

Monitoring the Device

The system includes dashboards and an Event Viewer that you can use to monitor the device and traffic that is passing through the device.

- Enable Logging to Obtain Traffic Statistics, on page 1
- Monitoring Traffic and System Dashboards, on page 4
- Monitoring Additional Statistics Using the Command Line, on page 7
- Viewing Events, on page 7

Enable Logging to Obtain Traffic Statistics

You can monitor a wide range of traffic statistics using the monitoring dashboards and the Event Viewer. However, you must enable logging to tell the system which statistics to collect. Logging generates various types of events that provide insight into the connections going through the system.

The following topics explain more about events and the information they provide, with special emphasis on connection logging.

Event Types

The system can generate the following types of events. You must generate these events to see related statistics in the monitoring dashboards.

Connection Events

You can generate events for connections as users generate traffic that passes through the system. Enable connection logging on access rules to generate these events. You can also enable logging on Security Intelligence policies and SSL decryption rules to generate connection events.

Connection events include a wide variety of information about a connection, including source and destination IP addresses and ports, URLs and applications used, and the number of bytes or packets transmitted. The information also includes the action taken (for example, allowing or blocking the connection), and the policies applied to the connection.

Intrusion Events

The system examines the packets that traverse your network for malicious activity that could affect the availability, integrity, and confidentiality of a host and its data. When the system identifies a possible intrusion, it generates an intrusion event, which is a record of the date, time, type of exploit, and contextual

information about the source of the attack and its target. Intrusion events are generated for any intrusion rule set to block or alert, regardless of the logging configuration of the invoking access control rule.

File Events

File events represent files that the system detected, and optionally blocked, in network traffic based on your file policies. You must enable file logging on the access rule that applies the file policy to generate these events.

When the system generates a file event, the system also logs the end of the associated connection regardless of the logging configuration of the invoking access control rule.

Malware Events

The system can detect malware in network traffic as part of your overall access control configuration. The Malware Defense can generate a malware event, containing the disposition of the resulting event, and contextual data about how, where, and when the malware was detected. You must enable file logging on the access rule that applies the file policy to generate these events.

The disposition of a file can change, for example, from clean to malware or from malware to clean. If Malware Defense queries the Secure Malware Analytics Cloud about a file, and the cloud determines the disposition has changed within a week of the query, the system generates retrospective malware events.

Security Intelligence Events

Security Intelligence events are a type of connection event generated by the Security Intelligence policy for each connection blocked or monitored by the policy. All Security Intelligence events have a populated Security Intelligence Category field.

For each of these events, there is a corresponding “regular” connection event. Because the Security Intelligence policy is evaluated before many other security policies, including access control, when a connection is blocked by Security Intelligence, the resulting event does not contain the information that the system would have gathered from subsequent evaluation, for example, user identity.

Configurable Connection Logging

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

Because the system can log a connection for multiple reasons, disabling logging in one place does not ensure that matching connections will not be logged.

You can configure connection logging in the following places.

- Access control rules and default action—Logging at the end of a connection provides the most information about the connection. You can also log the beginning of the connection, but these events have incomplete information. Connection logging is disabled by default, so you must enable it for each rule (and the default action) that targets traffic that you want to track.
- Security Intelligence policy—You can enable logging to generate Security Intelligence connection events for each blocked connection. 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.

- SSL Decryption rules and default action—You can configure logging at the end of a connection. 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.

Automatic Connection Logging

The system automatically saves the following end-of-connection events, regardless of any other logging configurations.

- The system automatically logs connections associated with intrusion events, unless the connection is handled by the access control policy's default action. You must enable logging on the default action to get intrusion events for matching traffic.
- The system automatically logs connections associated with file and malware events. This is for connection events only: you can optionally disable the generation of file and malware events.

