

Controlling Traffic with Reputation-Based Rules

Access control rules within *access control policies* exert granular control over network traffic logging and handling. Reputation-based conditions in access control rules allow you to manage which traffic can traverse your network, by contextualizing your network traffic and limiting it where appropriate. Access control rules govern the following types of reputation-based control:

- Application conditions allow you to perform *application control*, which controls application traffic based on not only individual applications, but also applications' basic characteristics: type, risk, business relevance, categories, and tags.
- URL conditions allow you to perform *URL filtering*, which controls web traffic based on individual websites, as well as websites' system-assigned category and reputation.

You can combine reputation-based conditions with each other and with other types of conditions to create an access control rule. These access control rules can be simple or complex, matching and inspecting traffic using multiple conditions. For detailed information on access control rules, see Tuning Traffic Flow Using Access Control Rules, page 6-1.

Note Security Intelligence-based traffic filtering, and some decoding and preprocessing occur **before** network traffic is evaluated by access control rules. You can also configure the *SSL inspection* feature to block or decrypt encrypted traffic before access control rules evaluate it.

Reputation-based access control requires the following licenses.

Table 8-1 License Requirements for Reputation-Based Access Control Rules

Requirement	Application Control	URL Filtering (cat. & rep.)	URL Filtering (manual)
license	Control	URL Filtering	Any

For information on adding reputation-based conditions to access control rules, see:

- Controlling Application Traffic, page 8-2
- Blocking URLs, page 8-7

The ASA FirePOWER module can perform other types of reputation-based control, but you do not configure these using access control rules. For more information, see:

- Blacklisting Using Security Intelligence IP Address Reputation, page 5-1 explains how to limit traffic based on the reputation of a connection's origin or destination as a first line of defense.
- Tuning Intrusion Prevention Performance, page 11-6 explains how to detect, track, store, analyze, and block the transmission of malware and other types of prohibited files.

Controlling Application Traffic

License: Control

When the ASA FirePOWER module analyzes IP traffic, it can identify and classify the commonly used applications on your network.

Understanding Application Control

Application conditions in access control rules allow you to perform this *application control*. Within a single access control rule, there are a few ways you can specify applications whose traffic you want to control:

- You can select individual applications, including custom applications.
- You can use system-provided *application filters*, which are named sets of applications organized according to the applications' basic characteristics: type, risk, business relevance, categories, and tags.
- You can create and use custom application filters, which group applications (including custom applications) in any way you choose.

Application filters allow you to quickly create application conditions for access control rules. They simplify policy creation and administration, and grant you assurance that the system will control web traffic as expected. For example, you could create an access control rule that identifies and blocks all high risk, low business relevance applications. If a user attempts to use one of those applications, the session is blocked.

In addition, Cisco frequently updates and adds additional detectors via system and vulnerability database (VDB) updates. By using filters based on application characteristics, you can ensure that the system uses the most up-to-date detectors to monitor application traffic.

Building Application Conditions

For traffic to match an access control rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.

In a single application condition, you can add a maximum of 50 items to the **Selected Applications and Filters** list. Each of the following counts as an item:

- One or more filters from the **Application Filters** list, individually or in custom combination. This item represents set of applications, grouped by characteristic.
- A filter created by saving an application search in the **Available Applications** list. This item represents a set of applications, grouped by substring match.
- An individual application from the **Available Applications** list.

In the module interface, filters added to a condition are listed above and separately from individually added applications.

Note that when you deploy an access control policy, for each rule with an application condition, the system generates a list of unique applications to match. In other words, you may use overlapping filters and individually specified applications to ensure complete coverage.

Note For encrypted traffic, the system can identify and filter traffic using only the applications tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic.

For more information, see the following sections:

- Matching Traffic with Application Filters, page 8-3
- Matching Traffic from Individual Applications, page 8-4
- Adding an Application Condition to an Access Control Rule, page 8-5
- Limitations to Application Control, page 8-6

Matching Traffic with Application Filters

License: Control

When building an application condition in an access control rule, use the **Application Filters** list to create a set of applications, grouped by characteristic, whose traffic you want to match.

