



Discovery Events

The following topics describe how to work with discovery events:

- [Requirements and Prerequisites for Discovery Events, on page 1](#)
- [Discovery and Identity Data in Discovery Events, on page 1](#)
- [Viewing Discovery Event Statistics, on page 2](#)
- [Viewing Discovery Performance Graphs, on page 5](#)
- [Using Discovery and Identity Workflows, on page 6](#)
- [History for Working with Discovery Events, on page 56](#)

Requirements and Prerequisites for Discovery Events

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Security Analyst

Discovery and Identity Data in Discovery Events

The system generates tables of events that represent the changes detected in your monitored network. You can use these tables to review the user activity on your network and determine how to respond. The *network discovery* and *identity* policies specify the kinds of data you want to collect, the network segments you want to monitor, and the specific hardware interfaces you want to use to do it.

You can use discovery and identity event tables to identify threats associated with hosts, applications, and users on your network. The system provides a set of predefined workflows that you can use to analyze the

events that your system generates. You can also create custom workflows that display only the information that matches your specific needs.

To collect and store network discovery and identity data for analysis, you must configure network discovery and identity policies. After you configure an identity policy, you must invoke it in your access control policy and deploy it to the devices you want to use to monitor traffic.

Your network discovery policy provides host, application, and non-authoritative user data. Your identity policy provides authoritative user data.

The following discovery event tables are located under the Analysis > Hosts and Analysis > Users menus.

Discovery Event Table	Populated With Discovery Data?	Populated With Identity Data?
Hosts	Yes	No
Host Indications of Compromise	Yes	No
Applications	Yes	No
Application Details	Yes	No
Servers	Yes	No
Host Attributes	Yes	No
Discovery Events	Yes	Yes
User Indications of Compromise	Yes	Yes
Active Sessions	Yes	Yes
User Activity	Yes	Yes
Users	Yes	Yes
Vulnerabilities	Yes	No
Third-Party Vulnerabilities	Yes	No

Viewing Discovery Event Statistics

The Discovery Statistics page displays a summary of the hosts, events, protocols, application protocols, and operating systems detected by the system.

The page lists statistics for the last hour and the total accumulated statistics. You can choose to view statistics for a particular device, or all devices. You can also view events that match the entries on the page by clicking the event, server, operating system, or operating system vendor listed within the summary.

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 > Discovery Statistics**.
- Step 2** From the **Select Device** list, choose the device whose statistics you want to view. Optionally, choose **All** to view statistics for all devices managed by the management center.
- Step 3** You have the following options:
- In the Statistics Summary, view general statistics as described in [The Statistics Summary Section, on page 3](#).
 - In the Event Breakdown, click the type of event you want to view. If no events appear, you may need to adjust the time range as described in [Changing the Time Window](#).
 - In the Protocol Breakdown, view the protocols currently in use by detected hosts.
 - In the Application Protocol Breakdown, click the name of the application protocol you want to view.
 - In the OS Breakdown, click the **OS Name** or **OS Vendor**.

Related Topics

- [The Event Breakdown Section, on page 4](#)
- [The Protocol Breakdown Section, on page 4](#)
- [The Application Protocol Breakdown Section, on page 5](#)
- [The OS Breakdown Section, on page 5](#)

The Statistics Summary Section

Descriptions of the rows of the Statistics Summary section follow.

Total Events

Total number of discovery events stored on the management center.

Total Events Last Hour

Total number of discovery events generated in the last hour.

Total Events Last Day

Total number of discovery events generated in the last day.

Total Application Protocols

Total number of application protocols from servers running on detected hosts.

Total IP Hosts

Total number of detected hosts identified by unique IP address.

Total MAC Hosts

Total number of detected hosts not identified by IP address.

Note that the Total MAC Hosts statistic remains the same whether you are viewing discovery statistics for all devices or for a specific device. This is so because managed devices discover hosts based on their IP addresses. This statistic gives the total of all hosts that are identified by other means and is independent of a given managed device.

Total Routers

Total number of detected nodes identified as routers.

Total Bridges

Total number of detected nodes identified as bridges.

Host Limit Usage

Total percentage of the host limit currently in use. The host limit is defined by the model of your management center. Note that the host limit usage only appears if you are viewing statistics for all managed devices.



