



Connection Logging

The following topics describe how to configure the system to log connections made by hosts on your monitored network:

- [About Connection Logging, on page 1](#)
- [Limitations of Connection Logging, on page 8](#)
- [Best Practices for Connection Logging, on page 8](#)
- [Requirements and Prerequisites for Connection Logging, on page 10](#)
- [Configure Connection Logging, on page 10](#)

About Connection Logging

The system can generate logs of the connections its managed devices detect. These logs are called *connection events*. Settings in rules and policies give you granular control over which connections you log, when you log them, and where you store the data. Special connection events, called *security-related connection events*, represent connections that were blocked by the reputation-based Security Intelligence feature.

Connection events contain data about the detected sessions. The information available for any individual connection event depends on several factors, but in general includes:

- Basic connection properties: timestamp, source and destination IP address, ingress and egress zones, the device that handled the connection, and so on
- Additional connection properties discovered or inferred by the system: applications, requested URLs, or users associated with the connection, and so on
- Metadata about why the connection was logged: which configuration handled the traffic, whether the connection was allowed or blocked, details about encrypted and decrypted connections, and so on

Log connections according to the security and compliance needs of your organization. When setting up connection logging, keep in mind that the system can log a connection for multiple reasons, and that disabling logging in one place does not mean that matching connections will not be logged.

The information in a connection event depends on several factors, including traffic characteristics, the configuration that ultimately handled the connection, and so on.



Note You can supplement the connection logs gathered by your managed devices with connection data generated from exported NetFlow records. This is especially useful if you have NetFlow-enabled routers or other devices deployed on networks that your managed devices cannot monitor.

Connections That Are Always Logged

Unless you disable connection event storage, the system automatically saves the following end-of-connection events to the Firewall Management Center database, regardless of any other logging configurations.

Connections Associated with Intrusions

The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action.

When an intrusion policy associated with the access control default action generates an intrusion event, the system does *not* automatically log the end of the associated connection. Instead, you must explicitly enable default action connection logging. This is useful for intrusion prevention-only deployments where you do not want to log any connection data.

However, if you enable beginning-of-connection logging for the default action, the system *does* log the end of the connection when an associated intrusion policy triggers, in addition to logging the beginning of the connection.

Connections Associated with File and Malware Events

The system automatically logs connections associated with file and malware events.



Note File events generated by inspecting NetBIOS-SSN (SMB) traffic do not immediately generate connection events because the client and server establish a persistent connection. The system generates connection events after the client or server ends the session.

Connections Associated with Intelligent Application Bypass

The system automatically logs bypassed and would-have-bypassed connections associated with IAB.

Monitored Connections

The system always logs the ends of connections for monitored traffic, even if the traffic matches no other rules and you do not enable default action logging. For more information, see [Logging for Monitored Connections, on page 4](#).

Other Connections You Can Log

So that you log only critical connections, enable connection logging on a per-rule basis. If you enable connection logging for a rule, the system logs all connections handled by that rule.

You can also log connections handled by policy default actions. Depending on the rule or default action (and for access control, a rule's inspection configuration), your logging options differ.

Prefilter Policy: Rules and Default Action

You can log connections (including entire plaintext, passthrough tunnels) that you fastpath or block with a prefilter policy.

Prefiltering uses outer-header criteria to handle traffic. For tunnels that you log, the resulting connection events contain information from the outer, encapsulation headers.

For traffic subject to further analysis, logging in the prefilter policy is disabled, although matching connections may still be logged by other configurations. The system performs all further analysis using inner headers, that is, the system independently handles and logs each connection within an allowed tunnel.

Decryption Policy: Rules and Default Action

You can log connections that match a decryption rule or decryption policy default action.

For blocked connections, the system immediately ends the session and generates an event. For monitored connections and connections that you pass to access control rules, the system generates an event when the session ends.

Access Control Policy: Security Intelligence Decisions

You can log a connection whenever it is blocked by the reputation-based Security Intelligence feature.

Optionally, and recommended in passive deployments, you can use a monitor-only setting for Security Intelligence filtering. This allows the system to further analyze connections that would have been blocked by Security Intelligence, but still log the match. Security Intelligence monitoring also allows you to create traffic profiles using Security Intelligence information.

When the system logs a connection event as the result of Security Intelligence filtering, it also logs a matching Security Intelligence event, which is a special kind of connection event that you can view and analyze separately, and that is also stored and pruned separately. So that you can identify the matching IP address in the connection, host icons beside blocked and monitored IP addresses look slightly different in the tables on the pages under the **Analysis > Connections > Events** menus.