Tips for Connection Logging

Keep the following tips in mind when considering your logging configuration and the evaluation of related statistics:

- When you allow traffic with an access control rule, you can use an associated intrusion or file policy (or both) to further inspect traffic and block intrusions, prohibited files, and malware before the traffic can reach its final destination. Note, however, that by default file and intrusion inspection is disabled for encrypted payloads. If the intrusion or file policies find reason to block a connection, the system immediately logs an end-of-connection event regardless of your connection log settings. Logging allowed connections provides the most statistical information on the traffic in your network.
- A trusted connection is one that is handled by a Trust access control rule or the default action in an access control policy. However, trusted connections are not inspected for discovery data, intrusions, or prohibited files and malware. Therefore, connection events for trusted connections contain limited information.
- For access control rules and access control policy default actions that block traffic, the system logs beginning-of-connection events. Matching traffic is denied without further inspection.
- 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.
- If you select the **Bypass Access Control policy for decrypted traffic (sysopt permit-vpn)** option when you configure remote access VPN connection profiles, or you otherwise enable the **sysopt connection permit-vpn** command, all site-to-site or remote access VPN traffic bypasses inspection and the access control policy. Thus, you will get no connection events for this traffic, and the traffic will not be reflected in any statistical dashboards.

Sending Events to an External Syslog Server

Besides viewing events through the Firewall Device Manager, which has a limited capacity to store events, you can selectively configure rules and policies to send events to an external syslog server. You can then use the features and additional storage of your selected syslog server platform to view and analyze event data.

To send events to an external syslog server, edit each rule, default action, or policy that enables connection logging and select a syslog server object in the log settings. To send intrusion events to a syslog server, configure the server in the intrusion policy settings. To send file/malware events to a syslog server, configure the server on **Device** > **System Settings** > **Logging Settings**.

For more information, see the help for each rule and policy type and also see Configuring Syslog Servers.

Evaluating Events Using Cisco Cloud-Based Services

In addition to using Event Viewer and your own syslog servers, you can send connection events, and high-priority intrusion, file, and malware events, to a Cisco cloud-based server. Cisco cloud-based services, such as Threat Response, can pull the events from that cloud server and you can then use those services to evaluate these events.

These cloud-based services are separate from the Firewall Threat Defense device and Firewall Device Manager. If you elect to use a service that requires you to send these events to the Cisco cloud, you must enable the connection on the **Device** > **System Settings** > **Cloud Services** page. See Sending Events to the Cisco Cloud.

Monitoring Traffic and System Dashboards

The system includes several dashboards that you can use to analyze the traffic going through the device and the results of your security policy. Use the information to evaluate the overall efficacy of your configuration and to identify and resolve network problems.

The dashboards for units in a high availability group show statistics for that device only. Statistics are not synchronized among the units.



Note The data used in traffic-related dashboards is collected from access control rules that enable connection or file logging, and other security policies that allow logging. The dashboards do not reflect traffic that matches rules for which no logging is enabled. Ensure that you configure your rules to log the information that matters to you. In addition, user information is available only if you configure identity rules to collect user identity. And finally, intrusion, file, malware, and URL category information is available only if you have a license for those features and configure rules that use the features.

Procedure

Step 1 Click **Monitoring** in the main menu to open the Dashboards page.

You can select predefined time ranges, such as the last hour or week, or define a custom time range with specific start and end times, to control the data shown in the dashboard graphs and tables.

Traffic-related dashboards include the following types of display:

- Top 5 bar graphs—These are shown in the **Network Overview** dashboard, and in the per-item summary dashboards you see if you click on an item in a dashboard table. You can toggle the information between a count of **Transactions** or **Data Usage** (total bytes sent and received). You can also toggle the display to show all transactions, allowed transactions, or denied transactions. Click the **View More** link to see the table associated with the graph.
- Tables—Tables show items of a particular type (for example, applications or URL categories) with that item's total transactions, allowed transactions, blocked transactions, data usage, and bytes sent and received. You can toggle the numbers between raw **Values** and **Percentages**, and show the top 10, 100, or 1000 entries. If the item is a link, click it to see a summary dashboard with more detailed information.

