

Encrypted Visibility Engine

Encrypted Visibility Engine (EVE) is used to identify client applications and processes utilizing Transport Layer Security (TLS) encryption. It enables visibility and allows administrators to take actions and enforce policy within their environments. The EVE technology can also be used to identify and stop malware.

- Encrypted Visibility Engine, on page 1
- How EVE works, on page 2
- Indications of compromise events, on page 3
- QUIC fingerprinting in EVE, on page 4
- Configure EVE, on page 4
- Configure EVE exception rules, on page 7
- Event enrichment, on page 11
- Use case - Block Traffic Based on the EVE Threat Confidence Score, on page 11

Encrypted Visibility Engine

The encrypted visibility engine (EVE) is used to provide more visibility into the encrypted sessions without the need to decrypt them. These insights into encrypted sessions are obtained by Cisco's open-source library that is packaged in Cisco's vulnerability database (VDB). The library fingerprints and analyzes incoming encrypted sessions and matches it against a set of known fingerprints. This database of known fingerprints is also available in the Cisco VDB.

Key capabilities

Important features of Encrypted Visibility Engine (EVE) include the following:

- Access control policy actions on traffic using information derived from EVE
- Vulnerability Database (VDB) integration with Cisco Secure Firewall for assigning applications to EVE-detected processes with high confidence values
- Custom application detector creation for mapping EVE-detected processes to user-defined applications and overriding built-in process confidence values
- Detection of the operating system type and version of clients that create Client Hello packets in encrypted traffic
- Quick UDP Internet Connections (QUIC) traffic fingerprinting and analysis with server name display in the URL field of the **Connection Events** page

- On any Secure Firewall release with Cisco Vulnerability Database (VDB) 428 or above, EVE supports fingerprinting and analysis of Secure Shell (SSH) and Session Traversal Utilities for NAT (STUN) protocols. Data related to these protocols is displayed in the existing widgets on the Management Center. If the traffic is labeled or randomized, EVE will provide the process name and other relevant EVE metadata fields to help with application identification and security enforcement. Also, AppID maps EVE process names from SSH and STUN sessions to known applications for policy enforcement.

For custom application detector configuration, see the **Configuring Custom Application Detectors** and **Specifying EVE Process Assignments** sections in the **Application Detection** chapter of the Cisco Secure Firewall Management Center Device Configuration Guide.

Note The encrypted visibility engine feature is supported only on Firewall Management Center-managed devices running Snort 3. This feature is not supported on Snort 2 devices and Firewall Device Manager-managed devices.

Attention To use EVE on Firewall Management Center, you must have a valid IPS license on your device. In the absence of a IPS license, the policy displays a warning and deployment is not allowed.

Note

- EVE can detect the operating system type and version of SSL sessions. Normal usage of the operating system, such as running applications and package management software, can trigger OS detection. To view client OS detection, in addition to enabling the EVE toggle button, you must enable **Hosts** under **Policies** > **Network Discovery**. To view a list of possible operating systems on the host IP address, click **Events & Logs** > **Hosts** > **Network Map**, and then choose the required host.
- After enabling EVE for your access control policy, ensure that you have turned on logging for the access control rules within that policy to display the expected results on the EVE dashboard whenever any specific rule conditions are met. For more information on how to turn on logging, see Create and Edit Access Control Rules.
- GRE encapsulation for Mercury detection is supported from Secure Firewall version 10.0.0.

How EVE works

Encrypted Visibility Engine (EVE) provides the ability to identify and control applications without enabling TLS decryption. By using fingerprints of known malicious processes, EVE technology can also be used to identify and block encrypted malicious traffic without outbound decryption.