Note If the host limit is reached and a host is deleted, the host will not reappear on the network map you purge discovery data.

Last Event Received

The date and time that the most recent discovery event occurred.

Last Connection Received

The date and time that the most recent connection was completed.

The Event Breakdown Section

The Event Breakdown section lists a count of each type of discovery event and host input event that occurred within the last hour, as well as a count of the total number of each event type stored in the database.

You can also use the Event Breakdown section to view details on discovery and host input events.

Related Topics

[Discovery and Host Input Events](#), on page 8

The Protocol Breakdown Section

The Protocol Breakdown section lists the protocols currently in use by detected hosts. It displays each detected protocol name, its “layer” in the protocol stack, and the total number of hosts that communicate using the protocol.

The Application Protocol Breakdown Section

The Application Protocol Breakdown section lists the application protocols that are currently in use by detected hosts. It lists the protocol name, the total number of hosts running the application protocol in the past hour, and the total number of hosts that have been detected running the protocol at any point.

You can also use the Application Protocol Breakdown section to view details on servers using the detected protocols.

Related Topics

[Server Data](#), on page 27

The OS Breakdown Section

The OS Breakdown section lists the operating systems currently running on the monitored network, along with their vendors and the total number of hosts running each operating system.

A value of `unknown` for the operating system name or version means that the operating system or its version does not match any of the system's fingerprints. A value of `pending` means that the system has not yet gathered enough information to identify the operating system or its version.

You can use the OS Breakdown section to view details on the detected operating systems.

Related Topics

[Host Data](#), on page 15

Viewing Discovery Performance Graphs

You can generate graphs that display performance statistics for managed devices with discovery events.

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.

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.

Before you begin

Edit the applicable network discovery policy to include applications, hosts, and users. (This may impact system performance.) See [Configuring Network Discovery Rules](#) and [Actions and Discovered Assets](#).

You must be an Admin or Maintenance user to perform this task.

Procedure

- Step 1** Choose **Overview > Summary > Discovery Performance**.
- Step 2** From the **Select Device** list, choose the management center or managed devices you want to include.
- Step 3** From the **Select Graph(s)** list, choose the type of graph you want to create as described in [Discovery Performance Graph Types, on page 6](#).
- Step 4** From the **Select Time Range** list, choose the time range you would like to use for the graph.

Step 5 Click **Graph** to graph the selected statistics.

Discovery Performance Graph Types

Descriptions of the available graph types follow.

Processed Events/Sec

Displays a graph that represents the number of events that the Data Correlator processes per second

Processed Connections/Sec

Displays a graph that represents the number of connections that the Data Correlator processes per second

Generated Events/Sec

Displays a graph that represents the number of events that the system generates per second

Mbits/Sec

Displays a graph that represents the number of megabits of traffic that are analyzed by the discovery process per second

Avg Bytes/Packet

Displays a graph that represents the average number of bytes included in each packet analyzed by the discovery process

K Packets/Sec

Displays a graph that represents the number of packets analyzed by the discovery process per second, in thousands

Using Discovery and Identity Workflows

The management center provides a set of event workflows that you can use to analyze the discovery and identity data that is generated for your network. The workflows are, along with the network map, a key source of information about your network assets.

The management center provides predefined workflows for discovery and identity data, detected hosts and their host attributes, servers, applications, application details, vulnerabilities, user activities, and users. You can also create custom workflows.

Procedure

Step 1 To access a predefined workflow:

- Discovery and Host Input Data — See [Viewing Discovery and Host Input Events](#), on page 14.

- Host Data — See [Viewing Host Data](#), on page 16.
- Host Attributes Data — See [Viewing Host Attributes](#), on page 21.
- Host or User Indications of Compromise Data — See [View and Work with Indications of Compromise Data](#), on page 24.
- Server Data — See [Viewing Server Data](#), on page 28.
- Application Data — See [Viewing Application Data](#), on page 31.
- Application Detail Data — See [Viewing Application Detail Data](#), on page 33.
- Active Session Data — See [Viewing Active Session Data](#), on page 48.
- User Data — See [Viewing User Data](#), on page 51.
- User Activity Data — See [Viewing User Activity Data](#), on page 53.
- Network Map — See [Viewing Network Maps](#).