Step 2 Click the **Dashboard** links in the table of contents to see dashboards for the following data:

- **Network Overview**—Shows summary information about the traffic in the network, including the access rules (policies) matched, users initiating traffic, applications used in connections, intrusion threats (signatures) matched, URL categories for URLs accessed, and the most frequent destinations for connections.
- **Users**—Shows the top users of your network. You must configure identity policies to see user information. If there is no user identity, the source IP address is included. You might see the following special entities:
  - **Failed Authentication**—The user was prompted to authenticate, but failed to enter a valid username/password pair within the maximum number of allowed attempts. Failure to authenticate does not itself prevent the user from accessing the network, but you can write an access rule to limit network access for these users.
  - **Guest**—Guest users are like Failed Authentication users, except that your identity rule is configured to call these users Guest. Guest users were prompted to authenticate and failed to do so within the maximum number of attempts.
  - **No Authentication Required**—The user was not prompted to authentication, because the user's connections matched identity rules that specified no authentication.
  - **Unknown**—There is no user mapping for the IP address, and there is no record of failed authentication yet. Typically, this means that no HTTP traffic has yet been seen from that address.
- **Applications**—Shows the top applications, such as HTTP, that are being used in the network. The information is available only for connections that are inspected. Connections are inspected if they match an “allow” rule, or a block rule that uses criteria other than zone, address, and port. Thus, application information is not available if the connection is trusted or blocked prior to hitting any rule that requires inspection.
- **Web Applications**—Shows the top web applications, such as Google, that are being used in the network. The conditions for collecting web application information are the same as those for the Application dashboard.
- **URL Categories**—Shows the top categories of web sites, such as Gambling or Educational Institutions, that are being used in the network based on the categorization of web sites visited. You must have at least one access control rule that uses URL category as a traffic matching criteria to get this information. The information will be available for traffic that matches the rule, or for traffic that has to be inspected to determine if it matches the rule. You will not see category (or reputation) information for connections that match rules that come before the first web-category access control rule.
- **Access And SI Rules**—Shows the top access rules and Security Intelligence rule-equivalents matched by network traffic.

- **Zones**—Shows the top security zone pairs for traffic entering and then exiting the device.
- **Destinations**—Shows the top destinations for network traffic.
- **Attackers**—Shows the top attackers, which are the source of connections that trigger intrusion events. You must configure intrusion policies on access rules to see this information.
- **Targets**—Shows the top targets of intrusion events, which are the victims of an attack. You must configure intrusion policies on access rules to see this information.
- **Threats**—Shows the top intrusion rules that have been triggered. You must configure intrusion policies on access rules to see this information.
- **File Logs**—Shows the top file types seen in network traffic. You must configure file policies on access rules to see this information.
- **Malware**—Shows the top Malware action and disposition combinations. You can drill down to see information on the associated file types. You must configure file policies on access rules to see this information.
  - Possible actions are: Malware Cloud Lookup, Block, Archive Block (Encrypted), Detect, Custom Detection, Cloud Lookup Timeout, Malware Block, Archive Block (Depth Exceeded), Custom Detection Block, TID block, Archive Block (Failed to Inspect).
  - Possible dispositions are: Malware, Unknown, Clean, Custom Detection, Unavailable.
- **SSL Decryption**—Shows the breakdown of encrypted vs. plain text traffic through the device, plus the breakdown of how encrypted traffic was decrypted according to SSL decryption rules.
- **System**— Shows an overall system view, including a display of interfaces and their status (mouse over an interface to see its IP addresses), overall average system throughput (in 5 minute buckets for up to one hour, and one hour buckets for longer periods), and summary information on system events, CPU usage, memory usage, and disk usage. You can restrict the throughput graph to show a specific interface rather than all interfaces.

  Note
  The information shown on the System dashboard is at the overall system level. If you log into the device CLI, you can use various commands to see more detailed information. For example, the **show cpu** and **show memory** commands include parameters for showing other details, whereas these dashboards show data from the **show cpu system** and **show memory system** commands.

Step 3 You can also click these links in the table of contents:

- **Events**—To view events as they occur. You must enable connection logging in individual access rules to see connection events related to those rules. Also, enable logging in the Security Intelligence policy and SSL decryption rules to see Security Intelligence events and additional connection event data. These events can help you resolve connection problems for your users.
- **Sessions**—To view and manage the Firewall Device Manager user sessions. For more information, see Managing the Firewall Device Manager User Sessions.

