

Intrusion Events

The following topics describe how to work with intrusion events.

- Intrusion events, on page 1
- Tools for reviewing and evaluating intrusion events, on page 2
- License requirements for intrusion events, on page 2
- Requirements and prerequisites for intrusion events, on page 3
- View intrusion events, on page 3
- Intrusion event workflow pages, on page 21
- View intrusion event statistics, on page 40
- View intrusion event performance graphs, on page 42
- View intrusion event graphs, on page 46
- History for intrusion events, on page 47

Intrusion events

An intrusion event is a security record that

- documents the date, time, and type of exploit when the system identifies a possible intrusion,
- includes contextual information about the source of the attack and its target,
- records a copy of the packet or packets that triggered the event for packet-based events, and
- is transmitted from managed devices to the Secure Firewall Management Center for aggregated analysis.

Network monitoring capabilities

The system can help you monitor your network for traffic that could affect the availability, integrity, and confidentiality of a host and its data. By placing managed devices on key network segments, you can examine the packets that traverse your network for malicious activity. The system has several mechanisms it uses to look for the broad range of exploits that attackers have developed.

When the system identifies a possible intrusion, it generates an *intrusion event* (sometimes called by a legacy term, "IPS event"), which is a record of the date, time, the type of exploit, and contextual information about the source of the attack and its target. For packet-based events, a copy of the packet or packets that triggered the event is also recorded. Managed devices transmit their events to the Secure Firewall Management Center

where you can view the aggregated data and gain a greater understanding of the attacks against your network assets.

You can also deploy a managed device as an inline, switched, or routed intrusion system, which allows you to configure the device to drop or replace packets that you know to be harmful.

Tools for reviewing and evaluating intrusion events

Tools for reviewing and evaluating intrusion events are a collection of features that

- provide overview capabilities for current activity on managed devices,
- enable automated alerting and correlation responses to specific intrusion events, and
- support both text-based and graphical reporting with customizable scheduling options.

Available intrusion event tools

These tools help you review intrusion events and evaluate their importance:

- An event summary page that gives you an overview of the current activity on your managed devices
- Text-based and graphical reports that you can generate for any time period you choose; you can also design your own reports and configure them to run at scheduled intervals
- An incident-handling tool that you can use to gather event data related to an attack; you can also add notes to help you track your investigation and response
- Automated alerting that you can configure for SNMP, email, and syslog
- Automated correlation policies that you can use to respond to and remediate specific intrusion events
- Predefined and custom workflows that you can use to drill down through the data to identify the events that you want to investigate further
- External tools for managing and analyzing data. You can send data to those tools using syslog or eStreamer. For more information, refer to Event Analysis Using External Tools

Additionally, you can use publicly-available information such as the predefined resources on the **Analysis** > **Advanced** > **Contextual Cross-Launch** page to learn more about malicious entities.

To search for a particular message string and retrieve documentation for the rule that generated an event, refer to https://www.snort.org/rule_docs/.

License requirements for intrusion events

This reference provides license requirement information for configuring intrusion events in your security deployment.

Threat Defense License

IPS

Requirements and prerequisites for intrusion events

Model support

Any

Supported domains

Any

User roles

- Admin
- Intrusion Admin

View intrusion events

The initial intrusion events view differs depending on the workflow you use to access the page. You can use one of the predefined workflows, which includes one or more drill-down pages, a table view of intrusion events, and a terminating packet view, or you can create your own workflow. You can also view workflows based on custom tables, which may include intrusion events.

An event view may be slow to display if it contains a large number of IP addresses and you have enabled the **Resolve IP Addresses** event view setting.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 You have these choices:

- Adjust time range: Adjust the time range for the event view as described in Change the time window.
- Change workflows: If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- Constrain: To narrow your view to the intrusion events that are important to your analysis, refer to Use intrusion event workflows, on page 22.
- Delete event: To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.
- Mark reviewed: To mark intrusion events reviewed, refer to Mark intrusion events reviewed, on page 17.
- View connection data: To view connection data associated with intrusion events, refer to View connection data associated with intrusion events, on page 17.

- View contents: To view the contents of the columns in the table as described in Intrusion event fields, on page 4.

Intrusion event data overview

An intrusion event captures the date, time, and type of exploit when the system identifies a possible intrusion and provides contextual information about the source of the attack and its target. It also includes a copy of the packet or packets that triggered the event for packet-based events.

You can view intrusion event data in the Secure Firewall Management Center web interface at **Analysis** > **Intrusions** > **Events** or emit data from certain fields as syslog messages for consumption by an external tool. Syslog fields are indicated in the list below; fields without a listed syslog equivalent are not available in syslog messages.

When searching intrusion events, keep in mind that the information available for any individual event can vary depending on how, why, and when system logged the event. For example, only intrusion events triggered on decrypted traffic contain TLS/SSL information.



Note In the Secure Firewall Management Center web interface, some fields in the table view of intrusion events are disabled by default. To enable a field for the duration of your session, expand the search constraints, then click the column name under **Disabled Columns**.

Intrusion event fields

This reference provides field definitions for intrusion events generated by security systems. These fields contain critical information for monitoring, analyzing, and responding to security threats.

Access Control Policy (Syslog: ACPolicy)

The access control policy associated with the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled.

Access Control Rule (Syslog: AccessControlRuleName)

The access control rule that invoked the intrusion policy that generated the event. `Default Action` indicates that the intrusion policy where the rule is enabled is not associated with a specific access control rule but, instead, is configured as the default action of the access control policy.

This field is empty (or, for syslog messages, omitted)if there is:

- No associated rule/default action: Intrusion inspection was associated with neither an access control rule nor the default action, for example, if the packet was examined by the intrusion policy specified to handle packets that must pass before the system can determine which rule to apply. (This policy is specified in the Advanced tab of the access control policy.)
- No associated connection event: The connection event logged for the session has been purged from the database, for example, if connection events have higher turnover than intrusion events.

Application Protocol (Syslog: ApplicationProtocol)

The application protocol, if available, which represents communications between hosts detected in the traffic that triggered the intrusion event.

Application Protocol Category and Tag

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

Application Risk

The risk associated with detected applications in the traffic that triggered the intrusion event: Very High, High, Medium, Low, and Very Low. Each type of application detected in a connection has an associated risk; this field displays the highest risk of those.

Business Relevance

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

Classification (Syslog: Classification)

The classification where the rule that generated the event belongs.

When searching this field, enter the classification number, or all or part of the classification name or description for the rule that generated the events you want to view. You can also enter a comma-separated list of numbers, names, or descriptions. Finally, if you add a custom classification, you can also search using all or part of its name or description.

Client (Syslog: Client)

The client application, if available, which represents software running on the monitored host detected in the traffic that triggered the intrusion event.

Client Category and Tag

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

Connection Counter (Syslog Only)

A counter that distinguishes one connection from another simultaneous connection. This field has no significance on its own.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Connection Instance ID (Syslog Only)

The Snort instance that processed the connection event. This field has no significance on its own.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Count

The number of events that match the information that appears in each row. Note that the Count field appears only after you apply a constraint that creates two or more identical rows. This field is not searchable.

CVE ID

This field is a search field only.

Search by the identification number associated with the vulnerability in MITRE's Common Vulnerabilities and Exposures (CVE) database (https://cve.mitre.org/).

Destination Continent

The continent of the receiving host involved in the intrusion event.

Destination Country

The country of the receiving host involved in the intrusion event.

Destination Host Criticality

The destination host criticality (value of the Host Criticality attribute for that corresponding host) when the event is generated.

Keep in mind that this field is not updated when the criticality of the host changes. However, new events will have the new criticality value.

Destination IP (Syslog: DstIP)

The IP address used by the receiving host involved in the intrusion event.

See also Initiator and responder, source and destination, and sender and receiver fields.

Destination Port / ICMP Code (Syslog: DstPort, ICMPCode)

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, this field displays the ICMP code.

Destination User

The username associated with the Responder IP of the connection event. This host may or may not be the host receiving the exploit. This value is typically known only for users on your network.

.

See also Initiator and responder, source and destination, and sender and receiver fields.

Device

The managed device where the access control policy was deployed.

DeviceUUID (Syslog Only)

The unique identifier of the firewall device that generated an event.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Domain

The domain of the device that detected the intrusion. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Egress Interface (Syslog: EgressInterface)

The egress interface of the packet that triggered the event. This interface column is not populated for a passive interface.

Egress Security Zone (Syslog: EgressZone)

The egress security zone of the packet that triggered the event. This security zone field is not populated in a passive deployment.

Egress Virtual Router

In networks using virtual routing, the name of the virtual router through which traffic exited the network.

Email Attachments

The MIME attachment file name that was extracted from the MIME Content-Disposition header. To display attachment file names, you must enable the SMTP preprocessor **Log MIME Attachment Names** option. Multiple attachment file names are supported.