Note that the mechanism for filtering applications within an access control rule is the same as that for creating reusable, custom application filters using the object manager; see Working with Application Filters, page 3-11. You can also save many filters you create on-the-fly in access control rules as new, reusable filters. You cannot save a filter that includes another user-created filter because you cannot nest user-created filters.

Understanding How Filters Are Combined

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

The system links multiple filters of the same filter type with an OR operation. For example, if you select the Medium and High filters under the Risks type, the resulting filter is:

```
Risk: Medium OR High
```

If the Medium filter contains 110 applications and the High filter contains 82 applications, the system displays all 192 applications in the **Available Applications** list.

The system links different types of filters with an AND operation. For example, if you select the Medium and High filters under the Risks type, and the Medium and High filters under the Business Relevance type, the resulting filter is:

```
Risk: Medium OR High
AND
Business Relevance: Medium OR High
```

In this case, the system displays only those applications that are included in both the Medium or High Risk type AND the Medium or High Business Relevance type.

Finding and Selecting Filters

To select filters, click the arrow next to a filter type to expand it, then select or clear the check box next to each filter whose applications you want to display or hide. You can also right-click a system-provided filter type (**Risks**, **Business Relevance**, **Types**, **Categories**, or **Tags**) and select **Check All** or **Uncheck All**.

To search for filters, click the **Search by name** prompt above the **Available Filters** list, then type a name. The list updates as you type to display matching filters.

After you are done selecting filters, use the **Available Applications** list to add those filters to the rule; see Matching Traffic from Individual Applications, page 8-4.

Matching Traffic from Individual Applications

License: Control

When building an application condition in an access control rule, use the **Available Applications** list to select the applications whose traffic you want to match.

Browsing the List of Applications

When you first start to build the condition the list is unconstrained, and displays every application the system detects, 100 at a time:

- To page through the applications, click the arrows underneath the list.
- To display a pop-up window with summary information about the application's characteristics, as well as Internet search links that you can follow, click the information icon ( ) next to an application.

Finding Applications to Match

To help you find the applications you want to match, you can constrain the **Available Applications** list in the following ways:

- To search for applications, click the **Search by name** prompt above the list, then type a name. The list updates as you type to display matching applications.
- To constrain the applications by applying a filter, use the **Application Filters** list (see Matching Traffic with Application Filters, page 8-3). The **Available Applications** list updates as you apply filters.

Once constrained, an **All apps matching the filter** option appears at the top of the **Available Applications** list. This option allows you to add all the applications in the constrained list to the **Selected Applications and Filters** list, all at once.

Note If you select one or more filters in the Application Filters list and also search the **Available Applications** list, your selections and the search-filtered **Available Applications** list are combined using an AND operation. That is, the **All apps matching the filter** condition includes all the individual conditions currently displayed in the **Available Applications** list as well as the search string entered above the **Available Applications** list.

Selecting Single Applications to Match in a Condition

After you find an application you want to match, click to select it. To select multiple applications, use the Shift and Ctrl keys, or right-click and select **Select All** to select all applications in the current constrained view.

In a single application condition, you can match a maximum of 50 applications by selecting them individually; to add more than 50 you must either create multiple access control rules or use filters to group applications.

Selecting All Applications Matching a Filter for a Condition

Once constrained by either searching or using the filters in the **Application Filters** list, the **All apps matching the filter** option appears at the top of the **Available Applications** list.

This option allows you to add the entire set of applications in the constrained **Available Applications** list to the **Selected Applications and Filters** list, at once. In contrast to adding applications individually, adding this set of applications counts as only one item against the maximum of 50, regardless of the number of individual applications that comprise it.

When you build an application condition this way, the name of the filter you add to the **Selected Applications and Filters** list is a concatenation of the filter types represented in the filter plus the names of up to three filters for each type. More than three filters of the same type are followed by an ellipsis (...). For example, the following filter name includes two filters under the Risks type and four under Business Relevance:

```
Risks: Medium, High Business Relevance: Low, Medium, High,...
```

Filter types that are not represented in a filter you add with **All apps matching the filter** are not included in the name of the filter you add. These filter types are set to *any*; that is, these filter types do not constrain the filter, so any value is allowed for these.