Monitoring Additional Statistics Using the Command Line

The Firewall Device Manager dashboards provide a wide variety of statistics related to the traffic going through the device and general system usage. However, you can get additional information on areas not covered by the dashboards using the CLI Console or by logging into the device CLI (see Logging Into the Command Line Interface (CLI)).

The CLI includes a variety of **show** commands to provide these statistics. You can also use the CLI for general troubleshooting, including commands such as **ping** and **traceroute**. Most **show** commands have companion **clear** commands to reset statistics to 0. (You cannot clear statistics from the CLI Console.)

You can find documentation for the commands in *Cisco Secure Firewall Threat Defense Command Reference*, http://www.cisco.com/c/en/us/td/docs/security/firepower/command_ref/b_Command_Reference_for_Firepower_Threat_Defense.html.

For example, you might find the following commands generally useful.

- **show nat** displays hit counts for your NAT rules.
- **show xlate** displays the actual NAT translations that are active.
- **show conn** provides information about current connections going through the device.
- **show dhcpd** provides information about the DHCP servers you configure on the interfaces.
- **show interface** provides usage statistics for each interface.

Viewing Events

You can view events that are generated from your security policies that enable logging. Events are also generated for intrusion and file policies that are triggered.

The event viewer table shows the events generated in real time. As new events are generated, older events are rolled out of the table.

Before you begin

Whether events of particular types are generated depends on the following in addition to connections that match the related policy:

- Connection events—An access rule must enable connection logging. You can also enable connection logging in the Security Intelligence policy and in SSL decryption rules.
- Intrusion events—An access rule must apply an intrusion policy.
- File and Malware events—An access rule must apply a file policy and enable file logging.
- Security Intelligence events—You must enable and configure the Security Intelligence policy, and enable logging.

Procedure

Step 1 Click **Monitoring** in the main menu.

Step 2 Select **Events** from the table of contents.

The event viewer organizes events on tabs based on event types. For more information, see Event Types, on page 1.

Step 3 Click the tab that shows the type of event you want to view.

You can do the following with the event list:

- Click **Pause** to stop the addition of new events so that you can more easily find and analyze an event. Click **Resume** to allow new events to appear.
- Select a different refresh rate (5, 10, 20, or 60 seconds) to control how fast new events are shown.
- Create a custom view that includes the columns you want. To create a custom view, either click the **+** button in the tab bar, or click **Add/Remove Columns**. You cannot change the pre-set tabs, so adding or removing columns creates a new view. For more information, see Configuring Custom Views, on page 8.
- To change the width of a column, click and drag the column heading divider to the desired width.
- Mouse over an event and click **View Details** to see complete information on an event. For a description of the various fields in an event, see Event Field Descriptions, on page 10.

Step 4 If necessary, apply a filter to the table to help you locate the desired events based on various event attributes.

To create a new filter, either manually type in the filter by selecting atomic elements from the drop-down list and entering the filter value, or build a filter by clicking a cell in the events table that includes a value on which you want to filter. You can click multiple cells in the same column to create an OR condition among the values, or click cells in different columns to create an AND condition among the columns. If you build the filter by clicking cells, you can also edit the resulting filter to fine-tune it. For detailed information about creating filter rules, see Filtering Events, on page 9.

Once you build the filter, do any of the following:

- To apply the filter and update the table to show only those events that match the filter, click the **Filter** button.
- To clear an entire filter that you have applied and return the table to a non-filtered state, click **Reset Filters** in the **Filter** box.
- To clear one of the atomic elements of a filter, mouse over the element and click the **X** for the element. Then, click the **Filter** button.

Configuring Custom Views

You can create your own custom views so that you can easily see the columns you want when viewing events. You can also edit or delete custom views, although you cannot edit or delete the pre-defined views.

Procedure

Step 1 Select **Monitoring** > **Events**.

Step 2 Do one of the following:

- To create a new view based on an existing custom (or pre-defined) view, click the tab for the view, then click the + button to the left of the tabs.
- To edit an existing custom view, click the tab for the view.

Note
To delete a custom view, simply click the **X** button in the view's tab. You cannot undo a delete.