Email Headers

This field is a search field only.

The data that was extracted from the email header.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option.

Email Recipient

The address of the email recipient that was extracted from the SMTP RCPT TO command. To display a value for this field, you must enable the SMTP preprocessor **Log To Addresses** option. Multiple recipient addresses are supported.

Email Sender

The address of the email sender that was extracted from the SMTP MAIL FROM command. To display a value for this field, you must enable the SMTP preprocessor **Log From Address** option. Multiple sender addresses are supported.

First Packet Time (Syslog Only)

The time the system encountered the first packet.

The following fields collectively uniquely identify the connection event associated with a particular intrusion event: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.

Generator

The component that generated the event.

See also information about the following intrusion event fields: GID, Message, and Snort ID.

GID (Syslog Only)

Generator ID; the ID of the component that generated the event.

See also information about the following intrusion event fields: Generator, Message, and Snort ID.

HTTP Hostname

The host name, if present, that was extracted from the HTTP request Host header. Note that request packets do not always include the host name.

To associate host names with intrusion events for HTTP client traffic, you must enable the HTTP Inspect preprocessor **Log Hostname** option.

In table views, this column displays the first fifty characters of the extracted host name. You can hover your pointer over the displayed portion of an abbreviated host name to display the complete name, up to 256 bytes. You can also display the complete host name, up to 256 bytes, in the packet view.

HTTP Response Code (Syslog: HTTPResponse)

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

HTTP URI

The normalized URI, if present, associated with the HTTP request packet that triggered the intrusion event. Note that request packets do not always include a URI.

To associate URIs with intrusion events for HTTP traffic, you must enable the HTTP Inspect preprocessor **Log URI** option.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly.

This column displays the first fifty characters of the extracted URI. You can hover your pointer over the displayed portion of an abbreviated URI to display the complete URI, up to 2048 bytes. You can also display the complete URI, up to 2048 bytes, in the packet view.

Impact

The impact level in this field indicates the correlation between intrusion data, network discovery data, and vulnerability information.

When searching this field, do not specify impact icon colors or partial strings. For example, do not use **`blue`**, **`level 1`**, or **`0`**. Valid case-insensitive values are:

- Impact 0, Impact Level 0
- Impact 1, Impact Level 1
- Impact 2, Impact Level 2
- Impact 3, Impact Level 3

- Impact 4, Impact Level 4
- Impact 5, Impact Level 5

Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

Ingress Interface (Syslog: IngressInterface)

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface.

Ingress Security Zone (Syslog: IngressZone)

The ingress security zone or tunnel zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment.

Ingress Virtual Router

In networks using virtual routing, the name of the virtual router through which traffic entered the network.

Inline Result (Syslog: InlineResult)

In workflow and table views, this field displays one of the following:

Table 1: Inline Result Field Contents in Workflow and Table Views

This Icon	Indicates
	The system dropped the packet that triggered the rule.
	IPS would have dropped the packet if you enabled the **Drop when Inline** intrusion policy option (in an inline deployment), or if a Drop and Generate rule generated the event while the system was pruning.
	IPS may have transmitted or delivered the packet to the destination, but the connection that contained this packet is now blocked.
No icon (blank)	The triggered rule was not set to Drop and Generate Events

The following table lists the possible reasons for the inline results — Would have dropped and Partially dropped.

Table 2:

Inline Result	Reason	Detailed Reason
Would Have Dropped	Interface in Passive or Tap mode	You have configured the interfaces in inline tap or passive mode.
	Intrusion Policy in "Detection" Inspection Mode	You have set the inspection mode in the intrusion policy to Detection.
	Connection Timed Out	The Snort inspection engine has suspended the inspection as the TCP/IP connection timed out.
Partially Dropped	Connection Closed (0x01)	While creating a new flow, if the allocated flows are more than the allowed number of flows, the Snort inspection engine prunes the least recently used flows.
	Connection Closed (0x02)	When reloading the Snort inspection engine causes a memory adjustment, the engine prunes the least recently used flows.
	Connection Closed (0x04)	When the Snort inspection engine is gracefully shutting down, the engine purges all the active flows.

In a passive deployment, the system does not drop packets, including when an inline interface is in tap mode, regardless of the rule state or the inline drop behavior of the intrusion policy.

When searching this field, enter either of the following:

- **`dropped`** to specify whether the packet is dropped in an inline deployment.
- **`would have dropped`** to specify whether the packet would have dropped if the intrusion policy had been set to drop packets in an inline deployment.
- **`partially dropped`** to specify whether the packet is transmitted or delivered to the destination, but the connection that contained this packet is now blocked.

Intrusion Policy (Syslog: IntrusionPolicy)

The intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event was enabled. You can choose an intrusion policy as the default action for an access control policy, or you can associate an intrusion policy with an access control rule.

IOC (Syslog: NumIOC)

Whether the traffic that triggered the intrusion event also triggered an indication of compromise (IOC) for a host involved in the connection.

When searching this field, specify **`triggered`** or **`n/a`**.

Message (Syslog: Message)

The explanatory text for the event. For rule-based intrusion events, the event message is pulled from the rule. For decoder- and preprocessor-based events, the event message is hard coded.

The Generator and Snort IDs (GID and SID) and the SID version (Revision) are appended in parentheses to the end of each message in the format of numbers separated by colons (GID:SID:version). For example `(1:36330:2)`.

MITRE

A count of techniques that you can click to bring up a modal, which shows the full list of MITRE tactics and techniques within that hierarchy.

MPLS Label (Syslog: MPLS_Label)

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

Network Analysis Policy (Syslog: NAPPolicy)

The network analysis policy, if any, associated with the generation of the event.

This field displays the first fifty characters of the extracted URI. You can hover your pointer over the displayed portion of an abbreviated URI to display the complete URI, up to 2048 bytes. You can also display the complete URI, up to 2048 bytes, in the packet view.

Original Client IP

The original client IP address that was extracted from an X-Forwarded-For (XFF), True-Client-IP, or custom-defined HTTP header.

To display a value for this field, you must enable the HTTP preprocessor **Extract Original Client IP Address** option in the network analysis policy. Optionally, in the same area of the network analysis policy, you can also specify up to six custom client IP headers, as well as set the priority order in which the system selects the value for the Original Client IP event field.

Priority (Syslog: Priority)

The event priority as determined by the Talos Intelligence Group. The priority corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other intrusion events, the priority is determined by the decoder or preprocessor. Valid values are high, medium, and low.

Protocol (Syslog: Protocol)

In the Secure Firewall Management Center web interface, this field is a search field only.

The name or number of the transport protocol used in the connection as listed in http://www.iana.org/assignments/protocol-numbers. This is the protocol associated with the source and destination port/ICMP column.

Reviewed By

The name of the user who reviewed the event. When searching this field, you can enter **`unreviewed`** to search for events that have not been reviewed.

Revision (Syslog Only)

The version of the signature that was used to generate the event.

See also information about the following intrusion event fields: Generator, GID, Message, SID, and Snort ID.

Rule Group

A count of non-MITRE rule groups that you can click to bring up a modal, which shows the full list of rule groups.

Security Context (Syslog: Context)

The metadata identifying the virtual firewall group through which the traffic passed. The system only populates this field for ASA FirePOWER in multiple context mode.

SID (Syslog Only)

The signature ID (also known as the Snort ID) of the rule that generated the event.

See also information about the following intrusion event fields: Generator, GID, Message, Revision, and Snort ID.

Snort ID

This field is a search field only.

(For the syslog field, see SID.)

When performing your search: Specify the Snort ID (SID) of the rule that generated the event or, optionally, specify the combination Generator ID (GID) and SID of the rule, where the GID and SID are separated with a colon (:) in the format GID:SID. You can specify any of the values in the following table:

Table 3: Snort ID Search Values

Value	Example
a single SID	10000
a SID range	10000-11000
greater than a SID	>10000
greater than or equal to a SID	>=10000
less than a SID	<10000
less than or equal to a SID	<=10000
a comma-separated list of SIDs	10000,11000,12000
a single GID:SID combination	1:10000
a comma-separated list of GID:SID combinations	1:10000,1:11000,1:12000
a comma-separated list of SIDs and GID:SID combinations	10000,1:11000,12000

The SID of the events you are viewing is listed in the Message column. For more information, see the description in this section for the Message field.

Source Continent

The continent of the sending host involved in the intrusion event.

Source Country

The country of the sending host involved in the intrusion event.

Source Host Criticality

The source host criticality (value of the Host Criticality attribute for that corresponding host) when the event is generated.

Keep in mind that this field is not updated when the criticality of the host changes. However, new events will have the new criticality value.

Source IP (Syslog: SrcIP)

The IP address used by the sending host involved in the intrusion event.

See also Initiator and responder, source and destination, and sender and receiver fields.

Source Port / ICMP Type (Syslog: SrcPort, ICMPType)

The port number on the sending host. For ICMP traffic, where there is no port number, this field displays the ICMP type.