Step 2 To access a custom workflow, choose **Analysis > Advanced > Custom Workflows**.

Step 3 To access a workflow based on a custom table, choose **Analysis > Advanced > Custom Tables**.

Step 4 Perform any of the following actions, which are common to all of the pages accessed in the network discovery workflows:

- Constrain Columns — To constrain the columns that display, click **Close** (✕) in the column heading that you want to hide. In the pop-up window that appears, click **Apply**.

Tip To hide or show other columns, check 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.

- Delete — To delete some or all items in the current constrained view, check the check boxes next to items you want to delete and click **Delete**, or click **Delete All**. These items remain deleted until the system's discovery function is restarted, when they may be detected again.

Caution Before you delete a non-VPN session on the **Analysis > Users > Active Sessions** page, verify that the session is actually closed. After you delete the active session, an applicable policy will not be able to detect the session on the device, and therefore the session will not be monitored or blocked even if the policy was configured to perform those actions.

Note For more information about VPN sessions on the **Analysis > Users > Active Sessions** page, see [Viewing Remote Access VPN Current Users](#).

Note You **cannot** delete Cisco (as opposed to third-party) vulnerabilities; you can, however, mark them reviewed.

- Drill Down — To drill down to the next page in the workflow, see [Using Drill-Down Pages](#).
- Navigate Current Page — To navigate within the current workflow page, see [Workflow Page Navigation Tools](#).
- Navigate within a Workflow — 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.

- Navigate to Other Workflows — To navigate to other event views to examine associated events, see [Inter-Workflow Navigation](#).
- Sort Data — To sort data in a workflow, click the column title. Click the column title again to reverse the sort order.
- View Host Profile — To view the host profile for an IP address, click **Host Profile** or, for hosts with active indications of compromise (IOC) tags, the **Compromised Host** that appears next to the IP address.
- View User Profile — To view user identity information, click the user icon that appears next to the **User Identity**, or for users associated with IOCs, **Red User**.

Related Topics

[Using Workflows](#)

[Purging Data from the Management Center Database](#)

Discovery and Host Input Events

The system generates discovery events that communicate the details of changes in your monitored network segments. *New* events are generated for newly discovered network features, and change events are generated for any change in previously identified network assets.

During its initial network discovery phase, the system generates new events for each host and any TCP or UDP servers discovered running on each host. Optionally, you can configure the system to use exported NetFlow records to generate these new host and server events.

In addition, the system generates new events for each network, transport, and application protocol running on each discovered host. You can disable detection of application protocols in discovery rules configured to monitor NetFlow exporters, but not in discovery rules configured to monitor managed devices. If you enable host or user discovery in non-NetFlow discovery rules, applications are automatically discovered.

After the initial network mapping is complete, the system continuously records network changes by generating change events. Change events are generated whenever the configuration of a previously discovered asset changes.

When a discovery event is generated, it is logged to the database. You can use the management center web interface to view, search, and delete discovery events. You can also use discovery events in correlation rules. Based on the type of discovery event generated as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and syslog, SNMP, and email alert responses when network traffic meets your criteria.

You can add data to the network map using the host input feature. You can add, modify, or delete operating system information, which causes the system to stop updating that information for that host. You can also manually add, modify, or delete application protocols, clients, servers, and host attributes or modify vulnerability information. When you do this, the system generates host input events.

Discovery Event Types

You can configure the types of discovery events the system logs in your network discovery policy. When you view the discovery events table, the event type is listed in the **Event** column. Descriptions of the discovery event types follow.

Additional MAC Detected for Host

This event is generated when the system detects a new MAC address for a previously discovered host.

This event is often generated when the system detects hosts passing traffic through a router. While each host has a different IP address, they all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an “ARP/DHCP detected” message within the event description in the event view.

Client Timeout

This event is generated when the system drops a client from the database due to inactivity.

Client Update

This event is generated when the system detects a payload (that is, a specific type of content, such as audio, video, or webmail) in HTTP traffic.

DHCP: IP Address Changed

This event is generated when the system detects that a host IP address has changed due to DHCP address assignment.