Step 3 Click the **Add/Remove Columns** link above the events table on the right, and select or deselect columns until the selected list includes only those columns to include in the view.

Click and drag columns between the available (but not used) and selected lists. You can also click and drag columns in the selected list to change the left-to-right order of the columns in the table. For a description of the columns, see Event Field Descriptions, on page 10.

When finished, click **OK** to save your column changes.

Note
If you change column selection while viewing a pre-defined view, a new view is created.

Step 4 If necessary, change column widths by clicking and dragging the column separators.

Filtering Events

You can create complex filters to limit the events table to the events that currently interest you. You can use the following techniques, alone or in combination, to build a filter:

Clicking columns

The easiest way to build a filter is to click on cells in the events table that contain the values on which you intend to filter. Clicking a cell updates the **Filter** field with a correctly-formulated rule for that value and field combination. However, using this technique requires that the existing list of events contains the desired values.

You cannot filter on all columns. If you can filter on the contents of a cell, it is underlined when you mouse over it.

Selecting atomic elements

You can also build a filter by clicking in the **Filter** field and selecting the desired atomic element from the drop-down list, then typing in the match value. These elements include event fields that are not shown as columns in the events table. They also include operators to define the relationship between the value you type in and the events to display. Whereas clicking columns always results in an "equals (=)" filter, when you select an element, you can also select "greater than (>)" or "less than (<)" for numeric fields.

Regardless of how you add an element to the **Filter** field, you can type into the field to adjust the operator or value. Click **Filter** to apply the filter to the table.

Operators for Event Filters

You can use the following operators in an event filter:

=	Equals. The event matches the specified value. You cannot use wildcards.
!=	Not equals. The event does not match the specified value. You must type in the ! (exclamation point) to build a not-equals expression.
>	Greater than. The event contains a value that is greater than the specified value. This operator is available for numeric values only, such as port and IP address.
<	Less than. The event contains a value that is less than the specified value. This operator is available for numeric values only.

Rules for Complex Event Filters

When building a complex filter that contains more than one atomic element, keep the following rules in mind:

- Elements of the same type have an OR relationship between all values for that type. For example, including Initiator IP=10.100.10.10 and Initiator IP=10.100.10.11 matches events that have either of these addresses as the traffic source.
- Elements of different types have an AND relationship. For example, including Initiator IP=10.100.10.10 and Destination Port/ICMP Type=80 matches events that have this source address AND destination port only. Events from 10.100.10.10 to a different destination port are not shown.
- Numeric elements, including IPv4 and IPv6 addresses, can specify ranges. For example, you could specify Destination Port=50-80 to capture all traffic for ports within this range. Use a hyphen to separate the start and end numbers. Ranges are not allowed for all numeric fields, for example, you cannot specify an IP address range in the Source element.
- You cannot use wildcards or regular expressions.

Event Field Descriptions

Events can contain the following information. You can see this information when you view event details. You can also add columns to the Event Viewer table to show the information that most interests you.

Following is a complete list of the available fields. Not every field applies to every type of event. Keep in mind that the information available for any individual event can vary depending on how, why, and when the system logged the connection.

Action

For connection or security intelligence events, the action associated with the access control rule or default action that logged the connection:

Allow

Explicitly allowed connections.

Trust

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

Block

Blocked connections. The **Block** action can be associated with Allow access rules under the following conditions:

- Connections where an exploit was blocked by an intrusion policy.
- Connections where a file was blocked by a file policy.
- Connections blocked by Security Intelligence.
- Connections blocked by an SSL policy.

Default Action

The connection was handled by the default action.

For file or malware events, the file rule action associated with the rule action for the rule the file matched, and any associated file rule action options.

Allowed Connection

Whether the system allowed the traffic flow for the event.

Application

The application detected in the connection.

Application Business Relevance

The business relevance associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Application Categories, Application Tag

Criteria that characterize the application to help you understand the application's function.

Application Risk

The risk associated with the application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of application detected in the connection has an associated risk; this field displays the highest of those.

Block Type

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Client Application, Client Version

The client application and version of that client detected in the connection.