Summary

The key components involved in how EVE works are:

- Client Hello inspection: EVE examines the initial TLS handshake packet to create client fingerprints
- Fingerprint database: Contains over 5,000 identified client processes mapped to applications

- Machine learning technology: Processes over one billion TLS fingerprints and over 10,000 malware samples daily
- Cisco Vulnerability Database (VDB): Delivers updated fingerprints to customers

Workflow

These stages describe how EVE identifies and processes client applications:

1. EVE inspects the Client Hello portion of the TLS handshake to identify client processes. The Client Hello is the initial data packet that is sent to the server. This gives a good indication of the client process on the host.
2. The system combines the fingerprint with other data such as destination IP address to provide the basis for EVE's application identification. By identifying specific application fingerprints during the TLS session establishment, the system can identify the client process and take appropriate action, such as allowing or blocking it.
3. If a fingerprint is not available in EVE's fingerprint database, it identifies the client application and estimates the threat score of the first flow using the destination details, such as IP address, port, and server name. At this point, the status of the fingerprints is randomized and the status can be viewed in the connection-based debugging messages or log messages. For subsequent flows with the same fingerprint, EVE skips reanalysis and marks the fingerprint status as unlabeled.
4. Through machine learning (ML) technology, Cisco processes over one billion TLS fingerprints and over 10000 malware samples daily to create and update EVE fingerprints. These updates are then delivered to customers using Cisco Vulnerability Database (VDB) package.

Indications of compromise events

An Indication of Compromise (IoC) event is a security detection mechanism that

- identifies connection events with a very high malware confidence level, as reported by EVE,
- triggers for encrypted sessions generated from a host using a malicious client, and
- provides information such as the IP address, MAC address, operating system information of the malicious host, and timestamp of the suspicious activity.

A session with an Encrypted Visibility Threat Confidence score of 'Very High' as seen in connection events generates an IoC event. You must enable **Hosts** from **Policies** > **+ Show more** > **Advanced** > **Network Discovery**. In the Firewall Management Center, view the IoC event existence from here:

- **Events & Logs** > **+ Show more** > **Hosts** > **Indications of Compromise**, and then **Analysis** > **Indications of Compromise**.
- **Events & Logs** > **Hosts** > **Network Map** > Choose the host that must be checked.

  You can view the process information of the session that generated the IoC on the **Connection Events** page. Click **Events & Logs** > **+ Show more** > **Connection** > **Events** to access the **Connection Events** page. Note that you must manually select the Encrypted Visibility fields and IoC field from the **Table View of Connection Events** tab.

The **Security-Related Connection Events** page displays connections that are blocked by EVE, as well as malicious connections with medium, high, and very high threat confidence levels. Click **Events & Logs** > **+ Show more** > **Connection** > **Security-Related Events** to access the **Security-Related Connection Events** page.

- Click **Events & Logs** > **Analysis** > **Unified Events** and choose the Encrypted Visibility fields and IoC field from the column picker option.

QUIC fingerprinting in EVE

QUIC fingerprinting in EVE is a Snort identification mechanism that

- detects applications over QUIC without enabling decryption,
- identifies malware without enabling decryption, and
- detects service applications for access control rule assignment based on the service detected over the QUIC protocol.

Configure EVE

Configure the Encrypted Visibility Engine (EVE) to detect client applications and monitor or block encrypted traffic based on threat confidence levels.

The Encrypted Visibility Engine (EVE) provides visibility into encrypted traffic by analyzing client application behavior and threat indicators. You can operate EVE in Monitor mode for detection only or Protect mode to actively block threats.

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

Step 3 Choose **Encrypted Visibility Engine** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 On the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.

Step 5 Choose the **Monitor** mode or the **Protect** mode.

- Choose the **Monitor** mode to detect client applications and monitor encrypted traffic.
- Choose the **Protect** mode to monitor and block encrypted traffic based on the threat confidence level of the client processes. You can use this mode to monitor and block malicious connections at two threat confidence levels:
  - **High**: Use this level to block connections with threat confidence levels ranging from High to Very High.
  - **Very High**: Use this level to block connections with threat confidence levels that are categorized as Very High.

Step 6 Click **Save** and then deploy the access control policy.

Note To manage exceptions, refer to Configure EVE exception rules, on page 7.

What to do next

Deploy configuration changes.

View Encrypted Visibility Engine events

This task allows you to access and view connection events generated by the Encrypted Visibility Engine (EVE) in the Firewall Management Center.

After enabling the **Encrypted Visibility Engine** and deploying your access control policy, you can start sending live traffic through your system. You can view the logged connection events in the **Unified Events** page.

Perform this procedure to access the connection events in the Firewall Management Center.

Before you begin

To view fingerprints generated from SSH and STUN protocol traffic, your Firewall Threat Defense device must be running version 10.1 or later.

Procedure

Step 1 Click **Events & Logs** > **Analysis** > **Unified Events**.