DHCP: IP Address Reassigned

This event is generated when a host is reusing an IP address; that is, when a host obtains an IP address formerly used by another physical host due to DHCP IP address assignment.

Hops Change

This event is generated when the system detects a change in the number of network hops between a host and the device that detects the host. This may happen if:

- The device sees host traffic through different routers and is able to make a better determination of the host’s location.
- The device detects an ARP transmission from the host, indicating that the host is on a local segment.

Host Deleted: Host Limit Reached

This event is generated when the host limit on the management center is exceeded and a monitored host is deleted from the network map.

Host Dropped: Host Limit Reached

This event is generated when the host limit on the management center is reached and a new host is dropped. Compare this with the previous event where old hosts are deleted from the network map when the host limit is reached.

To drop new hosts when the host limit is reached, go to **Policies > Network Discovery > Advanced** and set **When Host Limit Reached** to **Drop hosts**.

Host IOC Set

This event is generated when an IOC (Indications of Compromise) is set for a host and generates an alert.

Host Timeout

This event is generated when a host is dropped from the network map because the host has not produced traffic within the interval defined in the network discovery policy. Note that individual host IP addresses and MAC addresses time out individually; a host does not disappear from the network map unless all of its associated addresses have timed out.

If you change the networks you want to monitor in your network discovery policy, you may want to manually delete old hosts from the network map so that they do not count against your host limit.

Host Type Changed to Network Device

This event is generated when the system detects that a detected host is actually a network device.

Identity Conflict

This event is generated when the system detects a new server or operating system identity that conflicts with a current active identity for that server or operating system.

If you want to resolve identity conflicts by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

Identity Timeout

This event is generated when server or operating system identity data from an active source times out.

If you want to refresh identity data by rescanning the host to obtain newer active identity data, you can use Identity Conflict events to trigger an Nmap remediation.

MAC Information Change

This event is generated when the system detects a change in the information associated with a specific MAC address or TTL value.

This event often occurs when the system detects hosts passing traffic through a router. While each host has a different IP address, they will all appear to have the MAC address associated with the router. When the system detects the actual MAC address associated with the IP address, it displays the MAC address in bold text within the host profile and displays an “ARP/DHCP detected” message within the event description in the event view. The TTL may change because the traffic may pass through different routers or if the system detects the actual MAC address of the host.

NETBIOS Name Change

This event is generated when the system detects a change to a host's NetBIOS name. This event will only be generated for hosts using the NetBIOS protocol.

New Client

This event is generated when the system detects a new client.



Note To collect and store client data for analysis, make sure that you enable application detection in your discovery rules in the network discovery policy.

New Host

This event is generated when the system detects a new host running on the network.

This event can also be generated when a device processes NetFlow data that involves a new host. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover hosts.

New Network Protocol

This event is generated when the system detects that a host is communicating with a new network protocol (IP, ARP, and so on).

New OS

This event is generated when the system either detects a new operating system for a host, or a change in a host's operating system.

New TCP Port

This event is generated when the system detects a new TCP server port (for example, a port used by SMTP or web services) active on a host. This event is not used to identify the application protocol or the server associated with it; that information is transmitted in the TCP Server Information Update event.

This event can also be generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover applications.

New Transport Protocol

This event is generated when the system detects that a host is communicating with a new transport protocol, such as TCP or UDP.

New UDP Port

This event is generated when the system detects a new UDP server port running on a host.

This event can also be generated when a device processes NetFlow data involving a server on your monitored networks that does not already exist in the network map. To generate an event in this case, configure the network discovery rule that manages NetFlow data to discover applications.

TCP Port Closed

This event is generated when the system detects that a TCP port has closed on a host.

TCP Port Timeout

This event is generated when the system has not detected activity from a TCP port within the interval defined in the system's network discovery policy.

TCP Server Information Update

This event is generated when the system detects a change in a discovered TCP server running on a host.

This event may be generated if a TCP server is upgraded.

UDP Port Closed

This event is generated when the system detects that a UDP port has closed on a host.

UDP Port Timeout

This event is generated when the system has not detected activity from a UDP port within the interval defined in the network discovery policy.