Client Business Relevance

The business relevance associated with the client traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of client detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Client Category, Client Tag

Criteria that characterize the application to help you understand the application's function.

Client Risk

The risk associated with the client traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of client detected in the connection has an associated risk; this field displays the highest of those.

Connection

The unique ID for the traffic flow, internally generated.

Connection Blocktype Indicator

The type of block specified in the access control rule matching the traffic flow in the event: block or interactive block.

Connection Bytes

The total bytes for the connection.

Connection Time

The time for the beginning of the connection.

Connection Timestamp

The time the connection was detected.

Denied Connection

Whether the system denied the traffic flow for the event.

Destination Country and Continent

The country and continent of the receiving host.

Destination IP

The IP address used by the receiving host in an intrusion, file, or malware event.

Destination Port/ICMP Code; Destination Port; Destination Icode

The port or ICMP code used by the session responder.

Destination Security Group Tag, Destination Security Group Tag Name

The TrustSec security group tag number and name associated with the destination, if any.

Direction

The direction of transmission for a file.

Disposition

The file's disposition:

Malware

Indicates that the Secure Malware Analytics Cloud categorized the file as malware or the file’s threat score exceeded the malware threshold defined in the file policy. Local malware analysis can also mark files as malware.

Clean

Indicates that the Secure Malware Analytics Cloud categorized the file as clean, or that a user added the file to the clean list.

Unknown

Indicates that the system queried the Secure Malware Analytics Cloud, but the file has not been assigned a disposition; in other words, the Secure Malware Analytics Cloud has not categorized the file.

Custom Detection

Indicates that a user added the file to the custom detection list.

Unavailable

Indicates that the system could not query the Secure Malware Analytics Cloud. You may see a small percentage of events with this disposition; this is expected behavior.

N/A

Indicates that a Detect Files or Block Files rule handled the file and the system did not query the Secure Malware Analytics Cloud.

Egress Interface, Egress Security Zone

The interface and zone through which the connection exited the device.

Egress Virtual Router

The name of the virtual router, if any, to which the destination interface belongs.

Event, Event Type

The type of event.

Event Seconds, Event Microseconds

The time, in seconds or microseconds, when the event was detected.

File Category

The general categories of file type, for example: Office Documents, Archive, Multimedia, Executables, PDF files, Encoded, Graphics, or System Files.

File Event Timestamp

The time and date the file or malware file was created.

File Name

The name of the file.

File Rule Action

The action associated with file policy rule that detected the file, and any associated file rule action options.

File SHA-256

The SHA-256 hash value of the file.

File Size (KB)

The size of the file, in kilobytes. File size can be blank in cases where the system blocked the file before it was completely received.

File Type

The type of file, for example, HTML or MSEXE.

File/Malware Policy

The file policy associated with the generation of the event.

Filelog Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Firewall Policy Rule, Firewall Rule

The access control rule or default action that handled the connection.

First Packet

The date and time the first packet of the session was seen.

HTTP Referrer

The HTTP referrer, which represents the referrer of a requested URL for HTTP traffic detected in the connection (such as a website that provided a link to, or imported a link from, another URL).

HTTP Response

The HTTP status code sent in response to a client's HTTP request over a connection.

IDS Classification

The classification where the rule that generated the event belongs.

Ingress Interface, Ingress Security Zone

The interface and zone through which the connection entered the device.

Ingress Virtual Router

The name of the virtual router, if any, to which the source interface belongs.

Initiator Bytes, Initiator Packets

The total number of bytes or packets transmitted by the session initiator.

Initiator Country and Continent

The country and continent of the host that initiated the session. Available only if the initiator IP address is routable.

Initiator IP

The host IP address (and hostname, if DNS resolution is enabled) that initiated the session in a connection or Security Intelligence event.

Inline Result

Whether the system dropped or would have dropped the packet that triggered an intrusion event if operating in inline mode. Blank indicates that the triggered rule was not set to Drop and Generate Events

Intrusion Policy

The intrusion policy where the rule that generated the event was enabled.

IPS Blocktype Indicator

The action of the intrusion rule matching the traffic flow in the event.

Last Packet

The date and time the last packet of the session was seen.

