy to a device if these criteria exceed its resources.

General best practice for access control

General best practices

- Use a prefilter policy to provide early blocking for unwanted traffic, and to fastpath traffic that does not benefit from access control inspection. For more information, refer to Best practices for fastpath prefiltering.

- Access control rules are deployed as access control entries (ACE) on the device. To minimize the number of access control entries created per access control rule, and improve overall performance, enable object group search for each device. Object group search is a device setting, not an access control policy setting, so you must edit each device to ensure the feature is enabled. For more information, refer to Configure object group search.
- 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.
- Certain features, including URL filtering, application detection, and rate limiting, must allow some packets to pass in order for the system to identify the traffic.
- If you will use syslog or store events externally, avoid special characters in object names such as policy and rule names. Object names should not contain special characters, such as commas, that the receiving application may use as separators.

Licensing and feature support

- Although you can configure the system without licensing your deployment, many features require that you enable the appropriate licenses before you deploy.
- Some features are only available on certain device models. Warning icons and confirmation dialog boxes designate unsupported features.

Deployment considerations

- When you deploy an access control policy, its rules are not applied to existing connections. Traffic on an existing connection is not bound by the new policy that is deployed. In addition, the policy hit count is incremented only for the first packet of a connection that matches a policy. Thus, the traffic on an existing connection that could match a policy is omitted from the hit count. To have the policy rules effectively applied, clear the existing connection sessions, and then deploy the policy.
- For the system to affect traffic, you must deploy relevant configurations to managed devices using routed, switched, or transparent interfaces, or inline interface pairs.
- The system might prevent you from deploying inline configurations to passively deployed devices, including inline devices in tap mode. In other cases, the policy might deploy successfully, but attempts to block or alter traffic can produce unexpected results because passive devices do not enforce traffic actions. For example, the system might report multiple beginning-of-connection events for a connection that the policy attempts to block.

Best practices for creating, ordering, and implementing access control rules are detailed in Best practices for access control rules, on page 17 and subtopics.

Best practices for ordering rules

Access control rules are matched top down, first match wins. Thus, to get the results you expect, you must pay careful attention to proper sequencing.

- In general, place top-priority rules that must apply to all traffic near the top of the policy. Always order rules to suit your organization's needs.

- Specific rules should come before general rules, especially when the specific rules are exceptions to general rules. Otherwise, traffic will match the general rule first and never hit the applicable specific rule.
- 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. Rules based on these criteria do not require inspection to identify matching connections.
- Whenever possible, put specific drop rules near the top of the policy. This ensures the earliest possible decision on undesirable traffic.
- URL filtering, application-based, and geolocation-based rules and others that require inspection should come after rules that drop traffic based on layer-3/4 criteria only (such as IP address, security zone, and port number), but before rules that specify file and intrusion policies.
- For the most effective URL matching, place rules that include URL conditions before other rules, particularly if the URL rules are block rules and the other rules meet both of the following criteria:
  - They include application conditions.
  - The traffic to be inspected is encrypted.
- Put URL filtering rules above application rules, and follow application rules with micro-application rules and Common Industrial Protocol (CIP) sub-classification application filtering rules.
- In general, rules with application conditions 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. For more information, see Best practice for choosing between application matching and port matching, on page 16 and Recommendations for application control, on page 14.
- Rules that specify file policies and intrusion policies should come at the bottom of the rule order. These rules require resource-intensive deep inspection, and you should eliminate as many threats as possible using less-intensive methods first, for performance reasons, in order to minimize the number of potential threats that require deep inspection.

Application rule order

Place rules with application conditions 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.

For more information, refer to Best practice for choosing between application matching and port matching, on page 16 and Recommendations for application control, on page 14.

Rule preemption

Rule preemption is a policy evaluation condition that occurs when a rule will never match traffic because a rule earlier in the evaluation order matches the traffic first. A rule's conditions govern whether it preempts other rules.

Rule preemption scenarios

Any type of rule condition can preempt a subsequent rule. This table shows common preemption scenarios:

Table 1: Rule preemption examples

Rule 1 (preempting)	Rule 2 (preempted)	Rule preemption scenario
Allow Admin users	Block Admin users	The second rule cannot block Admin traffic because the first rule allows it.
Allow Source Network VLAN 22-33	Block Source Network VLAN 27	The VLAN range in the first rule includes the VLAN in the second rule, so the first rule preempts the second. To block VLAN 27, you must move that rule above the one that allows VLAN 22-33.
Allow Source Network 10.4.0.0/16	Allow Source Network 10.4.0.0/16, VLAN 2	The VLAN range in the first rule includes the VLAN in the second rule, so the first rule preempts the second. To block VLAN 27, you must move that rule above the one that allows VLAN 22-33.
Allow Source Network 10.10.10.0/24 to URL www.netflix.com	Allow Source Network 10.10.10.0/24 to URL www.netflix.com	A rule also preempts an identical subsequent rule where all configured conditions are the same.