UDP Server Information Update

This event is generated when the system detects a change in a discovered UDP server running on a host.

This event may be generated if a UDP server is upgraded.

VLAN Tag Information Update

This event is generated when the system detects a change in the VLAN tag attributed to a host.

Related Topics

[Host Input Event Types](#), on page 12

Host Input Event Types

When you view a table of discovery events, the event type is listed in the **Event** column.

Contrast host input events, which are generated when a user takes a specific action (such as manually adding a host), with discovery events, which are generated when the system itself detects a change in your monitored network (such as detecting traffic from a previously undetected host).

You can configure the types of host input events that the system logs by modifying your network discovery policy.

If you understand the information the different types of host input events provide, you can more effectively determine which events you want to log and alert on, and how to use these alerts in correlation policies. In addition, knowing the names of the event types can help you craft more effective event searches. Descriptions of the different types of host input events follow.

Add Client

This event is generated when a user adds a client.

Add Host

This event is generated when a user adds a host.

Add Protocol

This event is generated when a user adds a protocol.

Add Scan Result

This event is generated when the system adds the results of an Nmap scan to a host.

Add Port

This event is generated when a user adds a server port.

Delete Client

This event is generated when a user deletes a client from the system.

Delete Host/Network

This event is generated when a user deletes an IP address or subnet from the system.

Delete Protocol

This event is generated when a user deletes a protocol from the system.

Delete Port

This event is generated when a user deletes a server port or group of server ports from the system.

Host Attribute Add

This event is generated when a user creates a new host attribute.

Host Attribute Delete

This event is generated when a user deletes a user-defined host attribute.

Host Attribute Delete Value

This event is generated when a user deletes a value assigned to a host attribute.

Host Attribute Set Value

This event is generated when a user sets a host attribute value for a host.

Host Attribute Update

This event is generated when a user changes the definition of a user-defined host attribute.

Set Host Criticality

This event is generated when a user sets or modifies the host criticality value for a host.

Set Operating System Definition

This event is generated when a user sets the operating system for a host.

Set Server Definition

This event is generated when a user sets the vendor and version definitions for a server.

Set Vulnerability Impact Qualification

This event is generated when a vulnerability impact qualification is set.

When a vulnerability is disabled at a global level from being used for impact qualifications, or when a vulnerability is enabled at a global level, this event is generated.

Vulnerability Set Invalid

This event is generated when a user invalidates (or reviews) a vulnerability or vulnerabilities.

Vulnerability Set Valid

This event is generated when a user validates a vulnerability that was previously marked as invalid.

Related Topics

[Discovery Event Types](#), on page 8

Viewing Discovery and Host Input Events

Discovery events workflows allow you to view data from both discovery events and host input events. You can manipulate the event view depending on the information you are looking for.

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.

The page you see when you access events differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of discovery events and a terminating host view page. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1 Choose **Analysis > Hosts > Discovery Events**.

Step 2 You have the following options:

- Adjust the time range as described in [Changing the Time Window](#).

Note Events that were 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.

- Use a different workflow, including a custom workflow, by clicking (**switch workflow**).
- Perform basic workflow actions; see [Using Discovery and Identity Workflows](#), on page 6.
- Learn more about the contents of the columns in the table; see [Discovery Event Fields](#), on page 14.

Related Topics

[Using Discovery and Identity Workflows](#), on page 6

Discovery Event Fields

Descriptions of the fields that can be viewed and searched in the discovery events table follow.

Time

The time that the system generated the event.

Event

The discovery event type or host input event type.

IP Address

The IP address associated with the host involved in the event.

User

The last user to log into the host involved in the event before the event was generated. If only non-authoritative users log in after an authoritative user, the authoritative user remains the current user for the host unless another authoritative user logs in.

MAC Address

The MAC address of the NIC used by the network traffic that triggered the discovery event. This MAC address can be either the actual MAC address of the host involved in the event, or the MAC address of a network device that the traffic passed through.

MAC Vendor

The MAC hardware vendor of the NIC used by the network traffic that triggered the discovery event.

Port

The port used by the traffic that triggered the event, if applicable.

Description

The text description of the event.

Domain

The domain of the device that discovered the host. This field is only present if you have ever configured the management center for multitenancy.