MPLS Label

The Multiprotocol Label Switching label associated with the packet that triggered this intrusion event.

Malware Blocktype Indicator

The type of block specified in the file rule matching the traffic flow in the event: block or interactive block.

Message

For intrusion events, the explanatory text for the event. For malware or file events, any additional information associated with the malware event.

NAT Destination IP

For packets subject to Network Address Translation (NAT), the translated destination IP address.

NAT Destination Port

For packets subject to Network Address Translation (NAT), the translated destination port.

NAT Source IP

For packets subject to Network Address Translation (NAT), the translated source IP address.

NAT Source Port

For packets subject to Network Address Translation (NAT), the translated source port.

NetBIOS Domain

The NetBIOS domain used in the session.

Original Client Country and Continent

The country and continent of the original client host that initiated the session. Available only if the original client IP address is routable.

Original Client IP

The original IP address of the client that initiated an HTTP connection. This address is derived from the X-Forwarded-For (XFF) or True-Client-IP HTTP header fields or their equivalent.

Policy, Policy Revision

The access control policy, and its revision, that includes the access (firewall) rule associated with the event.

Priority

The event priority as determined by the Cisco Talos Intelligence Group (Talos): high, medium, or low.

Protocol

The transport protocol used in the connection.

Reason

The reason or reasons the connection was logged, in the situations explained in the following table. This field is otherwise empty.

Reason	Description
DNS Block	The system denied the connection without inspection, based on the domain name and Security Intelligence data. A reason of DNS Block is paired with an action of Block, Domain not found, or Sinkhole, depending on the DNS rule action.
DNS Monitor	The system would have denied the connection based on the domain name and Security Intelligence data, but you configured the system to monitor, rather than deny, the connection.
Elephant Flow	The connection is large enough to be considered an elephant flow, which is a flow that can be large enough to affect overall system performance. By default, elephant flows are those larger than 1GB/10 seconds. You can adjust the byte and time thresholds for identifying elephant flows in the device CLI using the **system support elephant-flow-detection** command.
File Block	The connection contained a file or malware file that the system prevented from being transmitted. A reason of File Block is always paired with an action of Block.
File Custom Detection	The connection contained a file on the custom detection list that the system prevented from being transmitted.
File Monitor	The system detected a particular type of file in the connection.
File Resume Allow	File transmission was originally blocked by a Block Files or Block Malware file rule. After a new access control policy allowing the file was deployed, the HTTP session automatically resumed.
File Resume Block	File transmission was originally allowed by a Detect Files or Malware Cloud Lookup file rule. After a new access control policy blocking the file was deployed, the HTTP session automatically stopped.
Intrusion Block	The system blocked or would have blocked an exploit (intrusion policy violation) detected in the connection. A reason of Intrusion Block is paired with an action of Block for blocked exploits and Allow for would-have-blocked exploits.
Intrusion Monitor	The system detected, but did not block, an exploit detected in the connection. This occurs when the state of the triggered intrusion rule is set to Generate Events.
IP Block	The system denied the connection without inspection, based on the IP address and Security Intelligence data. A reason of IP Block is always paired with an action of Block.
Pending Rule Match	This occurs if the connection ends before the system can determine which rule matches. These short-lived connections (such as DNS lookups) are typically allowed through the firewall, especially if you have an application-based rule early in the access control policy, which makes the system inspect a few packets to determine rule match.

Reason	Description
SSL Block	The system blocked an encrypted connection based on the SSL inspection configuration. A reason of SSL Block is always paired with an action of Block.
URL Block	The system denied the connection without inspection, based on the URL and Security Intelligence data. A reason of URL Block is always paired with an action of Block.

Receive Times

The date and time the event was generated.

Referenced Host

If the protocol in the connection is HTTP or HTTPS, this field displays the hostname that the respective protocol was using.

Responder Bytes, Responder Packets

The total number of bytes or packets transmitted by the session responder.

Responder Country and Continent

The country and continent of the host that responded to the session. Available only if the responder IP address is routable.

Responder IP

The host IP address (and hostname, if DNS resolution is enabled) of the session responder in a connection or Security Intelligence event.