Source User (Syslog: User)

The username associated with the IP address of the host that initiated the connection, which may or may not be the source host of the exploit. This user value is typically known only for users on your network.

If applicable, the username is preceded by `<realm>\`.

SSL Actual Action (Syslog: SSLActualAction)

In the Secure Firewall Management Center web interface, this field is a search field only.

The action the system applied to encrypted traffic:

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.

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Certificate Information

This field is a search field only.

The information stored on the public key certificate used to encrypt traffic, including:

- Subject/Issuer Common Name
- Subject/Issuer Organization
- Subject/Issuer Organization Unit
- Not Valid Before/After
- Serial Number
- Certificate Fingerprint
- Public Key Fingerprint

SSL Failure Reason

This field is a search field only.

The reason the system failed to decrypt encrypted traffic:

- Unknown
- No Match
- Success
- Uncached Session
- Unknown Cipher Suite
- Unsupported Cipher Suite
- Unsupported SSL Version
- SSL Compression Used
- Session Undecryptable in Passive Mode
- Handshake Error
- Decryption Error
- Pending Server Name Category Lookup
- Pending Common Name Category Lookup
- Internal Error
- Network Parameters Unavailable

- Invalid Server Certificate Handle
- Server Certificate Fingerprint Unavailable
- Cannot Cache Subject DN
- Cannot Cache Issuer DN
- Unknown SSL Version
- External Certificate List Unavailable
- External Certificate Fingerprint Unavailable
- Internal Certificate List Invalid
- Internal Certificate List Unavailable
- Internal Certificate Unavailable
- Internal Certificate Fingerprint Unavailable
- Server Certificate Validation Unavailable
- Server Certificate Validation Failure
- Invalid Action

Field values are displayed in the **SSL Status** field on the search workflow pages.

SSL Status

The action associated with the **SSL Actual Action** (Decryption rule, default action, or undecryptable traffic action) that logged the encrypted connection.

If the system fails to decrypt an encrypted connection, it displays the **SSL Actual Action** (undecryptable traffic action) taken, as well as the **SSL Failure Reason**. For example, if the system detects traffic encrypted with an unknown cipher suite and allows it without further inspection, this field displays `Do Not Decrypt (Unknown Cipher Suite)`.

Click the **Lock icon** to view certificate details.

When searching this field, enter one or more of the **SSL Actual Action** and **SSL Failure Reason** values to view encrypted traffic the system handled or failed to decrypt.

SSL Subject/Issuer Country

This field is a search field only.

A two-character ISO 3166-1 alpha-2 country code for the subject or issuer country associated with the encryption certificate.

Time

The date and time of the event. This field is not searchable.

VLAN ID (Syslog: VLAN_ID)

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

Web Application (Syslog: WebApplication)

The web application, which represents the content or requested URL for HTTP traffic detected in the traffic that triggered the intrusion event.

If the system detects an application protocol of HTTP but cannot detect a specific web application, the system supplies a generic web browsing designation instead.

Web Application Category and Tag

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

Zero Trust Origin User

Unique ID of the user device that requested access to the private application.

Intrusion event impact levels

To help you evaluate the impact an event has on your network, the Secure Firewall Management Center displays an impact level in the table view of intrusion events. For each event, the system adds an impact level icon whose color indicates the correlation between intrusion data, network discovery data, and vulnerability information.



Note Because no operating system information is available for hosts added to the network map from NetFlow data, the system cannot assign Vulnerable (impact level 1: red) impact levels for intrusion events involving those hosts. In such cases, use the host input feature to manually set the operating system identity for the hosts.

The following table describes the possible values for intrusion event impact levels:

Table 4: Impact levels

Impact Level	Vulnerability	Color	Description
Vulnerable (1)	Vulnerable	red	Either: • the source or the destination host is in the network map, and a vulnerability is mapped to the host • the source or destination host is potentially compromised by a virus, trojan, or other piece of malicious software
Potentially Vulnerable (2)	Potentially Vulnerable	orange	Either the source or the destination host is in the network map and one of these is true: • for port-oriented traffic, the port is running a server application protocol • for non-port-oriented traffic, the host uses the protocol

Impact Level	Vulnerability	Color	Description
Currently Not Vulnerable (3)	Currently Not Vulnerable	yellow	Either the source or the destination host is in the network map and one of these is true: • for port-oriented traffic (for example, TCP or UDP), the port is not open • for non-port-oriented traffic (for example, ICMP), the host does not use the protocol
Unknown Target (4)	Unknown Target	blue	Either the source or destination host is on a monitored network, but there is no entry for the host in the network map.
Unknown (0)	Unknown	gray	Neither the source nor the destination host is on a network that is monitored by network discovery.

View connection data associated with intrusion events

The system can log the connections where intrusion events are detected. Although this logging is automatic for intrusion policies associated with access control rules, you must manually enable connection logging to see associated connection data for the default action.

Viewing associated data is most useful when navigating between table views of events.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Events**.

Step 2 Choose the intrusion events using the check boxes in the table, then choose **Connections** from the **Jump to** drop-down list.

Tip You can view the intrusion events associated with particular connections in a similar way. For more information, refer to Navigate between workflows.

Mark intrusion events reviewed

If you are confident that an intrusion event is not malicious, you can mark the event reviewed.

If you have examined an intrusion event and are confident that the event does not represent a threat to your network security (for example, because you know that none of the hosts on your network are vulnerable to the detected exploit), you can mark the event reviewed.

Reviewed events are stored in the event database and are included in the event summary statistics, but no longer appear in the default intrusion event pages. Your name appears as the reviewer.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

If you perform a backup and then delete reviewed intrusion events, restoring your backup restores the deleted intrusion events but does not restore their reviewed status. You view those restored intrusion events under **Intrusion Events**, not under **Reviewed Events**.

Procedure

On a page that displays intrusion events, you have two options:

- To mark one or more intrusion events from the list of events, check the check boxes next to the events and click **Review**.
- To mark all intrusion events from the list of events, click **Review All**.

View previously reviewed intrusion events

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

Step 1 Choose **Analysis** > **Intrusions** > **Reviewed Events**.

Step 2 You have the following choices:

- Adjust the time range as described in Change the time window.
- If you are using a custom workflow that does not include the table view of intrusion events, choose any of the system-provided workflows by clicking **(switch workflow)** next to the workflow title.
- To learn more about the events that appear, refer to Intrusion event fields, on page 4.

Mark reviewed intrusion events unreviewed

You can return a reviewed event to the default intrusion events view by marking the event unreviewed.

In a multidomain deployment, if you mark an event reviewed, the system marks it reviewed in all domains that can view that event.

Procedure

On a page that displays reviewed events, you have two choices:

- To remove individual intrusion events from the list of reviewed events, check the check boxes next to specific events and click **Unreview**.
- To remove all intrusion events from the list of reviewed events, click **Unreview All**.

Preprocessor events

A preprocessor event is a security detection mechanism that

- reports the execution of specified preprocessor options by generating an event whenever a packet triggers that preprocessor option and the associated preprocessor rule is enabled,
- helps detect anomalous protocol exploits by analyzing packet characteristics and protocol behaviors, and
- differs from rule events in that the packet display does not include a detailed rule description for the event.

Preprocessor event characteristics

Preprocessors provide two functions: performing the specified action on the packet (for example, decoding and normalizing HTTP traffic) and reporting the execution of specified preprocessor options by generating an event whenever a packet triggers that preprocessor option and the associated preprocessor rule is enabled.

Preprocessor events differ from rule events in that the packet display does not include a detailed rule description for the event. Instead, the packet display shows the event message, the GID, the SID, the packet header data, and the packet payload. This allows you to analyze the packet's header information, determine if its header options are being used and if they can exploit your system, and inspect the packet payload. After the preprocessors analyze each packet, the rules engine executes appropriate rules against it (if the preprocessor was able to defragment it and establish it as part of a valid session) to further analyze potential content-level threats and report on them.

Double encoding detection example

You can enable the `Double Encoding` HTTP Inspect option and the associated preprocessor rule with the HTTP Inspect Generator (GID) 119 and the Snort ID (SID) 2 to generate an event when the preprocessor encounters IIS double-encoded traffic.

Attackers can craft overlapping IP fragments to cause a DoS attack on a host. The IP defragmentation preprocessor can detect this type of attack and generate an intrusion event for it.

Preprocessor generator IDs

Each preprocessor has its own Generator ID number, or GID, that indicates which preprocessor was triggered by the packet. Some of the preprocessors also have related SIDs, which are ID numbers that classify potential attacks. This helps you analyze events more effectively by categorizing the type of event much the way a

rule's Snort ID (SID) can offer context for packets triggering rules. You can list preprocessor rules by preprocessor in the Preprocessors filter group on the intrusion policy Rules page; you can also list preprocessor rules in the preprocessor and packet decoder sub-groupings in the Category filter group.