Additionally, you can use the **Security-Related Connection Events** page (**Events & Logs** > **+ Show more** > **Connection** > **Security-Related Events**) to view the connections that are blocked by Encrypted Visibility Engine, as well as malicious connections with medium, high, and very high threat confidence levels.

The Encrypted Visibility Engine can identify the client process that initiated a connection and the operating system in the client, and indicate if the process contains malware or not.

Step 2 In the **Unified Events** page, explicitly enable these columns that are added for the Encrypted Visibility Engine:

- **Encrypted Visibility Fingerprint**
- **Encrypted Visibility Process Name**
- **Encrypted Visibility Process Confidence Score**
- **Encrypted Visibility Threat Confidence**
- **Encrypted Visibility Threat Confidence Score**
- **Detection Type**
- **EVE Fingerprint**

- **EVE Process Name**
- **EVE Process Confidence Score**
- **EVE Threat Confidence**
- **EVE Threat Confidence Score**
- **Detection Type**

For information about these fields, refer to Connection and Security-Related Connection Event Fields in the *Cisco Secure Firewall Management Center Administration Guide*.

Note On the **Connection Events** page, if processes are assigned applications, the **Detection Type** column displays **Encrypted Visibility Engine**, indicating that the client application was identified by the Encrypted Visibility Engine. Without application assignments to process names, the **Detection Type** column displays **AppID**, indicating that the engine that identified the client application was AppID.

View EVE dashboard

You can use the EVE dashboard to monitor encrypted visibility data, view threat confidence levels, and analyze process identification results across your network.

Before you begin

- In an access control policy, the **Encrypted Visibility Engine (EVE)** must be enabled.
- To view the **Connections with Detected Process Names** and **Malicious Processes** widgets, your device must be running Firewall Threat Defense Version 7.7 or later.
- To view **Malicious Process Responder IPs**, **Malicious Process Contacted Domains**, and **Blocked Connections** data, your device must be running Firewall Threat Defense Version 10.0 or later.
- To view SSH and STUN protocol data in the EVE dashboard widgets, your device must be running Firewall Threat Defense Version 10.1 or later.

Follow these steps to view the EVE dashboards:

Procedure

Step 1 Go to **Insights & Reports** > **Dashboard**.

Step 2 In the **Summary Dashboard** window, click the **Encrypted Visibility Engine** tab.

Step 3 You can view the following dashboards:

- **Discovered Processes**: Displays top client processes used in your network and the connection count. You can click the process name in the table to see the filtered view of the **Connection Events** page, which is filtered by the process name.

- **Threat Confidence**: Displays connections by the confidence levels. You can click the Threat confidence level in the table to see the filtered view of the **Connection Events** page, which is filtered by the confidence level.
- **Connections with Detected Process Names**: Displays total count of connections in which EVE identified the client processes.
- **Malicious Processes**: Displays count of malicious client processes identified by EVE with high and very high threat confidence levels.
- **Malicious Process Responder IPs**: Displays the top destination IP addresses identified by EVE as malicious, categorized with high or very high threat confidence levels. Click on any responder IP address in the widget to navigate to the **Connection Events** page, which is filtered by the selected responder IP address.
- **Malicious Process Contacted Domains**: Displays the count of top domain names identified by EVE as malicious, categorized with high or very high threat confidence levels. Click on any domain name in the widget to navigate to the **Connection Events** page, which is filtered by the selected domain name.
- **Blocked Connections**: Displays the count of connections blocked by EVE.

Note

- If the Management Center is on Secure Firewall version 10.0.0 and there are no devices running version 10.0.0, data in the **Malicious Process Responder IPs**, **Malicious Process Contacted Domains**, and **Blocked Connections** widgets will not be populated.
- If the Management Center is on Secure Firewall version 10.0.0 and there is only one device running version 10.0.0, data in the **Malicious Process Responder IPs**, **Malicious Process Contacted Domains**, and **Blocked Connections** widgets will be populated from that one device only. The other widgets show data from devices running older versions.

Configure EVE exception rules

Configure Encrypted Visibility Engine (EVE) exception rules to ensure the continuity of trusted connections and services by bypassing EVE's block action. This allows trusted networks and processes to be exempted from EVE's block verdict based on the threat confidence level.