SI Category ID (Security Intelligence Category)

The name of the object that contained the blocked item, such as a network or URL object name, or the name of a feed category.

Signature

The signature ID for a file/malware event.

Source Country and Continent

The country and continent of the sending host. Available only if the source IP address is routable.

Source IP

The IP address used by the sending host in an intrusion, file, or malware event.

Source Port/ICMP Type; Source Port; Source Port Itype

The port or ICMP type used by the session initiator.

Source Security Group Tag, Source Security Group Tag Name

The TrustSec security group tag number and name associated with the source, if any.

SSL Actual Action

The actual action that the system applied to the connection. This can differ from the expected action. For example, a connection might match a rule that applies decryption, but could not be decrypted for some reason.

Action	Description
Block/Block with reset	Represents blocked encrypted connections.
Decrypt (Resign)	Represents an outgoing connection decrypted using a re-signed server certificate.
Decrypt (Replace Key)	Represents an outgoing connection decrypted using a self-signed server certificate with a substituted public key.
Decrypt (Known Key)	Represents an incoming connection decrypted using a known private key.
Default Action	Indicates the connection was handled by the default action.
Do not Decrypt	Represents a connection the system did not decrypt.

SSL Certificate Fingerprint

The SHA hash value used to authenticate the certificate.

SSL Certificate Status

This applies only if you configured a Certificate Status SSL rule condition. If encrypted traffic matches an SSL rule, this field displays one or more of the following server certificate status values:

- Self Signed
- Valid
- Invalid Signature
- Invalid Issuer
- Expired
- Unknown
- Not Valid Yet
- Revoked

If undecryptable traffic matches an SSL rule, this field displays Not Checked.

SSL Cipher Suite

The cipher suite used in the connection.

SSL Expected Action

The action specified in the SSL rule the connection matched.

SSL Flow Flags

The first ten debugging level flags for an encrypted connection.

SSL Flow Messages

The SSL/TLS messages exchanged between client and server during the SSL handshake, such as HELLO_REQUEST and CLIENT_HELLO. See http://tools.ietf.org/html/rfc5246 for more information about the messages exchanged in TLS connections.

SSL Policy

The name of the SSL Decryption policy applied to the connection.

SSL Rule

The name of the SSL Decryption rule applied to the connection.

SSL Session ID

The hexadecimal Session ID negotiated between the client and server during the SSL handshake.

SSL Ticket ID

A hexadecimal hash value of the session ticket information sent during the SSL handshake.

SSL URL Category

The URL category of the destination web server as determined during SSL decryption processing.

SSL Version

The SSL/TLS version used in the connection.

TCP Flags

The TCP flags detected in the connection.

Total Packets

The total number of packets transmitted in the connection, which is **Initiator Packets** + **Responder Packets**.

URL, URL Category, URL Reputation, URL Reputation Score

The URL requested by the monitored host during the session and its associated category, reputation, and reputation score, if available.

For DNS lookup request filtering, the category and reputation are for the FQDN shown in the DNS Query field. The URL field will be blank, because the category/reputation lookup is being done for a DNS request rather than a web request.

If the system identifies or blocks an SSL application, the requested URL is in encrypted traffic, so the system identifies the traffic based on an SSL certificate. For SSL applications, therefore, the URL indicates the common name contained in the certificate.

User

The user associated with the initiator IP address.

VLAN

The innermost VLAN ID associated with the packet that triggered the event.

Web App Business Relevance

The business relevance associated with the web application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of web application detected in the connection has an associated business relevance; this field displays the lowest (least relevant) of those.

Web App Categories, Web App Tag

Criteria that characterize the web application to help you understand the web application's function.

Web App Risk

The risk associated with the web application traffic detected in the connection: Very High, High, Medium, Low, or Very Low. Each type of web application detected in the connection has an associated risk; this field displays the highest of those.

Web Application

The web application, which represents the content or requested URL for HTTP traffic detected in the connection.

If the web application does not match the URL for the event, the traffic is probably referred traffic, such as advertisement traffic. If the system detects referred traffic, it stores the referring application (if available) and lists that application as the web application.