Best practice for rule actions and rule order

A rule's action determines how the system handles matching traffic. Improve performance by placing rules that do not perform or ensure further traffic handling before the resource-intensive rules that do. Then, the system can divert traffic that it might otherwise have inspected.

The following examples show how you might order rules in various policies, given a set of rules where none is more critical and preemption is not an issue.

If your rules include application conditions, also refer to Best practice for choosing between application matching and port matching, on page 16.

Optimum order for access control rules

Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Place access control rules that invoke deep inspection last.

1. Monitor—Rules that log matching connections, but take no other action on traffic. (However, see the important exception and caveat at Access control rule monitor action, on page 6.)
2. Trust, Block, Block with reset—Rules that handle traffic without further inspection.
3. Allow, Interactive Block (no deep inspection)—Rules that do not inspect traffic further, but allow discovery.
4. Allow, Interactive Block (deep inspection)—Rules associated with file or intrusion policies that perform deep inspection for prohibited files, malware, and exploits.

Best practices for simplifying and focusing rules

Follow these best practices to simplify configuration and ensure optimal device performance and rule clarity:

Simplify rule conditions

Minimize individual rule criteria. Use as few elements in rule conditions as possible. For example, use IP address blocks rather than individual IP addresses when defining network conditions.

Avoid redundant conditions

If a single condition is enough to match the traffic you want to handle, do not add more. Redundant conditions can make configuration unnecessarily complex and may negatively impact device performance, especially during cluster or high-availability unit re-join.

For example:

- Use security zones that represent multiple interfaces carefully. If you specify source and destination networks as conditions, and these are enough to match the traffic you are targeting, then specifying a security zone is not required.
- If you want to match a set of internal interfaces to any destination on the Internet (for example), then simply use a source security zone that includes your internal interfaces. No network or destination interface criteria are needed.

Combine elements for organization, not performance

The use of network objects containing many individual IP addresses offers organizational benefits only. It does not improve device performance compared to listing those IP addresses individually.

For recommendations related to application detection, refer to Best practice for choosing between application matching and port matching, on page 16.

Narrowly constrain resource-intensive rules, especially by interface

As much as possible, use rule conditions to narrowly define the traffic handled by resource-intensive rules. Focused rules are also important because rules with broad conditions can match many different types of traffic, and can preempt later, more specific rules. Examples of resource-intensive rules include:

- TLS/SSL rules that decrypt traffic—Not only the decryption, but further analysis of the decrypted traffic, requires resources. Narrow focus, and where possible, block or choose not to decrypt encrypted traffic.

  Certain Firewall Threat Defense models perform TLS/SSL encryption and decryption in hardware, which improves performance significantly. For more information, refer to TLS crypto acceleration.

- Access control rules that invoke deep inspection—Intrusion, file, and malware inspection requires resources, especially if you use multiple custom intrusion policies and variable sets. Make sure you only invoke deep inspection where required.
- If you specify security zones in a rule, the rule is deployed only to devices that have interfaces in the specified zone. So, if you want to focus a rule on only some devices that are assigned to the policy, ensure that you select a security zone that applies to the appropriate subset of devices. This ensures that unnecessary rules are not deployed to a device.

Maximum number of access control rules and intrusion policies

The maximum number of access control rules or intrusion policies that are supported by a device depends on many factors, including policy complexity, physical memory, and the number of processors on the device.

If you exceed the maximum supported by a device, you cannot deploy the access control policy and you must reevaluate.

Guidelines for intrusion policies

Consider these guidelines when configuring intrusion policies:

- In an access control policy, you can associate one intrusion policy with each Allow and Interactive Block rule, as well as with the default action. Every unique **pair** of intrusion policy and variable set counts as one policy.
- You may want to consolidate intrusion policies or variable sets so you can associate a single intrusion policy-variable set pair with multiple access control rules. On some devices you may find you can use only a single variable set for all your intrusion policies, or even a single intrusion policy-variable set pair for the whole device.

Access control rule management

Access control rule management is a security administration process that involves procedures for creating, modifying, and maintaining access control rules.