Access Control Policy: Rules and Default Action

You can log connections that match an access control rule or access control policy default action.

Related Topics

[How Rules and Policy Actions Affect Logging](#), on page 3

How Rules and Policy Actions Affect Logging

Connection events contain metadata about why the connection was logged, including which configurations handled the traffic. Where you can configure connection logging, rule actions, and policy default actions determine not only how the system inspects and handles matching traffic, but also when and how you can log details about matching traffic.

Related Topics

[Connection and Security-Related Connection Event Fields](#)

Logging for Fastpathed Connections

You can log fastpathed connections and non-encrypted tunnels, which includes traffic matching the following rules and actions in the prefilter policy:

- Tunnel rules—**Fastpath** action (logs the outer session)
- Prefilter rules—**Fastpath** action

Fastpathed traffic bypasses the rest of access control and QoS, so connection events for fastpathed connections contain limited information.

Logging for Monitored Connections

The system always logs the ends of connections for traffic matching the following configurations, even if the traffic matches no other rules and you do not enable default action logging:

- Security Intelligence—Block lists set to monitor (also generates a Security Intelligence event)
- SSL rules—**Monitor** action
- Access control rules—**Monitor** action

The system does not generate a separate event each time a single connection matches a Monitor rule. Because a single connection can match multiple Monitor rules, each connection event can include and display information on the first eight Monitor access control rules that the connection matches, as well as the first matching SSL Monitor rule.

Similarly, if you send connection events to an external syslog or SNMP trap server, the system does not send a separate alert each time a single connection matches a Monitor rule. Rather, the alert that the system sends at the end of the connection contains information on the Monitor rules the connection matched.

Logging for Trusted Connections

You can log the beginnings and ends of trusted connections, which includes traffic matching the following rules and actions:

- Access control rules—**Trust** action
- Access control default action—**Trust All Traffic**



Note Although you *can* log trusted connections, we recommend you do not do so because trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

TCP connections detected by a Trust rule on the first packet generate only an end-of-connection event. The system generates the event one hour after the final session packet.

Logging for Blocked Connections

You can log blocked connections, which includes traffic matching the following rules and actions:

- Tunnel rules—**Block**

- Prefilter rules—**Block**
- Prefilter default action—**Block all tunnel traffic**
- Security Intelligence—Block lists not set to Monitor (also generates a Security Intelligence event)
- Decryption rules—**Block** and **Block with reset**
- SSL default action—**Block** and **Block with reset**
- Access control rules—**Block**, **Block with reset**, and **Interactive Block**
- Access control default action—**Block All Traffic**

Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.



Caution

Firewall Threat Defense device logs end-of-connection events for TCP connections where the three-way handshake does not complete, including SYN-only, SYN-SYN/ACK (no ACK), and SYN-RST/ACK scenarios (except when TLS Identity is enabled). These logs provide visibility into incomplete or suspicious connection attempts. However, logging blocked TCP connections during a Denial of Service (DoS) attack can affect system performance and overwhelm the database with multiple similar events. Before you enable logging for a Block rule, consider whether the rule monitors traffic on an Internet-facing interface or other interface vulnerable to DoS attack.

Beginning vs End-of-Connection Logging for Blocked Connections

When you log a blocked connection, how the system logs it depends on why the connection was blocked; this is important to keep in mind when configuring correlation rules based on connection logs:

- For decryption rules and decryption policy default actions that block encrypted traffic, the system logs **end-of-connection** events. This is because the system cannot determine if a connection is encrypted using the first packet in the session.
- For other blocking actions, the system logs **beginning-of-connection** events. Matching traffic is denied without further inspection.

Logging Bypassed Interactive Blocks

Interactive blocking access control rules, which cause the system to display a warning page when a user browses to a prohibited website, allow you to configure end-of-connection logging. This is because if the user clicks through the warning page, the connection is considered a new, allowed connection which the system can monitor and log.

Therefore, for packets that match an Interactive Block or Interactive Block with Reset rule, the system can generate the following connection events:

- A beginning-of-connection event when a user's request is initially blocked and the warning page is displayed; this event has an associated action of `Interactive Block` or `Interactive Block with Reset`
- Multiple beginning- or end-of-connection events if the user clicks through the warning page and loads the originally requested page; these events have an associated action of `Allow` and a reason of `User Bypass`

The following figure shows an example of an interactive block followed by allow.