You can create an encrypted visibility engine (EVE) exception rule to ensure the continuity of trusted connections and services by bypassing the EVE's block action. You can add attributes such as process names, source and destination IP addresses or FQDNs, and destination dynamic objects to the exception rule. For example, you may want to bypass EVE's block verdict for trusted networks. All the connections in the bypassed networks are exempted from EVE's block verdict based on the threat confidence level.

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control**.

Step 2 Click **Edit** (✎) next to the access control policy you want to edit.

Step 3 Choose **Encrypted Visibility Engine** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 On the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button.

Step 5 Choose the **Protect** mode to monitor and block encrypted traffic based on the threat confidence level of the client processes.

You can use this mode to monitor and block malicious connections at two threat confidence levels:

- **High**: Use this level to block connections with threat confidence levels ranging from High to Very High.
- **Very High**: Use this level to block connections with threat confidence levels that are categorized as Very High.

Step 6 Click **Manage exceptions** to view and add exception rules.

Step 7 On the **Encrypted Visibility Engine (EVE) Exception List** window, click **+Add Exception Rules** and add the required attributes.

Note EVE exception rules can be configured only in the global domain. In the child domain, you can only view EVE exception rule details. You cannot add, edit, or delete EVE exception rules in the child domain.

a) Under the **Process Name** tab, enter an EVE-identified process name, and click **+Add** on the right side of the window.

You can add multiple process names to the same exception rule. The EVE exception list based on process names works only with EVE-identified process names, which are case and space sensitive.

b) Under the **Network Objects** tab, perform one of the following:

- Choose one or more network objects from the **Available Networks** list and add the same to the **Selected Source Network** or **Selected Destination Network** list.
- Under **Selected Source Network** or **Selected Destination Network**, manually enter the IP address, and click the **Add** (+) icon to add it to the list of selected networks.
- To create a new network object, click **+Create Network Object**.
  1. Enter a **Name** and an optional **Description**.
  2. Choose the required network type—**Host, Range, Network, or FQDN**. Enter the relevant IP address if you choose **Host, Range, or Network**. If you choose **FQDN**, enter the fully Qualified Domain Name(**FQDN**) and choose the required option from the **Lookup** drop-down list.
  3. If you want to allow configuration overrides, check the **Allow overrides** checkbox.
  4. Click **Add**.

c) Under the **Dynamic Attributes** tab, choose the required dynamic attributes from the **Available Dynamic Attributes** list and use the > button to add it to the **Selected Destination Dynamic Attributes** list.

For more information about creating dynamic objects or working with dynamic objects, see the *Create Dynamic Objects for the First Time* or *Work With Dynamic Objects* sections in Secure Firewall Management Center Device Configuration Guide.

d) To create a new dynamic attribute, click **+Create Dynamic Attribute.**

1. Enter a **Name** and an optional **Description**.
2. Click **Add**. You can configure this object using Cisco Secure Dynamic Attribute Connector (CSDAC) or Management Center APIs.

e) (Optional) In the **Comment** field available on all the tabs, you can enter a reason for adding the required network objects and dynamic attributes to the EVE exception rule.

Step 8 Click **Save** and then deploy the access control policy.

Note When a connection matches an exception rule, it bypasses the EVE's block verdict. You can view EVE's action in the **Connection Events** or **Unified Events** page. The **Reason** column header displays **EVE Exempted** for identification of such EVE-bypassed traffic.

Add exception rule from unified events

Add exception rules for connections that are blocked by EVE using the information in the Unified Events page.

Use the **Unified Events** page to add exception rules for connections that are blocked by EVE. The Firewall Management Center adds an exception rule to the **Encrypted Visibility Engine (EVE) exception list** object. Note that the exception rules added to this list are applicable for all the access control policies that have EVE enabled.

Before you begin

Exception list is supported only from threat defense version 7.6.0 or later.

Follow these steps to add exception rules for connections that are blocked by EVE:

Procedure

Step 1 Click **Events & Logs** > **Analysis** > **Unified Events**.

Step 2 In the **Reason** column with **Encrypted Visibility Block** as the reason, click the **Ellipsis**( ) icon inside the cell.

Step 3 Choose **Add EVE Exception Rule** from the drop-down list.