Note Events generated by standard text rules have a generator ID of 1 (Global domain or legacy GID) or 1000 - 2000 (descendant domains). For shared object rules, the events have a generator ID of 3. For both, the event's SID indicates which specific rule triggered.

Table 5: Generator IDs

ID	Component	Description
1	Standard Text Rule	The event was generated when the packet triggered a standard text rule (Global domain or legacy GID).
2	Tagged Packets	The event was generated by the Tag generator, which generates packets from a tagged session. This occurs when the `tag` rule option is used.
3	Shared Object Rule	The event was generated when the packet triggered a shared object rule.
102	HTTP Decoder	The decoder engine decoded HTTP data within the packet.
105	Back Orifice Detector	The Back Orifice Detector identified a Back Orifice attack associated with the packet.
106	RPC Decoder	The RPC decoder decoded the packet.
116	Packet Decoder	The event was generated by the packet decoder.
119, 120	HTTP Inspect Preprocessor	The event was generated by the HTTP Inspect preprocessor. GID 120 rules relate to server-specific HTTP traffic.
122	Portscan Detector	The event was generated by the portscan flow detector.
123	IP Defragmentor	The event was generated when a fragmented IP datagram could not be properly reassembled.
124	SMTP Decoder	The event was generated when the SMTP preprocessor detected an exploit against an SMTP verb.
125	FTP Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within FTP traffic.
126	Telnet Decoder	The event was generated when the FTP/Telnet decoder detected an exploit within telnet traffic.
128	SSH Preprocessor	The event was generated when the SSH preprocessor detected an exploit within SSH traffic.

ID	Component	Description
129	Stream Preprocessor	The event was generated during stream preprocessing by the stream preprocessor.
131	DNS Preprocessor	The event was generated by the DNS preprocessor.
133	DCE/RPC Preprocessor	The event was generated by the DCE/RPC preprocessor.
134	Rule Latency Packet Latency	The event was generated when rule latency suspended (134:1) or re-enabled (134:2) a group of intrusion rules, or when the system stopped inspecting a packet because the packet latency threshold was exceeded (134:3).
135	Rate-Based Attack Detector	The event was generated when a rate-based attack detector identified excessive connections to hosts on the network.
137	SSL Preprocessor	The event was generated by the TLS/SSL preprocessor.
138, 139	Sensitive Data Preprocessor	The event was generated by the sensitive data preprocessor.
140	SIP Preprocessor	The event was generated by the SIP preprocessor.
141	IMAP Preprocessor	The event was generated by the IMAP preprocessor.
142	POP Preprocessor	The event was generated by the POP preprocessor.
143	GTP Preprocessor	The event was generated by the GTP preprocessor.
144	Modbus Preprocessor	The event was generated by the Modbus SCADA preprocessor.
145	DNP3 Preprocessor	The event was generated by the DNP3 SCADA preprocessor.
148	CIP Preprocessor	The event was generated by the CIP SCADA preprocessor.
149	S7Commplus preprocessor	The event was generated by the S7Commplus SCADA preprocessor.
1000 - 2000	Standard Text Rule	The event was generated when the packet triggered a standard text rule (descendant domains).

Intrusion event workflow pages

Intrusion event workflow pages are a set of predefined workflows that

- provide event data populated interfaces to view and analyze intrusion events generated by preprocessor, decoder, and intrusion rules violations,
- step you through a series of pages to help pinpoint specific intrusion events for evaluation, and
- contain three different types of pages or event views for comprehensive analysis.

Workflow page types and functionality

The preprocessor, decoder, and intrusion rules that are enabled in the current intrusion policy generate intrusion events whenever the traffic that you monitor violates the policy.

The predefined intrusion event workflows contain three different types of pages, or event views:

- one or more drill-down pages
- the table view of intrusion events
- a packet view

Drill-down pages generally include two or more columns in a table (and, for some drill-down views, more than one table) that allow you to view one specific type of information.

When you "drill down" to find more information for one or more destination ports, you automatically select those events and the next page in the workflow appears. In this way, drill-down tables help you reduce the number of events you are analyzing at one time.

The initial *table view* of intrusion events lists each intrusion event in its own row. The columns in the table list information such as the time, the source IP address and port, the destination IP address and port, the event priority, the event message, and more.

When you select events on a table view, instead of selecting events and displaying the next page in the workflow, you add to what are called *constraints*. Constraints are limits that you impose on the types of events that you want to analyze.

For example, if you click **Close** (✕) in any column and clear **Time** from the drop-down list, you can remove Time as one of the columns. To narrow the list of events in your analysis, you can click the link for a value in one of the rows in the table view. For example, to limit your analysis to the events generated from one of the source IP addresses (presumably, a potential attacker), click the IP address in the **Source IP Address** column.

If you select one or more rows in a table view and then click **View**, the packet view appears. A *packet view* provides information about the packet that triggered the rule or the preprocessor that generated the event. Each section of the packet view contains information about a specific layer in the packet. You can expand collapsed sections to see more information.



Note Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view.

If the predefined workflows do not meet your specific needs, you can create custom workflows that display only the information you are interested in. Custom intrusion event workflows can include drill-down pages, a table view of events, or both; the system automatically includes a packet view as the last page. You can easily switch between the predefined workflows and your own custom workflows depending on how you want to investigate events.

Use intrusion event workflows

To avoid displaying the same intrusion events on different workflow pages, the time range pauses when you click a link at the bottom of the page to display another page of events, and resumes when you click to take any other action on the subsequent page.

Tip At any point in the process, you can save the constraints as a set of search criteria. For example, if you find that over the course of a few days your network is being probed by an attacker from a single IP address, you can save your constraints during your investigation and then use them again later. You cannot, however, save compound constraints as a set of search criteria.

Procedure

Step 1 Access an intrusion event workflow using **Analysis** > **Intrusions** > **Events**.

Step 2 Optionally, constrain the number of intrusion events that appear on the event views as described in Intrusion event drill-down page constraints, on page 24 or Intrusion event table view constraints, on page 25.

Step 3 Choose from the available workflow options.

Event analysis and intelligence

- To learn more about the columns that appear, refer to Intrusion event fields, on page 4.
- To view a host's profile, click **Host Profile** that appears next to the host IP address.
- To view geolocation details, click flag that appears in the Source Country or Destination Country columns.
- To view data in available sources external to your system, right-click an event value. The options you see depend on the data type and include public sources; other sources depend on the resources you have configured. For information, refer to Event investigation using web-based resources
- To gather general intelligence about an event, right-click an event value in the table and choose from a Cisco or third-party intelligence source. For example, you can get details about a suspicious IP address from Cisco Talos. The options you see depend on the data type and the integrations that are configured on your system. For more information, refer to Event investigation using web-based resources.

Time and data management

- To modify the time and date range for displayed events, refer to Change the time window.

Tip If no intrusion events appear on the event views, adjusting the specified time range might return results. If you specified an older time range, events in that time range might have been deleted. Adjusting the rule thresholding configuration might generate events.

Note Events generated outside the appliance's configured time window (whether global or event-specific) may appear in an event view if you constrain the event view by time. This may occur even if you configured a sliding time window for the appliance.

Event actions

- To delete events from the event database, check the check boxes next to events you want to delete, then click **Delete**, or click **Delete All**.
- To mark events reviewed to remove them from intrusion event pages, but not the event database, refer to Mark intrusion events reviewed, on page 17.
- To download a local copy of the packet (a packet capture file in libpcap format) that triggered each selected event, check the check boxes next to events triggered by the packets you want to download, then click **Download Packets**, or click **Download All Packets**. Captured packets are saved in libpcap format. This format is used by several popular protocol analyzers.

Navigation and workflow

- To sort events on the current workflow page or navigate within the current workflow page, refer to Use workflows.
- To navigate between pages in the current workflow, keeping the current constraints, click the appropriate page link at the top left of the workflow page.
- To navigate to other event views to view associated events, refer to Navigate between workflows.
- To temporarily use a different workflow, click **(switch workflow)**.
- To bookmark the current page so that you can quickly return to it, click **Bookmark This Page**.
- To view the Intrusion Events section of the Summary Dashboard, click **Dashboards**.
- To navigate to the bookmark management page, click **View Bookmarks**.
- To generate a report based on the data in the current view, refer to Create a report template from an event view.

Intrusion event drill-down page constraints

The table describes how to use drill-down pages to constrain events in intrusion event workflows.

Table 6: Constraining events on drill-down pages

To...	You can...
drill down to the next workflow page constraining on a specific value	click the value. For example, on the Destination Port workflow, to constrain the events to those with a destination of port 80, click **80/tcp** in the **DST Port/ICMP Code** column. The next page of the workflow, Events, appears and contains only port 80/tcp events.