These topics explain how to manage access control rules.

Add an access control rule category

Categories help organize rules into logical groups so you can understand and manage your policy more easily.

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 Firewall Management Center assigns rule numbers across sections and categories.

Procedure

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

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.

What to do next

You can now:

- Drag and drop rules into, or out of, the category.
- When creating a rule, select a category it should be part of.
- When editing a rule, reposition the rule to be in the category, or to move it out of the category.

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.

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.
  - **Copy Rule to Different Policy**, to copy the rule to another access control policy. A dialog box opens so that you can select the policy and chose the rule placement.
  - **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/Bulk Rule Actions** 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.

Table 2: Access control rule components

Rule component	Configuration description
Position	Specify the rule position (**Insert** for new rules, the **Reposition** icon next to the rule name for existing rules); see Access control rule order, on page 5.
Action	Choose a rule **Action**; see Access control rule actions, on page 6.
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.
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 Create a time range object.
Enable Rule	Whether the rule is active. Disabled rules are not applied to connections. You can disable rules to temporarily turn them off, for example, during troubleshooting.
Security Protections (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.
Conditions	Select the objects you want to add to 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 25 for more information.
Comments	Open the comments list at the bottom of the dialog box, enter your comment, and click **Add Comment/Post** to add a comment.

Step 4 Click **Add** or **Apply** to save the 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 Configure a time zone for a device.

Deploy configuration changes; see Deploy configuration changes.

Related Topics

Best practices for access control rules, on page 17

Access control rule conditions

Access control rule conditions are traffic matching criteria that 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

Security/tunnel zone rule conditions are traffic matching mechanisms that

- use security zones and tunnel zones to select traffic for a rule
- segment your network to help you manage and classify traffic flow by grouping interfaces across multiple devices
- 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.

Network criteria selection

- 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 these 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 these 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).

Traffic matching configuration:

- 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.

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 tag rule conditions

VLAN tags rule conditions are filtering mechanisms that

- control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic
- use the innermost VLAN tag to filter VLAN traffic, with the exception of prefilter policy which uses the outermost VLAN tag, and
- apply only to inline sets and do not match traffic on firewall interfaces.

VLAN tag rule conditions configuration details

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

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 this 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

A user rule condition is a traffic matching mechanism that

- matches traffic based on whether the user initiating the connection or the group to which they belong
- can be configured for users in Microsoft Active Directory realms only, and
- enables policy enforcement based on user identity and group membership.

Special identities users

In addition to configuring users and groups for configured realms, you can set policies for the special identities users:

- Failed Authentication: User who 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 such as users who 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 Associate prefilter, decryption, and identity policies with an access control.

User rule condition example

Configure a Block rule to prohibit anyone in the Finance group from accessing a network resource.

Application rule conditions

Application rule conditions are application awareness and control mechanisms that

- identify and classify applications through IP traffic analysis
- enable control with organized filters
- require one enabled detector for each condition.

Application filters and control benefits

System-provided *application filters* help perform control by organizing applications according to basic characteristics, such as type, risk, business relevance, category, and tags. You can also create reuseable, user-defined filters using combinations of system-provided filters, or custom combinations of applications.

At least one detector must be enabled for each application rule condition in the policy. If no detector is enabled, the system automatically enables all available system-provided detectors, or the most recently modified user-defined detector. For more information about application detectors, see Application detectors.

You can utilize both application filters and individually specified applications to ensure complete coverage. Understand the application's requirements to effectively order access control rules.

Application filters help you to quickly configure application control effectively. 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.

The system characterizes detected applications using several criteria which form effective filters.

Table 3: 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 applications 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.
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 `high bandwidth` and `displays ads`.

Configure 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. It is recommended that you 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) in the access control policy's advanced settings for access control rules to perform application control.

Follow these steps to configure application conditions and filters:

Procedure

Step 1 In the access policy rule editor, click **Applications**.

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. A **lock** icon 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/Selection** to clear your current choices.

Step 4 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 (concept)

A port, protocol, and ICMP code rule condition is a traffic matching criterion that

- matches traffic based on source and destination ports
- supports TCP, UDP, ICMP, and other protocols that do not use ports, and
- controls traffic flow in access control and other policy rules.

Port condition types

Port conditions match traffic based on the source and destination ports. Depending on the rule type, "port" can represent any of these types:

- TCP and UDP—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—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—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 method used to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Use application filtering criteria when possible; however, note that application filtering is unavailable in prefilter rules.