Connection Events ([switch workflow](#))

[Connections with Application Details](#) > [Table View of Connection Events](#)

No Search Constraints ([Edit Search](#))

Jump to... ▼

		First Packet	Last Packet	Action	Reason	Initiator IP
↓	☐	2018-09-17 09:57:45	2018-09-17 09:58:21	Allow		
↓	☐	2018-09-17 09:57:43	2018-09-17 09:57:43	Interactive Block		

Logging for Allowed Connections

You can log allowed connections, which includes traffic matching the following rules and actions:

- SSL rules—**Decrypt** action
- SSL rules—**Do not decrypt** action
- SSL default action—**Do not decrypt**
- Access control rules—**Allow** action
- Access control default action—**Network Discovery Only** and any intrusion prevention option

Enabling logging for these configurations ensures the connection is logged, while also permitting (or specifying) the next phase of inspection and traffic handling. SSL logging is always end-of-connection; access control configurations also allow beginning-of-connection logging.

Although the **Analyze** action in tunnel and prefilter rules also allows connections to continue with access control, logging is disabled for rules with this action. Matching connections may still be logged by other configurations. Allowed tunnels might have their encapsulated sessions evaluated and logged individually.

When you allow traffic with an access control rule or default action, you can use an associated intrusion policy to further inspect traffic and block intrusions. For access control rules, you can also use a file policy to detect and block prohibited files, including malware. Unless you disable connection event storage, the system automatically logs most allowed connections associated with intrusion, file, and malware events. For detailed information, see [Connections That Are Always Logged, on page 2](#).

Connections with encrypted payloads are not subject to deep inspection, so connection events for encrypted connections contain limited information.

File and Malware Event Logging for Allowed Connections

When a file policy detects or blocks a file, it logs one of the following events to the Firewall Management Center database:

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

You can disable this logging on a per-access-control-rule basis. You can also disable file and malware event storage entirely.



Note We recommend you leave file and malware event logging enabled.

Beginning vs End-of-Connection Logging

You can log a connection at its beginning or its end, with the following exceptions for blocked traffic:

- **Blocked traffic**—Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events for blocked traffic. There is no unique end of connection to log.
- **Blocked encrypted traffic**—When you enable connection logging in a decryption policy, the system logs end-of-connection rather than beginning-of-connection events. This is because the system cannot determine if a connection is encrypted using the first packet in the session, and thus cannot immediately block encrypted sessions.

To optimize performance, log either the beginning or the end of any connection, but not both. Monitoring a connection for any reason forces end-of-connection logging. For a single non-blocked connection, the end-of-connection event contains all of the information in the beginning-of-connection event, as well as information gathered over the duration of the session.

The following table details the differences between beginning and end-of-connection events, including the advantages to logging each.

Table 1: Comparing Beginning and End-of-Connection Events

	Beginning-of-Connection Events	End-of-Connection Events
Can be generated...	When the system detects the beginning of a connection (or, after the first few packets if event generation depends on application or URL identification).	When the system: • Detects the close of a connection. • Does not detect the end of a connection after a period of time. • Can no longer track the session due to memory constraints.
Can be logged for...	All connections except those blocked by the decryption policy.	Most connections.

	Beginning-of-Connection Events	End-of-Connection Events
Contain...	Only information that can be determined in the first packet (or the first few packets, if event generation depends on application or URL identification).	All information in the beginning-of-connection event, plus information determined by examining traffic over the duration of the session; for example, the total amount of data transmitted or the timestamp of the last packet in the connection. Note The connection event does not count the amount of data transmitted after the threat defense returns a snort verdict for the connection or if you fastpath the connection.
Are useful...	If you want to log: - Blocked connections. - Only the beginning of a connection because the end-of-connection information does not matter to you.	If you want to: - Log encrypted connections handled by a decryption policy. - Perform any kind of detailed analysis on, or trigger correlation rules using, information collected over the duration of the session. - View connection summaries (aggregated connection data) in custom workflows, view connection data in graphical format, or create and use traffic profiles.

Limitations of Connection Logging

You cannot log:

- The outer session of a plaintext, passthrough tunnel whose encapsulated connections are inspected by access control.
- TCP connections if the three-way handshake is not completed, to avoid denial-of-service attacks against your firewalls. To monitor or debug failed connections, you can use the packet capture feature ([Packet Capture Overview](#)).

Best Practices for Connection Logging

Use the following best practices to ensure that you log *only* the connections you want to log.

So that you log only critical connections, enable connection logging on a per-access-control-rule basis.

Connections that are always logged

The system automatically logs the following:

- Some connections associated with detected files, malware, intrusions, and Intelligent Application Bypass (IAB).

For more information, see [Connections That Are Always Logged, on page 2](#).

- Monitored connections.

For more information, see [Logging for Monitored Connections, on page 4](#).