To...	You can...
drill down to the next workflow page constraining on selected events	select the check boxes next to the events you want to view on the next workflow page, then click **View**. For example, on the Destination Port workflow, to constrain the events to those with destination ports 20/tcp and 21/tcp, select the check boxes next to the rows for those ports and click **View**. The next page of the workflow, Events, appears and contains only port 20/tcp and 21/tcp events. Note that if you constrain on multiple rows and the table has more than one column (not including a Count column), you build what is called a compound constraint. Compound constraints ensure that you do not include more events in your constraint than you mean to. For example, if you use the Event and Destination workflow, each row that you select on the first drill-down page creates a compound constraint. If you pick event 1:100 with a destination IP address of 10.10.10.100 and you also pick event 1:200 with a destination IP address of 192.168.10.100, the compound constraint ensures that you do not also select events with 1:100 as the event type and 192.168.10.100 as the destination IP address or events with 1:200 as the event type and 10.10.10.100 as the destination IP address.
drill down to the next workflow page keeping the current constraints	click **View All**.

Intrusion event table view constraints

The table describes how to use the table view.

Table 7: Constraining events on the table view of events

To...	You can...
constrain the view to events with a single attribute	click the attribute. For example, to constrain the view to events with a destination of port 80, click **80/tcp** in the **DST Port/ICMP Code** column.
remove a column from the table	click **Close** (✕) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**. If you want to hide or show other columns, select or clear the appropriate check boxes before you click **Apply**. To add a disabled column back to the view, click the **expand arrow** to expand the search constraints, then click the column name under **Disabled Columns**.
view the packets associated with one or more events	either: • click the **down arrow** next to the event whose packets you want to view. • select one or more events whose packets you want to view, and, at the bottom of the page, click **View.** • at the bottom of the page, click **View All**to view the packets for all events that match the current constraints.

Using the intrusion event packet view

A packet view provides information about the packet data associated with the rule that generated an intrusion event.

Tip The packet view on a Secure Firewall Management Center does not contain packet information when the **Transfer Packets** option is disabled for the device detecting the event.

When the **Transfer Packets** option is enabled, the device sends event metadata and packet data associated with the intrusion event to the Secure Firewall Management Center. The packet view does not necessarily contain the complete original packet as transmitted on the wire or the complete traffic session. For intrusion events generated by Snort inspection, the displayed packet data can contain only the Protocol Data Unit (PDU) or packet content associated with the rule match.

The packet view indicates why a specific packet data unit was captured by providing information about the intrusion event that it triggered, including the event's time stamp, message, classification, priority, and, if the event was generated by a standard text rule, the rule that generated the event. The packet view also provides general information about the displayed packet data, such as its size.

In addition, the packet view has sections that describe the available data link, network, and transport layer information, as well as the bytes included in the event packet data. If the system decrypted traffic, you can view the available decrypted bytes. You can expand collapsed sections to display detailed information.

Note Because each portscan event is triggered by multiple packets, portscan events use a special version of the packet view.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 On the table view of intrusion events, choose packets to view as described in Intrusion event table view constraints, on page 25.

Step 2 Optionally, if you chose more than one event, you can page through the packets in the packet view by using the page numbers at the bottom of the page.

Step 3 You also have these options:

- Adjust: To modify the date and time range in the packet views, refer to Change the time window.
- Configure: To configure the intrusion rule that triggered the event, click the arrow next to Actions and continue as described in Configure intrusion rules within the packet view, on page 30.
- Delete: To delete an event from the database, click **Delete** to delete the event whose packet you are viewing or click **Delete All** to delete all the events whose packets you previously selected.
- Download: To download a local copy of the packet (a packet capture file in libpcap format) that triggered the event, click **Download Packet**to save a copy of the captured packet for the event you are viewing or click **Download All Packets** to save copies of the captured packets for all the events whose packets

you previously selected. The captured packet is saved in libpcap format. This format is used by several popular protocol analyzers.

Note You cannot download a portscan packet because single portscan events are based on multiple packets; however, the portscan view provides all usable packet information. You must have at least 15% available disk space in order to download.

- Mark reviewed: To mark an event reviewed to remove it from event views, but not the event database, click **Review** to mark the event whose packet you are viewing or click **Review All** to mark all the events whose packets you previously selected. For more information, refer to Mark intrusion events reviewed, on page 17.
- View additional information: To expand or collapse a page section, click the arrow next to the section. For details, refer to Event information fields, on page 27, Frame information fields, on page 33, and Data link layer information fields, on page 34.
- View network layer information: Refer to View network layer information, on page 34.
- View packet byte information: Refer to View packet byte information, on page 40.
- View transport layer information: Refer to View transport layer information, on page 37

Event information fields

On the packet view, you can view information about the packet in the event information section.

Event

The event message. For rule-based events, this corresponds to the rule message. For other events, this is determined by the decoder or preprocessor.

The ID for the event is appended to the message in the format (GID:SID:Rev). GID is the generator ID of the rules engine, the decoder, or the preprocessor that generated the event. SID is the identifier for the rule, decoder message, or preprocessor message. Rev is the revision number of the rule.

Timestamp

The time that the packet was captured, in UTC time zone.

Classification

The event classification. For rule-based events, this corresponds to the rule classification. For other events, this is determined by the decoder or preprocessor.

Priority

The event priority. For rule-based events, this corresponds to either the value of the `priority` keyword or the value for the `classtype` keyword. For other events, this is determined by the decoder or preprocessor.

Ingress security zone

The ingress security zone of the packet that triggered the event. Only this security zone field is populated in a passive deployment.

Egress security zone

The egress security zone of the packet that triggered the event. This field is not populated in a passive deployment.

Domain

The domain where the managed device belongs. This field is only present if you have ever configured the Firewall Management Center for multitenancy.

Device

The managed device where the access control policy was deployed.

Ingress interface

The ingress interface of the packet that triggered the event. Only this interface column is populated for a passive interface.

Egress interface

For an inline set, the egress interface of the packet that triggered the event.

Source/Destination IP

The host IP address or domain name where the packet that triggered the event (source) originated, or the target (destination) host of the traffic that triggered the event.

Source Port/ICMP type

Source port of the packet that triggered the event. For ICMP traffic, where there is no port number, the system displays the ICMP type.

Destination Port/ICMP code

The port number for the host receiving the traffic. For ICMP traffic, where there is no port number, the system displays the ICMP code.

Email headers

The data that was extracted from the email header. Note that email headers do not appear in the table view of intrusion events, but you can use email header data as a search criterion.

To associate email headers with intrusion events for SMTP traffic, you must enable the SMTP preprocessor **Log Headers** option. For rule-based events, this row appears when email data is extracted.

HTTP hostname

The host name, if present, extracted from the HTTP request Host header. This row displays the complete host name, up to 256 bytes. You can expand the complete host name if it is longer than a single row.

To display host names, you must enable the HTTP Inspect preprocessor **Log Hostname** option.

Note that HTTP request packets do not always include a host name. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

HTTP URI

The normalized URI, if present, associated with the HTTP request packet that triggered the intrusion event. This row displays the complete URI, up to 2048 bytes. You can expand the complete URI if it is longer than a single row.

To display the URI, you must enable the HTTP Inspect preprocessor **Log URI** option.

Note that HTTP request packets do not always include a URI. For rule-based events, this row appears when the packet contains the HTTP host name or the HTTP URI.

To see the associated HTTP URI in intrusion events triggered by HTTP responses, you should configure HTTP server ports in the **Perform Stream Reassembly on Both Ports** option; note, however, that this increases resource demands for traffic reassembly.

Intrusion policy

The intrusion policy, if present, where the intrusion, preprocessor, or decoder rule that generated the intrusion event was enabled. You can choose an intrusion policy as the default action for an access control policy or associate an intrusion policy with an access control rule.

Access control policy

The access control policy that includes the intrusion policy where the intrusion, preprocessor, or decoder rule that generated the event is enabled.

Access control rule

The access control rule associated with an intrusion rule that generated the event. Default Action indicates that the intrusion policy where the rule is enabled is not associated with an access control rule but, instead, is configured as the default action of the access control policy.

Rule

For standard text rule events, the rule that generated the event.

Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Because rule data may contain sensitive information about your network, administrators may toggle users' ability to view rule information in the packet view with the View Local Rules permission in the user role editor.

Actions

For standard text and custom rule events, expand **Actions** to take any of these actions on the rule that triggered the event:

- edit the rule
- view documentation for the revision of the rule; for standard text rules only, after clicking **View Documentation** under Actions, you can click **Rule Documentation** in the documentation pop-up window to view more-specific rule details.