Note EVE exception rules can be configured only in the global domain. In the child domain, you can only view EVE exception rule details. You cannot add, edit, or delete EVE exception rules in the child domain.

Step 4 In the **Encrypted Visibility Engine** window that is displayed, the rule is automatically added to the bottom of the exception list. You can review and make changes to the added rule before saving and deploying the configuration.

Upgrade EVE exception rules

On Secure Firewall version 7.7 and earlier, Encrypted Visibility Engine (EVE) exception rules are configured for each policy separately. From Secure Firewall version 10.0.0, the EVE exception list is part of the global domain. As a result, the EVE exception rules are configured in the global domain and applied to all the policies on which EVE is enabled, to block traffic.

EVE exception rule upgrade behavior

When you are upgrading the Management Center from version 7.7 to 10.0.0, all the EVE exception rules from the leaf domains that contain leaf domain network objects, are identified and stored. After the upgrade is complete:

- All EVE exception rules from global domain policies, as well as rules from leaf domain policies that reference global domain objects or inline IP addresses, are consolidated into a single global EVE exception list. Some policies may now include EVE exception rules that were not present before the upgrade as a result.
- All policies that contain EVE exception rules are marked as out-of-date.

If the exception rules from leaf domains contain leaf domain network or dynamic objects, these rules are removed during the upgrade process. The upgrade script log file records all the merged and deleted exception rules and includes the corresponding access control policy and domain for each rule. The log file is located at */var/log/sf/Cisco_Secure_FW_Mgmt_Center_Upgrade10.0.0/800_post/1114_eve_rules.pl.log*.

When you deploy the configuration for the first time after the upgrade, a warning message on the Management Center lists all the deleted EVE exception rules. The warning message also states that there could be possible traffic impact if the rules are not reconfigured in the global EVE exception list. The warning message appears only when you deploy the configuration for the first time after the upgrade is complete.

For Secure Firewall devices running version 7.7 and earlier that are mapped to Management Center running version 10.0.0, only **Very High** threat confidence connection events are sent to the **Security-Related Connection Events** table. For Secure Firewall devices running version 10.0.0, EVE **Blocked** and **Medium+** EVE threat confidence connection events are sent to the **Security-Related Connection Events** table.

Change management support during EVE upgrade:

- When you upgrade the Management Center to version 10.0.0, all active change management tickets that contain access control policies on which EVE is enabled will have their EVE exception rules automatically merged with the global EVE exception list.
- The merging of EVE exception rules with the global EVE exception list occurs regardless of the ticket's approval state. This ensures that no exception rules are lost during the upgrade.

EVE ticket preview generation behavior:

- If a change management ticket contains a policy that is locked and it contains only EVE-related modifications, such as EVE settings or exception rules, the EVE ticket preview will not be automatically regenerated after the upgrade.

- If the ticket contains other policy modifications in addition to EVE-related modifications, the EVE ticket preview will be generated normally.

Event enrichment

Event enrichment is a security analysis feature that

- provides context enrichment for MITRE ATT&CK from the Talos taxonomy and the Encrypted Visibility Engine (EVE),
- communicates both Talos and EVE enrichments using the Talos taxonomy, and
- requires EVE to be enabled for EVE enrichment to work.

For more information about enabling EVE, see Configure EVE, on page 4.

On the **Connection Events** page, you can view these column headers that are added as part of enriched eventing content. You must explicitly enable these columns:

- **MITRE ATT&CK**
- **Other Enrichment**

For information about these fields, see *Connection and Security-Related Connection Event Fields* in *Cisco Secure Firewall Management Center Administration Guide*.

Use case - Block Traffic Based on the EVE Threat Confidence Score

EVE assigns a threat confidence level to encrypted traffic based on its assessment of the likelihood that the associated client process is malicious. Administrators can use these threat confidence levels to monitor traffic or automatically block connections that meet or exceed a configured threshold.

Sample Business Scenario

A large corporate network uses Snort 3 as its primary intrusion detection and prevention system. In a rapidly evolving threat landscape, adoption of robust network security measures is necessary and important. The security team uses EVE to enhance encrypted traffic inspection without the need to implement full man-in-the-middle (MITM) decryption. The EVE technology uses fingerprints of known malicious processes to identify and stop malware.

Prerequisite

- Ensure that you have a valid Intrusion Prevention System (IPS) license and Snort 3 is the detection engine.