You can add multiple instances of **All apps matching the filter** to an application condition, with each instance counting as a separate item in the **Selected Applications and Filters** list. For example, you could add all high risk applications as one item, clear your selections, then add all low business relevance applications as another item. This application condition matches applications that are high risk OR have low business relevance.

Adding an Application Condition to an Access Control Rule

License: Control

For traffic to match an access control rule with an application condition, the traffic must match one of the filters or applications that you add to a **Selected Applications and Filters** list.

You can add a maximum of 50 items per condition, and filters added to a condition are listed above and separately from individually added applications. When building an application condition, warning icons indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To control application traffic:

Step 1 In the access control policy where you want to control traffic by application, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, click the **Applications** tab.

Step 3 Optionally, enable content restriction features by clicking the dimmed icons for Safe Search ( ) or YouTube EDU ( ) and setting related options; for additional configuration requirements, see Using Access Control Rules to Enforce Content Restriction, page 13-1.

In most cases, enabling content restriction populates the condition's **Selected Applications and Filters** list with the appropriate values. The system does not automatically populate the list if applications or filters related to content restriction are already present in the list when you enable content restriction.

Continue with the procedure to refine your application and filter selections, or skip to saving the rule.

Step 4 Optionally, use filters to constrain the list of applications displayed in the **Available Applications** list.

Select one or more filters in the **Application Filters** list. For more information, see Matching Traffic with Application Filters, page 8-3.

Step 5 Find and select the applications you want to add from the **Available Applications** list.

You can search for and select individual applications, or, when the list is constrained, **All apps matching the filter**. For more information, see Matching Traffic from Individual Applications, page 8-4.

Step 6 Click **Add to Rule** to add the selected applications to the **Selected Applications and Filters** list.

You can also drag and drop selected applications and filters. Filters appear under the heading *Filters*, and applications appear under the heading *Applications*.

Tip Before you add another filter to this application condition, click **Clear All Filters** to clear your existing selections.

Step 7 Optionally, click the add icon ( ) above the **Selected Applications and Filters** list to save a custom filter comprised of all the individual applications and filters currently in the list.

Use the object manager to manage this on-the-fly-created filter; see Working with Application Filters, page 3-11. Note that you cannot save a filter that includes another user-created filter; you cannot nest user-created filters.

Step 8 Save or continue editing the rule.

You must deploy the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Limitations to Application Control

License: Control

Keep the following points in mind when performing application control.

Speed of Application Identification

The system cannot perform application control before:

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

This identification should occur within 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted. If one of these first packets matches all other conditions in an access control rule containing an application condition but the identification is not complete, the access control policy allows the packet to pass. This behavior allows the connection to be established so that applications can be identified. For your convenience, affected rules are marked with an information icon ( ).

The allowed packets are inspected by the access control policy's *default* intrusion policy (not the *default action* intrusion policy nor the almost-matched rule's intrusion policy). For more information, see Setting the Default Intrusion Policy for Access Control, page 20-1.

After the system completes its identification, the system applies the access control rule action, as well as any associated intrusion and file policy, to the remaining session traffic that matches its application condition.

Handling Encrypted Traffic

The system can identify and filter unencrypted application traffic that becomes encrypted using StartTLS, such as SMTPS, POPS, FTPS, TelnetS, and IMAPS. In addition, it can identify certain encrypted applications based on the Server Name Indication in the TLS client hello message, or the server certificate subject distinguished name value.

These applications are tagged **SSL Protocol**. Applications without this tag can only be detected in unencrypted or decrypted traffic.

Handling Application Traffic Packets Without Payloads

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 Traffic

To create a rule to act on traffic referred by a web server, such as advertisement traffic, add a condition for the referred application rather than the referring application.

Controlling Application Traffic That Uses Multiple Protocols (Skype)

The system can detect multiple types of Skype application traffic. When building an application condition to control Skype traffic, select 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. For more information, see Matching Traffic with Application Filters, page 8-3.