- add a comment to the rule
- change the state of the rule
- set a threshold for the rule
- suppress the rule

If the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Configure intrusion rules within the packet view

Within the packet view of an intrusion event, you can take several actions on the rule that triggered the event. Note that if the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 You have the following choices:

- Comment: For standard text rule events, click **Rule Comment** to add a text comment to the rule that generated the event. This allows you to provide additional context and information about the rule and the exploit or policy violation it identifies. You can also add and view rule comments in the intrusion rules editor.
- Disable: To disable this rule, click one of the following options:
  - **Disable this rule in the current Snort 2 policy (<policy_name>)**
  - **Disable this rule in all locally created Snort 2 policies**

  If this event is generated by a standard text rule, you can disable the rule, if necessary. You can set the rule in all policies that you can edit locally. Alternately, you can set the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system.

Note You cannot disable shared object rules from the packet view, nor can you disable rules in the default policies.

- Drop packets and generate an event: To set the rule to drop packets that trigger it and generate an event, click one of the following options:
  - **Set this rule to drop the triggering packet and generate an event in the current Snort 2 policy (<policy_name>)**
  - **Set this rule to drop the triggering packet and generate an event in all locally created Snort 2 inline policies**

  If your managed device is deployed inline on your network, you can set the rule that triggered the event to drop packets that trigger the rule in all policies that you can edit locally. Alternately, you can set the

rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system. Note also that this option appears only when **Drop when Inline** is enabled in the current policy.

- Edit: For standard text rule events, click **Edit** (to edit the Snort 2 rule) or **Edit Snort 3 Rule** to modify the rule that generated the event. If the event is based on a shared object rule, a decoder, or a preprocessor, the rule is not available.

Note If you edit a system-provided rule (as opposed to a custom standard text rule), you actually create a new local rule. Make sure you set the local rule to generate events and also disable the original rule in the current intrusion policy. Note, however, that you cannot enable local rules in the default policies.

- Generate events: Click **Set this rule to generate events in all locally created Snort 2 policies** to set the rule to generate events.

  If this event is generated by a standard text rule, you can set the rule to generate events in all policies that you can edit locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by the system.

Note You cannot set shared object rules to generate events from the packet view, nor can you disable rules in the default policies.

- Set suppression options: Expand **Set Suppression Options** and continue as described in Set suppression options within the packet view, on page 32.

  You can use this option to suppress the rule that triggered this event in all policies that you can edit locally. Alternately, you can suppress the rule only in the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Cisco.

- Set threshold options: Expand **Set Thresholding Options** and continue as described in Set threshold options within the packet view, on page 32.

  You can use this option to create a threshold for the rule that triggered this even in all policies that you can edit locally. Alternately, you create a threshold only for the current policy (that is, the policy that generated the event) if you can edit the current policy locally.

  Note that the current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default intrusion policy provided by the system.

- View documentation: Click **View Documentation** to learn more about the rule that generated the event. Optionally, then click **Rule Documentation** to view more-specific rule details.

Set threshold options within the packet view

You can set threshold options in all policies that you can edit locally or, when it can be edited locally, only in the current policy (that is, the policy that caused the event to be generated).

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Thresholding Options** and choose one of the two possible options:

- **in the current Snort 2 policy (<policy_name>)**
- **in all locally created Snort 2 policies**

Step 3 Choose the type of threshold you want to set:

- Click **limit** to limit notification to the specified number of event instances per time period.
- Click **threshold** to provide notification for each specified number of event instances per time period.
- Click **both** to provide notification once per time period after a specified number of event instances.

Step 4 Click the appropriate threshold to indicate whether you want the event instances tracked by **Source** or **Destination** IP address.

Step 5 In the **Count** field, enter the number of event instances you want to use as your threshold.

Step 6 In the **Seconds** field, enter a number between 1 and 86400 that specifies the time period for which event instances are tracked.

Step 7 If you want to override any current thresholds for this rule in existing intrusion policies, check the **Override any existing settings for this rule** check box.

Step 8 Click **Save Thresholding**.

Set suppression options within the packet view

This task is performed within the packet view of an intrusion event that was generated by an intrusion rule.

Before you begin

Follow these steps to set suppression options within the packet view:

Procedure

Step 1 Within the packet view of an intrusion event that was generated by an intrusion rule, expand **Actions** in the Event Information section.

Step 2 Expand **Set Suppression Options** and click one of the two possible options:

- **in the current Snort 2 policy (<policy_name>)**

- **in all locally created Snort 2 policies**

Note The current policy option appears only when you can edit the current policy; for example, you can edit a custom policy, but you cannot edit a default policy provided by Cisco.

Step 3 Choose one of the following **Track By** options:

- Click **Source** to suppress events generated by packets originating from a specified source IP address.
- Click **Destination** to suppress events generated by packets going to a specified destination IP address.
- Click **Rule** to completely suppress events for the rule that triggered this event.

Step 4 In the **IP address or CIDR block** field, enter the IP address or CIDR block/prefix length you want to specify as the source or destination IP address.

Step 5 Click **Save Suppression**.

Frame information fields

On the packet view, click the arrow next to **Frame** to view information about the captured frame. The packet view may display a single frame or multiple frames. Each frame provides information about an individual network packet. You would see multiple frames, for example, in the case of tagged packets or packets in reassembled TCP streams.

Frame information fields provide these details:

Frame identification and timing

- **Frame n**: The captured frame, where *n* is 1 for single-frame packets and the incremental frame number for multi-frame packets. The number of captured bytes in the frame is appended to the frame number.
- **Arrival Time**: The date and time the frame was captured.
- **Time delta from previous captured frame**: For multi-frame packets, the elapsed time since the previous frame was captured.
- **Time delta from previous displayed frame**: For multi-frame packets, the elapsed time since the previous frame was displayed.
- **Time since reference or first frame**: For multi-frame packets, the elapsed time since the first frame was captured.
- **Frame Number**: The incremental frame number.

Frame characteristics

- **Frame Length**: The length of the frame in bytes.
- **Capture Length**: The length of the captured frame in bytes.
- **Frame is marked**: Whether the frame is marked (true or false).
- **Protocols in frame**: The protocols included in the frame.

Data link layer information fields

On the packet view, click the arrow next to the data link layer protocol (for example, **Ethernet II**) to view the data link layer information about the packet, which contains the 48-bit media access control (MAC) addresses for the source and destination hosts. It may also display other information about the packet, depending on the hardware protocol.



Note Note that this example discusses Ethernet link layer information; other protocols may also appear.

The packet view reflects the protocol used at the data link layer. This reference describes the information you might see for an Ethernet II or IEEE 802.3 Ethernet packet in the packet view.

Destination

The MAC address for the destination host.



Note Ethernet can also use multicast and broadcast addresses as the destination address.

Source

The MAC address for the source host.

Type

For Ethernet II packets, the type of packet that is encapsulated in the Ethernet frame; for example, IPv6 or ARP datagrams. Note that this item only appears for Ethernet II packets.

Length

For IEEE 802.3 Ethernet packets, the total length of the packet, in bytes, not including the checksum. Note that this item only appears for IEEE 802.3 Ethernet packets.

View network layer information

When analyzing network packets, you need to examine the network layer protocol details to understand packet routing, addressing, and protocol-specific information. This view provides comprehensive data about IP packets and other network layer protocols.

Procedure

On the packet view, click the arrow next to the network layer protocol (for example, **Internet Protocol**) to view more detailed information about network layer information related to the packet.

Note Note that this example discusses IP packets; other protocols may also appear.

IPv4 network layer information fields

This listing describes protocol-specific information that might appear in an IPv4 packet.

Version

The Internet Protocol version number.

Header length

The number of bytes in the header, including any IP options. An IP header with no options is 20 bytes long.

Differentiated services field

The values for differentiated services that indicate how the sending host supports Explicit Congestion Notification (ECN):

- `0x0`: does not support ECN-Capable Transport (ECT)
- `0x1` and `0x2`: supports ECT
- `0x3`: Congestion Experienced (CE)

Total length

The length of the IP packet, in bytes, minus the IP header.

Identification

The value that uniquely identifies an IP datagram sent by the source host. This value is used to trace fragments of the same datagram.

Flags

The values that control IP fragmentation, where:

values for the Last Fragment flag indicate whether there are more fragments associated with the datagram:

- `0`: there are no more fragments associated with the datagram
- `1`: there are more fragments associated with the datagram

values for the Don't Fragment flag control whether the datagram can be fragmented:

- `0`: the datagram can be fragmented
- `1`: the datagram must **not** be fragmented

Fragment offset

The value for the fragment offset from the beginning of the datagram.

Time to live (ttl)

The remaining number of hops that the datagram can make between routers before the datagram expires.

Protocol

The transport protocol that is encapsulated in the IP datagram; for example, ICMP, IGMP, TCP, or UDP.

Header checksum

The indicator for whether the IP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an intrusion evasion attempt.

Source/Destination

The IP address or domain name for the source (or destination) host.

Note that to display the domain name, you must enable IP address resolution.

Click the address or domain name to view the context menu, then select **Whois** to do a whois search on the host, **View Host Profile** to view host information, or choose an option to add the address to a global Block list or Do-Not-Block list.