Connections to never log

Do not enable logging for the following:

- Access control rules with a Trust action.

Trusted connections are not subject to deep inspection or discovery, so connection events for trusted connections contain limited information.

- Do not enable logging for Block rules in passive deployments. To log connections that the system *would have blocked* if your devices were deployed inline, use a Monitor rule instead of a Block rule.

Only devices deployed inline (that is, using routed, switched, or transparent interfaces, or inline interface pairs) can block traffic. Because blocked connections are not actually blocked in passive deployments, the system may report multiple beginning-of-connection events for each blocked connection.

- Traffic you're not interested in. Examples follow:
 - Specific allowed traffic, such as DNS requests to a trusted DNS host.
 - Infrastructure traffic that is not related to your service offering.

(As previously mentioned, you can still monitor this traffic for threats.)

As discussed in [Connections That Are Always Logged, on page 2](#), even if you disable logging for the preceding, intrusion events, malware, and IAB are still logged.

Avoid logging what's being logged elsewhere

If another device or service is logging connection data for a network segment, disable logging for that segment's data in the Firewall Management Center. Examples follow:

- If a router logs connection events on the same network segment as the Firewall Management Center, avoid logging the same connections on the Firewall Management Center unless you need those connection events for something else, such as correlation policies or traffic profiles.
- If you use Secure Network Analytics to leverage NetFlow records reported from switches and routers to identify potential behavioral anomalies and suspicious traffic patterns, you can disable connection logging for rules monitoring those segments and instead rely on Secure Network Analytics for behavioral analytics for those parts of your network.

For more information, consult the [Secure Network Analytics documentation](#).

Log either the beginning or end of the connection (not both)

If you have a choice between beginning and end-of-connection logging, enable end-of-connection logging. This is because end-of-connection logs information from beginning-of-connection events, as well as information gathered over the duration of the session.

Log the beginning of connections *only* if you want to log blocked connections, or if end-of-connection information does not matter to you.

For more information, see [Beginning vs End-of-Connection Logging, on page 7](#).

Logging for blocked traffic

Because blocked traffic is immediately denied without further inspection, usually you can log only beginning-of-connection events.

For more information, see [Logging for Blocked Connections, on page 4](#).

Log events to an external location

If your company's security policy permits it, you can save disk space on your Firewall Management Center by streaming logs to an external source using any of the following:

- eStreamer, which enables you to stream logs from a Firewall Management Center to a custom-developed client application. For more information, see the *Secure Firewall Management Center Event Streamer Integration Guide*.
- Syslog or SNMP trap, which are referred to as *alert responses*. For more information, see [Secure Firewall Management Center Alert Responses](#).

Control what is displayed in connection events

To specify the number of rows displayed in connection events, click your username in the upper right of the Firewall Management Center, click **User Preferences**, and then click the **Event View Settings** tab. The maximum you can set is 1000 events per page.

Requirements and Prerequisites for Connection Logging

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Configure Connection Logging

The following sections describe how to set up connection logging to match various rules and conditions.

Logging Connections with Tunnel and Prefilter Rules

The prefilter policy applies to Secure Firewall Threat Defense devices only.

Before you begin

- Set the rule action to **Block** or **Fastpath**. Logging is disabled for the **Analyze** action, which allows connections to continue with access control, where other configurations determine their handling and logging.
- Logging is performed on inner flows, not on the encapsulating flow.

Procedure

-
- Step 1** In the prefilter policy editor, click **Edit** (✎) next to the rule where you want to configure logging. If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Step 2** Click **Logging**.
- Step 3** Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.
To optimize performance, log either the beginning or the end of any connection, but not both. Because blocked traffic is immediately denied without further inspection, you cannot log end-of-connection events for Block rules.
- Step 4** Specify where to send connection events:
- Step 5** Click **Save** to save the rule.
- Step 6** Click **Save** to save the prefilter policy.
-

What to do next

- Deploy configuration changes.

Logging Decryptable Connections with TLS/SSLDecryption Rules

Procedure

-
- Step 1** In the decryption policy editor, click **Edit** (✎) next to the rule where you want to configure logging. If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Step 2** Click **Logging**.
- Step 3** Check the **Log at End of Connection** check box.

For monitored traffic, end-of-connection logging is required.

Step 4 Specify where to send connection events.

Send events to the event viewer if you want to perform Firewall Management Center-based analysis on these connection events. For monitored traffic, this is required.

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

Step 6 Click **Save** to save the decryption policy.

What to do next

- Deploy configuration changes.

Logging Connections with Security Intelligence

The Security Intelligence policy requires the Threat Smart License or Protection Classic License.