Use application filtering for FTP and similar applications that open separate channels for control and 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 configuring port conditions can match non-TCP traffic, restrictions apply:

- Access control rules—For Classic devices, 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, add only network-based conditions: zone, IP address, port, and VLAN tag. 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 only TCP port conditions.
- 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. To match any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Recommendation: use URL conditions for website access control

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

For complete information, see URL filtering with category and reputation.

Dynamic attributes rule conditions

Dynamic attributes rule conditions allow you to specify access control rules based on real-time properties such as network objects, device types, locations, and security tags. Dynamic attributes can be used as source criteria and destination criteria in access control rules. These criteria help match traffic more flexibly and accurately.

Dynamic attributes configuration and behavior

You can use these types of dynamic attributes to match connections in DNS rules:

- SGT objects (Source only): These contain tags defined manually or in Cisco ISE. For more information, refer to Source and destination security group tag (SGT) matching and Security group tags .
- Location IP objects (Source only): Defined by Cisco ISE.
- Device type objects (Source only): Defined by Cisco ISE, also referred to as endpoint profile objects.

When you configure dynamic attributes for a DNS rule, objects of the same type within a source or destination list use OR logic, while objects of different types use AND logic. For 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.

When you open the **Dynamic Attributes** tab, all **Dynamic Objects** and **Security Group Tags** appear in the list. To remove objects from the list, deselect the options. To find an object, type in the search box.

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.

Enable and disable access control rules

Enable or disable access control rules to control whether they are used to evaluate network traffic.

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 rules in an access control policy, disabled rules are grayed out but can still be modified. You can also change a rule's status in the rule editor.

Procedure

Step 1 In the access control policy editor, right-click the rule you want to change 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.

Copy 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 Rule to Different Policy**.
- To copy multiple rules, select their checkboxes, then select **Copy Rules 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 bottom of either the Mandatory or Default sections.

Step 4 Click **Copy**.

What to do next

- Deploy configuration changes; see Deploy configuration changes.

Move 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 these conditions before you proceed:

- When moving an access control rule to a prefilter policy the layer 7 (L7) parameters, such as application or URL filterin, 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 this 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 this table.

Action in the access control rule	Enabled logging configurations in the prefilter rule
Allow	None of the check boxes are checked.
Block Block with reset Interactive Block Interactive Block with reset	• Log at Beginning of Connection • Event Viewer • Syslog Server • SNMP Trap
Trust	• 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 options:

- To move a single rule, right-click the rule and select **Move Rule to Prefilter Policy**.

- To move multiple rules, select their checkboxes, then select **Move Rules 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, either **At the bottom** or **At the top**.

Step 3 Click **Move**.

What to do next

- Deploy configuration changes; see Deploy configuration changes.

Position an access control rule

Position rules correctly within an access control policy to ensure the first match wins behavior produces the intended security outcomes.

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, on page 17.

Procedure

Step 1 Do one of these:

- 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. This action is final, you are not prompted to confirm it.
- 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. You are prompted for where you want the rule to go.
- Existing rule when editing the rule—Click the **Reposition Rule** icon next to the rule name.

Step 2 When editing or repositioning the rule, choose where you want to move or insert the rule, then click **Move**, **Confirm**, or **Reposition** (depending on what you are doing).

- 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 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.

Add comments to an access control rule

Adding comments to access control rules helps you keep track of the purpose of a rule, and changes made to the 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

These topics provide examples of access control rules.

How to control access using security zones

Configure security zones and access control rules to inspect incoming traffic for intrusions and malware while allowing unrestricted internet access for internal hosts.

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.

Control application usage

Block web-based applications, particularly anonymizers or proxies, by creating access control rules that identify and restrict traffic based on application attributes

Enterprise networks increasingly rely on web-based delivery platforms and rich media applications, often using the same web protocols. Firewall Threat Defense inspects connections to determine the application, enabling access control rules targeting applications rather than just targeting specific TCP/UDP ports.

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** > **Access Control heading** > **Access Control** and edit the access control policy.

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

a) Name the rule (for example, Block_Anonymizers).

b) For **Action**, select **Block**.