IPv6 network layer information fields

This listing describes protocol-specific information that might appear in an IPv6 packet.

Traffic class

An experimental 8-bit field in the IPv6 header for identifying IPv6 packet classes or priorities similar to the differentiated services functionality provided for IPv4. When unused, this field is set to zero.

Flow label

A optional 20-bit IPv6 hexadecimal value 1 to FFFFF that identifies a special flow such as non-default quality of service or real-time service. When unused, this field is set to zero.

Payload length

A 16-bit field identifying the number of octets in the IPv6 payload, which is comprised of all of the packet following the IPv6 header, including any extension headers.

Next header

An 8-bit field identifying the type of header immediately following the IPv6 header, using the same values as the IPv4 Protocol field.

Hop limit

An 8-bit decimal integer that each node that forwards the packet decrements by one. The packet is discarded if the decremented value reaches zero.

Source

The 128-bit IPv6 address for the source host.

Destination

The 128-bit IPv6 address for the destination host.

View transport layer information

Use this procedure when analyzing network packets to examine transport layer details including TCP, UDP, and ICMP protocol information and payload data.

Procedure

Step 1 On the packet view, click the arrow next to the transport layer protocol (for example, **TCP**, **UDP**, or **ICMP**).

Step 2 Optionally, click **Data** when present to view the first twenty-four bytes of the payload for the protocol immediately above it in the Packet Information section of the packet view.

Step 3 View the contents of the transport layer for TCP, UDP, and ICMP protocols as described in TCP packet view fields, on page 37, UDP packet view fields, on page 38, or ICMP packet view fields, on page 39.

Note These examples discuss TCP, UDP, and ICMP packets; other protocols may also appear.

TCP packet view fields

This section describes the protocol-specific information for a TCP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Sequence number

The value for the first byte in the current TCP segment, keyed to initial sequence number in the TCP stream.

Next sequence number

In a response packet, the sequence number of the next packet to send.

Acknowledgment number

The TCP acknowledgment, which is keyed to the sequence number of the previously accepted data.

Header length

The number of bytes in the header.

Flags

The six bits that indicate the TCP segment's transmission state:

- `U`: The urgent pointer is valid
- `A`: The acknowledgment number is valid
- `P`: The receiver should push data
- `R`: Reset the connection
- `S`: Synchronize sequence numbers to start a new connection
- `F`: The sender has finished sending data

Window size

The amount of unacknowledged data, in bytes, that the receiving host will accept.

Checksum

The indicator for whether the TCP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit or may be being used in an in evasion attempt.

Urgent pointer

The position, if present, in the TCP segment where the urgent data ends. Used in conjunction with the `U` flag.

Options

The values, if present, for TCP options.

UDP packet view fields

This section describes the protocol-specific information for a UDP packet.

Source port

The number that identifies the originating application protocol.

Destination port

The number that identifies the receiving application protocol.

Length

The combined length of the UDP header and data.

Checksum

The indicator for whether the UDP checksum is valid. If the checksum is invalid, the datagram may have been corrupted during transit.

ICMP packet view fields

The Internet Control Message Protocol (ICMP) packets contain several protocol-specific fields. These fields help identify the type of message, additional codes for certain message types, and ensure packet integrity.

Type

The type field indicates the kind of ICMP message:

Echo and error reporting

- 0: Echo reply
- 3: Destination unreachable
- 8: Echo request
- 11: Time exceeded
- 12: Parameter problem

Routing and flow control

- 4: Source quench
- 5: Redirect
- 9: Router advertisement
- 10: Router solicitation

Timestamp and address management

- 13: Timestamp request
- 14: Timestamp reply
- 17: Address mask request
- 18: Address mask reply

Obsolete messages

- 15: Information request (obsolete)
- 16: Information reply (obsolete)

Code

This field provides additional context for certain ICMP message types. For example, ICMP message types 3, 5, 11, and 12 use codes as defined in RFC 792 to further specify the message's meaning.

Checksum

This field indicates whether the ICMP packet is valid. If the checksum is invalid, the datagram may have been corrupted while being transmitted.

View packet byte information

Use this procedure when you need to analyze intrusion-event packet contents at the byte level. The displayed bytes represent the packet data transferred for the intrusion event and do not necessarily represent the complete original packet as transmitted on the wire or the complete traffic session. If the system decrypted the traffic, you can also view the available decrypted packet bytes.

Procedure

On the packet view, click the arrow next to **Packet Bytes** to view hexadecimal and ASCII versions of the bytes included in the intrusion-event packet data. If the system decrypted the traffic, you can also view the available decrypted packet bytes.

Internally sourced intrusion events

An internally sourced intrusion event is a security incident that originates from internal sources and indicates a compromised host on your network.

Investigation indicators

If the source IP address is on your network, this is a sign that you should investigate this host.

View intrusion event statistics

The Intrusion Event Statistics page provides you with a quick summary of the current state of your appliance and any intrusion events generated for your network.

Each of the IP addresses, ports, protocols, event messages, and so on shown on the page is a link. Click any link to view the associated event information. For example, if one of the top 10 destination ports is `80 (http)/tcp`, clicking that link displays the first page in the default intrusion events workflow, and lists the events targeting that port. Note that only the events (and the managed devices that generate events) in the current time range appear. Also, intrusion events that you have marked reviewed continue to appear in the statistics. For example, if the current time range is the past hour but the first event was generated five hours ago, when you click the **First Event** link, the resulting event pages will not show the event until you change the time range.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Statistics**.

Step 2 From the two selection boxes at the top of the page, choose the zones and devices whose statistics you want to view, or choose **All Security Zones** and **All Devices** to view statistics for all the devices that are collecting intrusion events.

Step 3 Click **Get Statistics**.

Tip To view data from a custom time range, click the link in the upper right page area and follow the directions in Change the time window.

Host statistics

The Host Statistics section of the intrusion event statistics page provides information about the appliance itself. On the Secure Firewall Management Center, this section also provides information about any managed devices.

This information includes these metrics:

- **Time**: The current time on the appliance.
- **Uptime**: The number of days, hours, and minutes since the appliance itself was restarted. On the Secure Firewall Management Center, the uptime also shows the last time each managed device was rebooted, the number of users logged in, and the load average.
- **Disk Usage**: The percentage of the disk that is being used.
- **Memory Usage**: The percentage of system memory that is being used.
- **Load Average**: The average number of processes in the CPU queue for the past 1 minute, 5 minutes, and 15 minutes.

Event overview

The event overview section of the intrusion event statistics page provides an overview of the information in the intrusion event database.

These statistics include the following:

- **Events**: The number of events in the intrusion event database.
- **Events in Time Range**: The currently selected time range as well as the number and percentage of events from the database that fall within the time range.
- **First Event**: The event message for the first event in the event database.
- **Last Event**: The event message for the last event in the event database.

Note If you select a managed device while viewing intrusion event data on the Secure Firewall Management Center, the Event Overview section for that device appears instead.

Event statistics

The Event Statistics section of the Intrusion Event Statistics page provides more specific information about the information in the intrusion event database.

This information includes details on:

- the top 10 event types
- the top 10 source IP addresses
- the top 10 destination IP addresses
- the top 10 destination ports
- the protocols, ingress and egress security zones, and devices with the greatest number of events

Note In a multidomain deployment, the system builds a separate network map for each leaf domain. As a result, a leaf domain can contain an IP address that is unique within its network, but identical to an IP address in another leaf domain. When you view event statistics in an ancestor domain, the system may display multiple instances of that repeated IP address. At first glance, they might appear to be duplicate entries. However, if you drill down to the host profile information for each IP address, the system shows that they belong to different leaf domains.

View intrusion event performance graphs

Graphs can be generated to reflect number of intrusion events per second, number of megabits per second, average number of bytes per packet, the percent of packets uninspected by Snort, and the number of packets blocked as the result of TCP normalization. These graphs can show statistics for the last hour, last day, last week, or last month of operation.

Note New data is accumulated for statistics graphs every five minutes. Therefore, if you reload a graph quickly, the data may not change until the next five-minute increment occurs. Each graph displays *average* values in the intervals shown (day, hour, or five minutes) for the selected time period (last month, week, day, or hour). Decimal values are displayed when the average is less than one.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Performance**.

Step 2 From the **Select Device** list, choose the devices whose data you want to view.

Step 3 From the **Select Graph(s)** list, choose the type of graph you want to create as described in Intrusion event performance statistics graph types, on page 43.

Step 4 From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph**.

Step 6 To save the graph, right-click it and follow the instructions for your browser to save the image.

Intrusion event performance statistics graph types

The table lists the available graph types for intrusion event performance statistics. Graph types may display differently if they are populated with data affected by the network analysis policy **Inline Mode** setting. If **Inline Mode** is disabled, graph types marked with an asterisk (`*`) in the web interface (a `yes` in the column below) populate with data about the traffic the system would have modified or dropped if **Inline Mode** were enabled.