Procedure

Step 1 In the access control policy editor, click **Security Intelligence**.

Step 2 Click the **Logging** () icon to enable Security Intelligence logging using the following criteria:

- By IP address—Click the logging icon next to **Networks**.
- By URL—Click the logging icon next to **URLs**.
- By Domain Name—Click the logging icon next to the **DNS Policy** drop-down list.

If the logging icon is disabled, settings are inherited from an ancestor policy, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.

Step 3 Check the **Log Connections** check box.

Step 4 Specify where to send connection and Security-Related connection events.

Send events to the Security Cloud Control if you want to perform Security Cloud Control-based analysis, or if you set a Block list to monitor-only.

Step 5 Click **OK** to set logging options.

Step 6 Click **Save** to save the policy.

What to do next

- Deploy configuration changes.

Logging Connections with Access Control Rules

Depending on your choices for the rule action and deep inspection options, your logging options differ; see [How Rules and Policy Actions Affect Logging, on page 3](#).

Procedure

- Step 1** In the access control policy editor, click **Edit** (✎) next to the rule where you want to configure logging.
- If **View** (👁) appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Step 2** Click **Logging**.
- Step 3** Specify whether you want to **Log at Beginning of Connection** or **Log at End of Connection**.
- To optimize performance, log either the beginning or the end of any connection, but not both.
- Step 4** (Optional) Check the **Log Files** check box to log file and malware events associated with the connection.
- It is recommended to leave this option enabled.
- Step 5** Specify where to send the connection events:
- **Event Viewers:** Send events to the Firewall Management Center. When using cloud management, send events to the cloud-delivered Firewall Management Center and to an on-premises Firewall Management Center if you have configured it to perform event analytics only. You can view the events in the event viewer of either product.
 - **Syslog Server:** Send connection events to the syslog server configured in the Logging tab in Access Control Policy, unless overridden.
- Show Overrides:** Displays the options to override the settings configured in the access control policy.
- **Override Severity:** When you choose this option and select a severity for the rule, connection events for this rule will have the selected severity regardless of the severity configured in the Logging tab in Access Control Policy.
 - **Override Default Syslog Destination:** Send the syslog generated for the connection event for this rule to destination specified in this alert.
- **SNMP Trap:** Connection events are sent to the selected SNMP trap.
- Step 6** Click **Confirm**.
- Step 7** Click **Apply** to save the rule.
-

What to do next

- Deploy configuration changes.

Logging Connections with a Policy Default Action

A policy's default action determines how the system handles traffic that matches none of the rules in the policy (except Monitor rules in access control and decryption policies, which match and log—but do not handle or inspect—traffic).

Logging settings for the decryption policies default action also govern how the system logs undecryptable sessions.

Before you begin

- For prefilter default action logging, set the default action to **Block all tunnel traffic**. Logging is disabled for the **Allow all tunnel traffic** action, which allows connections to continue with access control, where other configurations determine their handling and logging.

Procedure

-
- Step 1** In the policy editor, click the **Default Logging and Inspection**  next to the **Default Action** drop-down list.
- Step 2** Specify when you want to log matching connections:
- Log at Beginning of Connection—Not supported for SSL default actions.
 - Log at End of Connection—Not supported if you choose the access control **Block All Traffic** default action or the prefilter **Block all tunnel traffic** default action.
- To optimize performance, log either the beginning or the end of any connection, but not both.
- If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration. In an access control policy, the configuration may also be inherited from an ancestor policy.
- Step 3** Specify where to send connection events.
- Send events to the event viewer if you want to perform Firewall Management Center-based analysis on these connection events.
- Step 4** Click **Apply**.
- Step 5** Click **Save** to save the policy.
-

What to do next

- Deploy configuration changes.

Limiting Logging of Long URLs

End-of-connection events for HTTP traffic record the URL requested by monitored hosts. Disabling or limiting the number of stored URL characters may improve system performance. Disabling URL logging (storing zero characters) does not affect URL filtering. The system filters traffic based on requested URLs even though the system does not record them.

Procedure

-
- Step 1** In the access control policy editor, click **More > Advanced Settings**, then click **Edit** (✎) next to **General Settings**.
- If **View** (👁) appears instead, the configuration is inherited from an ancestor policy, belongs to an ancestor domain, or you do not have permission to modify the configuration. If the configuration is unlocked, uncheck **Inherit from base policy** to enable editing.
- Step 2** Enter the **Maximum URL characters to store in connection events**.
- Step 3** Click **OK**.
- Step 4** Click **Save** to save the policy.
-

What to do next

- Deploy configuration changes.