Device

The name of the managed device that generated the event. For new host and new server events based on NetFlow data, this is the managed device that processed the data.

Related Topics

[Event Searches](#)

Host Data

The system generates an event when it detects a host and collects information about it to build the host profile. You can use the management center web interface to view, search, and delete hosts.

While viewing hosts, you can create traffic profiles and compliance allow lists based on selected hosts. You can also assign host attributes, including host criticality values (which designate business criticality) to groups of hosts. You can then use these criticality values, allow lists, and traffic profiles within correlation rules and policies.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see [Differences between NetFlow and Managed Device Data](#).

Viewing Host Data

You can use the management center to view a table of hosts that the system has detected. Then, you can manipulate the view depending on the information you are looking for.

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.

The page you see when you access hosts differs depending on the workflow you use. Both predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1

Access the host data:

- If you are using the predefined workflow, choose **Analysis > Hosts > Hosts**.
- If you are using a custom workflow that does not include the table view of hosts, click **(switch workflow)**, then choose **Hosts**.

Step 2

You have the following options:

- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [Host Data Fields, on page 16](#).
 - Right-click an item in the table to see options. (Not every column offers options.)
 - Assign a host attribute to specific hosts; see [Setting Host Attributes for Selected Hosts, on page 23](#).
 - Create traffic profiles for specific hosts, see [Creating a Traffic Profile for Selected Hosts, on page 20](#).
 - Create a compliance allow list based on specific hosts, see [Creating a Compliance Allow List Based on Selected Hosts, on page 21](#).
-

Host Data Fields

When the system discovers a host, it collects data about that host. That data can include the host's IP addresses, the operating system it is running, and more. You can view some of that information in the table view of hosts.

Descriptions of the fields that can be viewed and searched in the hosts table follow below.

Last Seen

The date and time any of the host's IP addresses was last detected by the system. The Last Seen value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system generates a new host event for any of the host's IP addresses.

For hosts with operating system data updated using the host input feature, the Last Seen value indicates the date and time when the data was originally added.

IP Address

The IP addresses associated with the host.

MAC Address

The host's detected MAC address of the NIC.

The MAC Address field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Address field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

MAC Vendor

The host's detected MAC hardware vendor of the NIC.

The MAC Vendor field appears in the Table View of Hosts, which you can find in the Hosts workflow. You can also add the MAC Vendor field to:

- custom tables that include fields from the Hosts table
- drill-down pages in custom workflows based on the Hosts table

When searching this field, enter `virtual_mac_vendor` to match events that involve virtual hosts.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-specified criticality value assigned to the host.

NetBIOS Name

The NetBIOS name of the host. Only hosts running the NetBIOS protocol will have a NetBIOS name.

VLAN ID

VLAN ID used by the host.

Hops

The number of network hops from the device that detected the host to the host.

Host Type

The type of host. Can be any of the following: host, mobile device, jailbroken mobile device, router, bridge, NAT device, and load balancer.

The methods the system uses to distinguish network devices include:

- the analysis of Cisco Discovery Protocol (CDP) messages, which can identify network devices and their type (Cisco devices only)
- the detection of the Spanning Tree Protocol (STP), which identifies a device as a switch or bridge
- the detection of multiple hosts using the same MAC address, which identifies the MAC address as belonging to a router
- the detection of TTL value changes from the client side, or TTL values that change more frequently than a typical boot time, which identify NAT devices and load balancers

If a device is not identified as a network device, it is categorized as a host.

When searching this field, enter `!host` to search for all network devices.

Hardware

The hardware platform for a mobile device.

OS

One of the following:

- The operating system (name, vendor, and version) either detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple identities, it displays those identities in a comma-separated list.

This field appears when you invoke the hosts event view from the Custom Analysis widget on the dashboard. It is also a field option in custom tables based on the Hosts table.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Conflict

This field is search only.

OS Vendor

One of the following:

- The vendor of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple vendors, it displays those vendors in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Name

One of the following:

- The operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple names, it displays those names in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

OS Version

One of the following:

- The version of the operating system detected on the host or updated using Nmap or the host input feature
- `unknown` if the operating system does not match any known fingerprint
- `pending` if the system has not yet gathered enough information to identify the operating system

If the system detects multiple versions, it displays those versions in a comma-separated list.