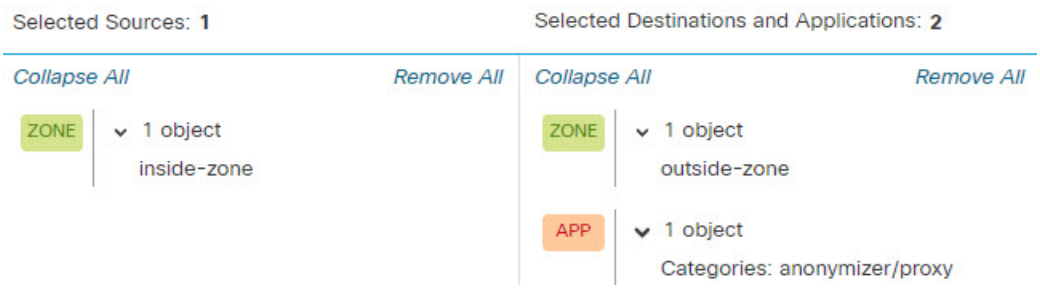


c) If you have configured security zones and want to restrict traffic from inside to outside, click the **Zones** tab, and choose your inside zone as the source zone, and the outside zone as the destination zone.

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.

For purposes of this example, choose **anonymizers/proxy** from the **Categories** list and add it to **Destinations and Applications**. The match criteria should now look similar to the following graphic.



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.

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.

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.

Security Intelligence policies allow you to preemptively block traffic to or from known bad IP addresses and URLs. This task demonstrates how to add an intrusion policy that allows traffic from the internal 192.168.1.0/24 network to external destinations, assuming block rules exist to eliminate unwanted connections, and adds Security Intelligence to block known threats.

Before you begin

- Apply the IPS license to any managed device that uses this rule.
- Ensure you have created security zones for inside and outside interfaces.
- Create the network object for the inside network.

Procedure

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

a) Edit the access control policy and add a new rule.

b) Name the rule (for example, Inside_Outside) and set the rule action to **Allow**.



c) Choose **Balanced Security and Connectivity** as the **Intrusion Policy**. 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.

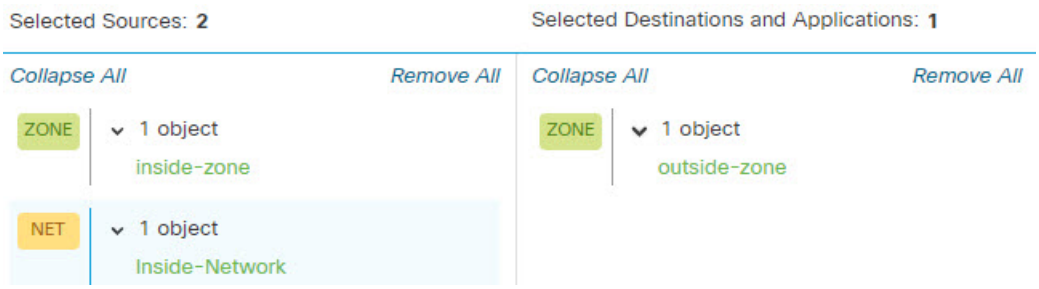
A discussion of variable sets is beyond the scope of this example. Please read the chapters on intrusion policy for detailed information about variable sets and custom policies.



d) Select the **Zones** tab, and add your inside security zone to the source criteria, and outside zone to the destination criteria.

e) 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:



f) Click **Logging** and enable logging at the beginning or end of the connection, or both, as desired.

g) Click **Apply** to save the rule, and then **Save** to save the updated policy.

h) 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 Networks list until you get to the global feeds, 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

This history table provides a summary of feature changes by release for the access control policy rules.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Maximum objects per match criteria per access control rule is 200.	7.3	Any	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. We updated the access control policy to allow the increased object limit.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
Search for access control rule comments	6.7	Any	The **Search Rules** bar now offers the option to search for comments. New/modified pages: Access control rules page, **Search Rules** text entry field. Supported platforms: Firewall Management Center
Copy or move rules between access control and prefilter policies	6.7	Any	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. New/modified pages: Access control policy page; the right-click menu for the selected rules provides additional options to copy and move. Supported platforms: Firewall Management Center
Bulk edit of certain settings in access control rules	6.6	Any	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. New/modified pages: Access control rules page. Supported platforms: Firewall Management Center
Enhanced searching on configured rules	6.6	Any	Enhanced searching on configured rules. New/modified pages: Access control rules page. Supported platforms: Firewall Management Center
Time ranges for rule application	6.6	Any	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. New/modified pages: • 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. Supported platforms: Firewall Threat Defense devices only
View object details from access control rule pages	pre-6.6	Any	To see information about an object in a list of rules or from the rule configuration dialog, right-click the object. New/modified pages: **Policies** > **Access Control** > **Access control**, and **Add Rule** page. Supported platforms: Firewall Management Center