Blocking URLs

License: feature dependent

URL conditions in access control rules allow you to limit the websites that users on your network can access. This feature is called *URL filtering*. There are two ways you can use access control to specify URLs you want to block (or, conversely, allow):

- With any license, you can manually specify individual URLs or groups of URLs to achieve granular, custom control over web traffic.
- With a URL Filtering license, you can also control access to websites based on the URL's general classification, or *category*, and risk level, or *reputation*. The system displays this category and reputation data in connection logs, intrusion events, and application details.

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

When you block a website, you can either allow the user's browser its default behavior, or you can display a generic system-provided or custom page. You can also give users a chance to bypass a website block by clicking through a warning page.

Table 8-2 License Requirements for URL Filtering

Requirement	Category & Reputation-Based	Manual
license	URL Filtering	Any

For more information, see:

- Performing Reputation-Based URL Blocking, page 8-8
- Performing Manual URL Blocking, page 8-10
- Limitations to URL Detection and Blocking, page 8-11
- Allowing Users to Bypass URL Blocks, page 8-12
- Displaying a Custom Web Page for Blocked URLs, page 8-14

Performing Reputation-Based URL Blocking

License: URL Filtering

With a URL Filtering license, you can control your users' access to websites based on the category and reputation of requested URLs, which the ASA FirePOWER module obtains from the Cisco cloud:

- The URL *category* is a general classification for the URL. For example, ebay.com belongs to the **Auctions** category, and monster.com belongs to the **Job Search** category. A URL can belong to more than one category.
- The URL *reputation* represents how likely the URL is to be used for purposes that might be against your organization's security policy. A URL's risk can range from **High Risk** (level 1) to **Well Known** (level 5).

Note Before access control rules with category and reputation-based URL conditions can take effect, you **must** enable communications with the Cisco cloud. This allows the ASA FirePOWER module to retrieve URL data. For more information, see Enabling Cloud Communications, page 44-2.

Advantages to Reputation-Based URL Blocking

URL categories and reputations allow you to quickly create URL conditions for access control rules. For example, you could create an access control rule that identifies and blocks all **High Risk** URLs in the **Abused Drugs** category. If a user attempts to browse to any URL with that category and reputation combination, the session is blocked.

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

Some examples include:

- If a rule blocks all gaming sites, as new domains get registered and classified as **Gaming**, the system can block those sites automatically.
- If a rule blocks all malware sites, and a blog page gets infected with malware, the cloud can recategorize the URL from **Blog** to **Malware** and the system can block that site.
- If a rule blocks high-risk social networking sites, and somebody posts a link on their profile page that contains links to malicious payloads, the cloud can change the reputation of that page from **Benign sites** to **High Risk** so the system can block it.

Note that if the cloud does not know the category or reputation of a URL, or if the ASA FirePOWER module cannot contact the cloud, that URL does **not** trigger access control rules with category or reputation-based URL conditions. You cannot assign categories or reputations to URLs manually.

Building URL Conditions

You can add a maximum of 50 items to the **Selected URLs** to match in a single URL condition. Each URL category, optionally qualified by reputation, counts as a single item. Note that you can also use literal URLs and URL objects in URL conditions, but you cannot qualify these items with a reputation. For more information, see Performing Manual URL Blocking, page 8-10.

Note that you cannot qualify a literal URL or URL object with a reputation.

When building a URL condition, warning icons indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To control traffic by requested URL using category and reputation data:

Step 1 In the access control policy where you want to control traffic by URL, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, select the URLs tab.

The URLs tab appears.

Step 3 Find and select the categories of URL you want to add from the **Categories and URLs** list. To match web traffic regardless of category, select **Any** category.

To search for categories to add, click the **Search by name or value** prompt above the **Categories and URLs** list, then type the category name. The list updates as you type to display matching categories.

To select a category, click it. To select multiple categories, use the Shift and Ctrl keys.

Tip Although you can right-click and **Select All** categories, adding all categories this way exceeds the 50-item maximum for an access control rule. Instead, use **Any**.

Step 4 Optionally, qualify your category selections by clicking a reputation level from the **Reputations** list. If you do not specify a reputation level, the system defaults to **Any**, meaning all levels.