High-Level Workflow

1. EVE analyzes the incoming traffic and gives a verdict on the probability of incoming traffic being malware or not.
2. If EVE detects incoming traffic to be malware with a certain level of confidence, you can configure EVE to block that traffic.
3. The packets are first checked for malware probability or threat score, and the threat score is compared with the block threshold that you have set.
4. If the threat level is higher than the configured threshold, EVE blocks the traffic.
5. If the threat level is lower than the configured threshold, EVE takes no action.

Configure Block Thresholds in EVE

This procedure shows how to block potentially malicious traffic, based on the EVE threat confidence level.

Procedure

Step 1 Choose **Policies** > **Security policies** > **Access Control**.

Step 2 Click the **Edit** (✎) next to the access control policy you want to edit.

Step 3 Choose **Encrypted Visibility Engine** from the **More** drop-down arrow at the end of the packet flow line.

Step 4 In the **Encrypted Visibility Engine** page, enable the **Encrypted Visibility Engine (EVE)** toggle button..

Step 5 Choose the **Protect** mode to monitor and block encrypted traffic based on the EVE threat confidence levels.

Step 6 Choose **Very High** to block connections with threat confidence levels that are categorized as Very High.

Step 7 Click **Save**.

What to do next

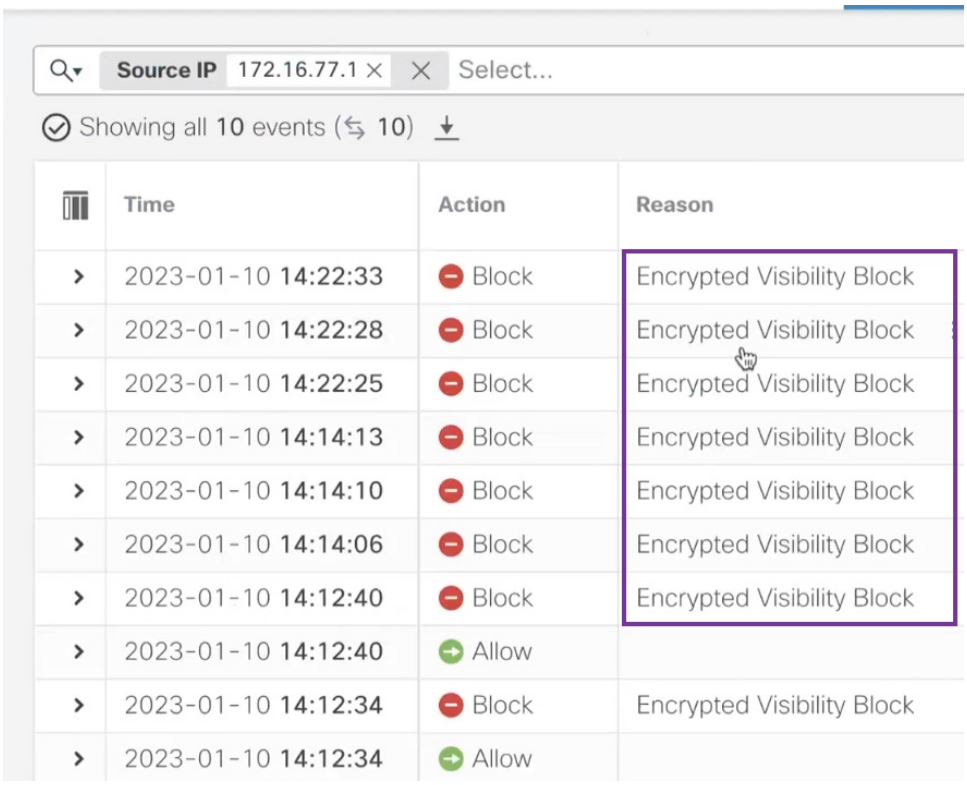
Deploy configuration changes.

View EVE Events

Procedure

Step 1 To verify the block action, choose **Events & Logs > Analysis > Unified Events**.

Step 2 If you have configured EVE to block traffic, the **Reason** field shows **Encrypted Visibility Block**.



Step 3 The following is an example of the **EVE Process Name** as **test_malware**, **EVE Threat Confidence** as **Very High**, and **EVE Threat Confidence Score** as **90** percent.