Table 8: Intrusion event performance graph types

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
Avg Bytes/Packet	n/a	the average number of bytes included in each packet.	no
ECN Flags Normalized in TCP Traffic/Packet	enable **Explicit Congestion Notification** and select **Packet**	the number of packets for which ECN flags have been cleared on a per-packet basis regardless of negotiation.	yes
ECN Flags Normalized in TCP Traffic/Session	enable **Explicit Congestion Notification** and select **Stream**	the number of times that ECN flags have been cleared on a per-stream basis when ECN use was not negotiated.	yes
Events/Sec	n/a	the number of events per second generated on the device.	no
ICMPv4 Echo Normalizations	enable **Normalize ICMPv4**	the number of ICMPv4 packets for which the 8-bit Code field in Echo (Request) or Echo Reply messages were cleared.	yes
ICMPv6 Echo Normalizations	enable **Normalize ICMPv6**	the number of ICMPv6 packets for which the 8-bit Code field in Echo (Request) or Echo Reply messages was cleared.	yes
IPv4 DF Flag Normalizations	enable **Normalize IPv4** and **Normalize Don't Fragment Bit**	the number of IPv4 packets for which the single-bit Don't Fragment subfield of the IPv4 Flags header field was cleared.	yes

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
IPv4 Options Normalizations	enable **Normalize IPv4**	the number of IPv4 packets for which the option octet was set to 1 (No Operation).	yes
IPv4 Reserved Flag Normalizations	enable **Normalize IPv4** and **Normalize Reserved Bit**	the number of IPv4 packets for which the single-bit Reserved subfield of the IPv4 Flags header field was cleared.	yes
IPv4 Resize Normalizations	enable **Normalize IPv4**	the number of IPv4 packets with excessive-length payload that have been truncated to the datagram length specified in the IP header.	yes
IPv4 TOS Normalizations	enable **Normalize IPv4** and **Normalize TOS Bit**	the number of IPv4 packets for which the one-byte Differentiated Services (DS) field (formerly known as the Type of Service (TOS) field) was cleared.	yes
IPv4 TTL Normalizations	enable **Normalize IPv4** , **Maximum TTL** , and **Reset TTL**	the number of IPv4 Time to Live normalizations.	yes
IPv6 Options Normalizations	enable **Normalize IPv6**	the number of IPv6 packets for which the Option Type field in the Hop-by-Hop Options or Destination Options extension header was set to `00` (Skip and continue processing).	yes
IPv6 TTL Normalizations	enable **Normalize IPv6** , **Minimum TTL** , and **Reset TTL**	the number of IPv6 Hop Limit (TTL) normalizations.	yes
Mbits/Sec	n/a	the number of megabits per second of traffic that passes through the device.	no
Packet Resized to Fit MSS Normalizations	enable **Trim Data to MSS**	the number of packets for which the payload was longer than the TCP Data field, so the payload was trimmed to the Maximum Segment Size.	yes
Packet Resized to Fit TCP Window Normalizations	enable **Trim Data to Window**	the number of packets for which the TCP Data field was trimmed to fit the receiving host's TCP window.	yes
Percent Packets Dropped	n/a	the average percentage of uninspected packets across all selected devices. For example, if you select two devices, then an average of 50% may indicate that one device has a 90% drop rate and the other has a 10% drop rate. It may also indicate that both devices have a drop rate of 50%. The graph only represents the total % drop when you select a single device.	no
RST Packets With Data Stripped Normalizations	enable **Remove Data on RST**	the number of packets for which data was removed from a TCP reset (RST) packet.	yes

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
SYN Packets With Data Stripped Normalizations	enable **Remove Data on SYN**	the number of packets for which data was removed from SYN packets when the TCP operating system was not Mac OS.	yes
TCP Header Padding Normalizations	enable **Normalize/Clear Option Padding Bytes**	the number of TCP packets in which option padding bytes were set to 0.	yes
TCP No Option Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of packets from which the Time Stamp option was stripped.	yes
TCP NS Flag Normalizations	enable **Explicit Congestion Notification** and select **Packet**	the number of ECN Nonce Sum (NS) option normalizations.	yes
TCP Options Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of options (excluding MSS, Window Scale, Time Stamp, and explicitly allowed options) for which the option field is set to No Operation (TCP Option 1).	yes
TCP Packets Blocked By Normalizations	enable **Normalize TCP Payload** (segment reassembly must fail)	the number of packets dropped because the TCP segments could not be properly reassembled.	yes
TCP Reserved Flags Normalizations	enable **Normalize/Clear Reserved Bits**	the number of TCP packets where the Reserved bits have been cleared.	yes
TCP Segment Reassembly Normalizations	enable **Normalize TCP Payload** (segment reassembly must be successful)	the number of packets for which the TCP Data field was normalized to ensure consistency in retransmitted data (any segments that cannot be properly reassembled are dropped).	yes
TCP SYN Option Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of options for which the Maximum Segment Size or Window Scale option was set to No Operation (TCP Option 1) because the SYN control bit was not set.	yes
TCP Timestamp ECR Normalizations	enable **Allow These TCP Options** and set to an option other than `any`	the number of packets for which the Time Stamp Echo Reply (TSecr) option field was cleared because the Acknowledgment (ACK) control bit was not set.	yes
TCP Urgent Pointer Normalizations	enable **Normalize Urgent Pointer**	the number of packets for which the two-byte TCP header Urgent Pointer field was greater than the payload length and was set to the payload length.	yes
Total Blocked Packets	configure **Inline Mode** or **Drop when Inline**	the total number of dropped packets, including rule, decoder, and preprocessor drops.	no
Total Injected Packets	configure **Inline Mode**	the number of packets that were resized before being retransmitted.	no

To generate data for...	You must...	Which represents...	Affected by Inline Mode?
Total TCP Filtered Packets	configure TCP Stream Preprocessing	the number of packets skipped by the stream because of TCP port filtering.	no
Total UDP Filtered Packets	configure UDP Stream Preprocessing	the number of packets skipped by the stream because of UDP port filtering.	no
Urgent Flag Cleared Normalizations	enable **Clear URG if Urgent Pointer is Not Set**	the number of packets for which the TCP header URG control bit was cleared because the urgent pointer was not set.	yes
Urgent Pointer and Urgent Flag Cleared Normalizations	enable **Clear Urgent Pointer/URG on Empty Payload**	the number of packets for which the TCP header Urgent Pointer field and the URG control bit have been cleared because there was no payload.	yes
Urgent Pointer Cleared Normalizations	enable **Clear Urgent Pointer if URG=0**	the number of packets for which the 16-bit TCP header Urgent Pointer field was cleared because the urgent (URG) control bit was not set.	yes

View intrusion event graphs

The system provides graphs that show you intrusion event trends over time. You can generate intrusion event graphs over time ranging from the last hour to the last month, for one or all managed devices.

In a multidomain deployment, you can view data for the current domain and for any descendant domains. You cannot view data from higher level or sibling domains.

Procedure

Step 1 Choose **Overview** > **Summary** > **Intrusion Event Graphs**.

Step 2 Under **Select Device**, choose **all** to include all devices, or choose the specific device you want to include in the graph.

Step 3 Under **Select Graph(s)**, choose the type of graph you want to generate:

- Top 10 Destination Ports
- Top 10 Source IP Addresses
- Top 10 Event Messages

Step 4 Under **Select Time Range**, choose the time range for the graph:

- Last Hour
- Last Day
- Last Week
- Last Month

Step 5 Click **Graph**.

History for intrusion events

This reference provides the historical changes and feature additions for intrusion events across different software versions.

Feature	Minimum Firewall Management Center	Minimum Firewall Threat Defense	Details
IPS Events Datastore Replacement	7.1	Any	• Intrusion incidents, the intrusion event clipboard, and default custom tables (that use the intrusion event columns - **Intrusion Events with Source Criticality** and **Intrusion Events with Destination Criticality**) are deprecated. You can no longer add events to the clipboard using the **Copy** and **Copy All** buttons. Deprecated pages: • **Analysis** > **Intrusions** > **Clipboard** • **Analysis** > **Intrusions** > **Incidents** • Two new fields are added to the main intrusion event table - **Source Host Criticality** and **Destination Host Criticality**. Supported platforms: Secure Firewall Management Center
Unique identifier for connection event in syslogs	6.4.0.4	Any	These syslog fields collectively uniquely identify a connection event and appear in syslogs for intrusion events: DeviceUUID, First Packet Time, Connection Instance ID, and Connection Counter.
IntrusionPolicy field is now included in syslog	6.4	Any	Intrusion event syslogs now specify the intrusion policy that triggered the event.
New intrusion event search field: CVE ID	6.4	Any	You can now search by MITRE's Common Vulnerabilities and Exposures identification number Modified screens: **Analysis** > **Intrusions** > **Events** > **Edit Search** Supported Platforms: All.