You can only select one reputation level. When you choose a reputation level, the access control rule behaves differently depending on its purpose:

- If the rule blocks or monitors web access (the rule action is **Block**, **Block with reset**, **Interactive Block**, **Interactive Block with reset**, or **Monitor**) selecting a reputation level also selects all reputations more severe than that level. For example, if you configure a rule to block or monitor **Suspicious sites** (level 2), it also automatically blocks or monitors **High risk** (level 1) sites.
- If the rule allows web access, whether to trust or further inspect it (the rule action is **Allow** or **Trust**), selecting a reputation level also selects all reputations less severe than that level. For example, if you configure a rule to allow **Benign sites** (level 4), it also automatically allows **Well known** (level 5) sites.

If you change the rule action for a rule, the system automatically changes the reputation levels in URL conditions according to the above points.

Step 5 Click **Add to Rule** or drag and drop the selected items to add them to the **Selected URLs** list.

Step 6 Save or continue editing the rule.

You must deploy the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Performing Manual URL Blocking

License: Any

To supplement or selectively override URL filtering by category and reputation, you can control web traffic by manually specifying individual URLs or groups of URLs. This allows you to achieve granular, custom control over allowed and blocked web traffic. You can also perform this type of URL filtering without a special license.

To manually specify URLs to allow or block in an access control rule, you can type in a single literal URL. Or, you can configure URL conditions using URL objects, which are reusable and associate a name with a URL or IP address.

Tip After you create a URL object, you can use it not only to build access control rules, but also to represent URLs in various other places in the system's module interface. You can create these objects using the object manager; you can also create URL objects on-the-fly while you are configuring access control rules. For more information, see Working with URL Objects, page 3-10.

Manually Specifying URLs in URL Conditions

Although manual entry gives you precise control over allowed and blocked web traffic, you cannot qualify a manually specified URL with a reputation. Additionally, you must make sure that your rules do not have unintended consequences. To determine whether network traffic matches a URL condition, the system performs a simple substring match. If the value of a URL object or manually typed URL matches any part of a URL requested by a monitored host, the URL condition of the access control rule is satisfied.

Therefore, when manually specifying URLs in URL conditions, including in URL objects, carefully consider other traffic that might be affected. For example, if you allow all traffic to example.com, your users could browse to URLs including:

- http://example.com/
- http://example.com/newexample
- http://www.example.com/

As another example, consider a scenario where you want to explicitly block ign.com (a gaming site). However, substring matching means that blocking *ign.com* also blocks veris*ign.com*, which might not be your intent.

Manually Blocking Encrypted Web Traffic

URL conditions in access control rules:

- disregard the encryption protocol of web traffic (HTTP vs HTTPS)

  For example, access control rules treat traffic to http://example.com/ the same as traffic to https://example.com/. To configure an access control rule that matches only HTTP or HTTPS traffic, add an application condition to the rule. For more information, see Blocking URLs, page 8-7.
- match HTTPS traffic based on the subject common name in the public key certificate used to encrypt the traffic, and also disregard subdomains within the subject common name

  Do not include subdomain information when manually filtering HTTPS traffic.

When building a URL condition, warning icons indicate invalid configurations. For details, see Troubleshooting Access Control Policies and Rules, page 4-12.

To control web traffic by manually specifying URLs to allow or block:

Step 1 In the access control policy where you want to control traffic by URL, create a new access control rule or edit an existing rule.

For detailed instructions, see Creating and Editing Access Control Rules, page 6-2.

Step 2 In the rule editor, select the URLs tab.

The URLs tab appears.

Step 3 Find and select the URL objects and groups you want to add from the **Categories and URLs** list:

- To add a URL object on the fly, which you can then add to the condition, click the add icon ( ) above the **Categories and URLs** list; see Working with URL Objects, page 3-10.
- To search for URL objects and groups to add, click the **Search by name or value** prompt above the **Categories and URLs** list, then type either the name of the object, or the value of a URL or IP address in the object. The list updates as you type to display matching objects.

To select an object, click it. To select multiple objects, use the Shift and Ctrl keys. Although you can right-click and **Select All** URL objects and categories, adding URLs this way exceeds the 50-item maximum for an access control rule.

Step 4 Click **Add to Rule** or to add the selected items to the **Selected URLs** list.

You can also drag and drop selected items.

Step 5 Add any literal URLs that you want to specify manually. You **cannot** use wildcards (*) in this field.

Click the **Enter URL** prompt below the **Selected URLs** list; then type a URL or IP address and click **Add**.

Step 6 Save or continue editing the rule.

You must deploy the access control policy for your changes to take effect; see Deploying Configuration Changes, page 4-11.

Limitations to URL Detection and Blocking

License: Any

Keep the following points in mind when performing URL detection and blocking.

Speed of URL Identification

The system cannot filter URLs before:

- a monitored connection is established between a client and server
- the system identifies the HTTP or HTTPS application in the session
- the system identifies the requested URL (for encrypted sessions, from either the client hello message or the server certificate)

This identification should occur within 3 to 5 packets, or after the server certificate exchange in the SSL handshake if the traffic is encrypted. If one of these first packets matches all other conditions in an access control rule containing a URL condition but the identification is not complete, the access control policy allows the packet to pass. This behavior allows the connection to be established so that URLs can be identified. For your convenience, affected rules are marked with an information icon ( ).

The allowed packets are inspected by the access control policy's *default* intrusion policy (not the *default action* intrusion policy nor the almost-matched rule's intrusion policy). For more information, see Setting the Default Intrusion Policy for Access Control, page 20-1.

After the system completes its identification, the system applies the access control rule action, as well as any associated intrusion and file policy, to the remaining session traffic that matches its URL condition.

Handling Encrypted Web Traffic

When evaluating encrypted web traffic using access control rules with URL conditions, the system:

- disregards the encryption protocol; an access control rule matches both HTTPS and HTTP traffic if the rule has a URL condition but not an application condition that specifies the protocol
- matches HTTPS traffic based on the subject common name in the public key certificate used to encrypt the traffic, and disregards subdomains within the subject common name
- does not display an HTTP response page, even if you configured one

Search Query Parameters in URLs

The system does not use search query parameters in the URL to match URL conditions. For example, consider a scenario where you block all shopping traffic. In that case, using a web search to search for amazon.com is not blocked, but browsing to amazon.com is.

Allowing Users to Bypass URL Blocks

License: Any

When you block a user's HTTP web request using an access control rule, setting the rule action to **Interactive Block** or **Interactive Block with reset** gives that user a chance to bypass the block by clicking through a warning *HTTP response page*. You can display a generic system-provided response page or you can enter custom HTML.

By default, the system allows users to bypass blocks for 10 minutes (600 seconds) without displaying the warning page on subsequent visits. You can set the duration to as long as a year, or you can force the user to bypass the block every time.

If the user does not bypass the block, matching traffic is denied without further inspection; you can also reset the connection. On the other hand, if a user bypasses the block, the system allows the traffic. Allowing this traffic means that you can continue to inspect unencrypted payloads for intrusions, malware and prohibited files. Note that users may have to refresh after bypassing the block to load page elements that did not load.

Note that you configure the interactive HTTP response page separately from the response page you configure for Block rules. For example, you could display the system-provided page to users whose sessions are blocked without interaction, but a custom page to users who can click to continue. For more information, see Displaying a Custom Web Page for Blocked URLs, page 8-14.

If you block web traffic decrypted by the SSL inspection feature, the system encrypts the response page and sends it at the end of the reencrypted SSL stream.

Tip To quickly disable interactive blocking for all rules in an access control policy, display neither the system-provided page nor a custom page. This causes the system to block all connections that match an Interactive Block rule without interaction.

To allow users to bypass a website block:

Step 1 Create an access control rule that matches web traffic with a URL condition.

See Performing Reputation-Based URL Blocking, page 8-8 and Performing Manual URL Blocking, page 8-10.

Step 2 Make sure the access control rule action is **Interactive Block** or **Interactive Block with reset.**

See Using Rule Actions to Determine Traffic Handling and Inspection, page 6-7.

Step 3 Assume users will bypass the block and choose inspection and logging options for the rule accordingly. As with Allow rules:

- You can associate either type of Interactive Block rule with a file and intrusion policy. For more information, see Controlling Traffic Using Intrusion and File Policies, page 11-1.
- Logging options for interactively blocked traffic are identical to those in allowed traffic, but keep in mind that if a user does not bypass the interactive block, the system can log only beginning-of-connection events.

  Note that when the system initially warns the user, it marks any logged beginning-of-connection event with the Interactive Block or Interactive Block with reset action. If the user bypasses the block, additional connection events logged for the session have an action of Allow. For more information, see Logging Connections Based on Access Control Handling, page 36-9.

Step 4 Optionally, set the amount of time that elapses after a user bypasses a block before the system displays the warning page again.

See Setting the User Bypass Timeout for a Blocked Website, page 8-13.

Step 5 Optionally, create and use a custom page to display to allow users to bypass a block.

See Displaying a Custom Web Page for Blocked URLs, page 8-14.

Setting the User Bypass Timeout for a Blocked Website

License: Any

By default, the system allows a user to bypass interactive blocks for 10 minutes (600 seconds) without displaying the warning page on subsequent visits. You can set the duration to as long as a year, or to zero to force the user to bypass the block every time. This limit applies to every Interactive Block rule in the policy. You cannot set the limit per rule.

To customize the length of time before a user bypass expires:

Step 1 Select **Configuration > ASA FirePOWER Configuration > Policies > Access Control Policy**.

The Access Control Policy page appears.

Step 2 Click the edit icon ( ) next to the access control policy you want to configure.

The access control policy editor appears.

Step 3 Select the **Advanced** tab.

Advanced settings for the access control policy appear.

Step 4 Click the edit icon ( ) next to General Settings.

The General Settings pop-up window appears.

Step 5 In the **Allow an Interactive Block to bypass blocking for (seconds)** field, type the number of seconds that must elapse before the user bypass expires.

You can specify any number of seconds from zero to 31536000 (one year). Specifying zero forces your users to bypass the block every time.

Step 6 Click **OK**.

Advanced settings for the access control policy appear.

Step 7 Click **Store ASA FirePOWER Changes**.

You must deploy the access control policy for your changes to take effect. For more information, see Deploying Configuration Changes, page 4-11.

Displaying a Custom Web Page for Blocked URLs

License: Any

When the system blocks a user's HTTP web request, what the user sees in a browser depends on how you block the session, using the access control rule's action. You should select:

- **Block** or **Block with reset** to deny the connection. A blocked session times out; the system resets Block with reset connections. However, for both blocking actions, you can override the default browser or server page with a custom page that explains that the connection was denied. The system calls this custom page an *HTTP response page*.
- **Interactive Block** or **Interactive Block with reset** if you want to display an *interactive HTTP response page* that warns users, but also allows them to click a button to continue or refresh the page to load the originally requested site. Users may have to refresh after bypassing the response page to load page elements that did not load.

You can either display a generic system-provided response page, or you can enter custom HTML. When you enter custom text, a counter shows how many characters you have used.

In each access control policy, you configure the interactive HTTP response page separately from the response page you use to block traffic without interaction, that is, using a Block rule. For example, you could display the system-provided page to users whose sessions are blocked without interaction, but a custom page to users who can click to continue.

Reliable display of HTTP response pages to your users depends on your network configuration, traffic loads, and size of the page. If you build a custom response page, keep in mind that a smaller page is more likely to display successfully.

To configure HTTP response pages:

Step 1 Edit the access control policy monitoring your web traffic; see Editing Access Control Policies, page 4-7.

Step 2 Click the HTTP Responses tab.

Step 3 For the **Block Response Page** and the **Interactive Block Response Page**, choose responses from the drop-down lists. For each page, you have the following choices:

- **System-provided**—Displays a generic response. Click the view icon ( ) to view the code for this page.
- **Custom**—Create a custom response page.

  A pop-up window appears, prepopulated with system-provided code that you can replace or modify. When you are done, save your changes. Note that you can edit a custom page by clicking the edit icon ( ).
- **None**—Disables the response page and blocks sessions without interaction or explanation. Note that selecting this option for interactively blocked sessions prevents users from clicking to continue; the session is blocked without interaction.

Step 4 Click **Store ASA FirePOWER Changes**.

You must redeploy the configuration for your changes to take effect. For more information, see Deploying Configuration Changes, page 4-11.