When searching this field, enter `n/a` to include hosts where the operating system has not yet been identified.

Source Type

The type of source used to establish the host's operating system identity:

- User: `user_name`
- Application: `app_name`
- Scanner: `scanner_type` (Nmap or scanner added through network discovery configuration)
- `Firepower` for operating systems detected by the system

The system may reconcile data from multiple sources to determine the identity of an operating system.

Confidence

One of the following:

- the percentage of confidence that the system has in the identity of the operating system running on the host, for hosts detected by the system
- 100%, for operating systems identified by an active source, such as the host input feature or Nmap scanner
- `unknown`, for hosts for which the system cannot determine an operating system identity, and for hosts added to the network map based on NetFlow data

When searching this field, enter `n/a` to include hosts added to the network map based on NetFlow data.

Notes

The user-defined content of the Notes host attribute.

Domain

The domain associated with the host. This field is only present if you have ever configured the management center for multitenancy.

Device

Either the managed device that detected the traffic or the device that processed NetFlow or host input data.

If this field is blank, either of the following conditions is true:

- The host was added to the network map by a device that is not explicitly monitoring the network where the host resides, as defined in the network discovery policy.
- The host was added using the host input feature and has not also been detected by the system.

Count

The number of events that match the information that appears in each row. This field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

[Event Searches](#)

[Operating System Identity Conflicts](#)

Creating a Traffic Profile for Selected Hosts

A traffic profile is a profile of the traffic on your network, based on connection data collected over a timespan that you specify. After you create a traffic profile, you can detect abnormal network traffic by evaluating new traffic against your profile, which presumably represents normal network traffic.

You can use the Hosts page to create a traffic profile for a group of hosts that you specify. The traffic profile will be based on connections detected where one of the hosts you specify is the initiating host. Use the sort and search features to isolate the hosts for which you want to create a profile.

Before you begin

You must be an Admin user to perform this task.

Procedure

-
- Step 1** On a table view in the hosts workflow, check the check boxes next to the hosts for which you want to create a traffic profile.
 - Step 2** At the bottom of the page, click **Create Traffic Profile**.
 - Step 3** Modify and save the traffic profile according to your specific needs.

Related Topics

[Introduction to Traffic Profiles](#)

Creating a Compliance Allow List Based on Selected Hosts

Compliance allow lists allow you to specify which operating systems, clients, and network, transport, or application protocols are allowed on your network.

You can use the Hosts page to create a compliance allow list based on the host profiles of a group of hosts that you specify. Use the sort and search features to isolate the hosts that you want to use to create an allow list.

Before you begin

You must be an Admin user to perform this task.

Procedure

-
- Step 1** On a table view in the hosts workflow, check the check boxes next to the hosts for which you want to create an allow list.
 - Step 2** At the bottom of the page, click **CreateAllow List**.
 - Step 3** Modify and save the allow list according to your specific needs.

Related Topics

[Introduction to Compliance Allow Lists](#)

Host Attribute Data

The system collects information about the hosts it detects and uses that information to build host profiles. However, there may be additional information about the hosts on your network that you want to provide to your analysts. You can add notes to a host profile, set the business criticality of a host, or provide any other information that you choose. Each piece of information is called a *host attribute*.

You can use host attributes in host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also set attribute values in response to a correlation rule.

Related Topics

[Viewing Host Attributes](#), on page 21

[Configuring Set Attribute Remediations](#)

Viewing Host Attributes

You can use the management center to view a table of hosts detected by the system, along with their host attributes. Then, you can manipulate the view depending on the information you are looking for.

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.

The page you see when you access host attributes differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of host attributes that lists all detected hosts and their attributes, and terminates in a host view page, which contains a host profile for every host that meets your constraints.

You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

- Step 1** Access the host attributes data:
- If you are using the predefined workflow, choose **Analysis > Hosts > Host Attributes**.
 - If you are using a custom workflow that does not include the table view of host attributes, click (**switch workflow**), then choose **Attributes**.
- Step 2** You have the following options:
- Use a different workflow, including a custom workflow, by clicking (**switch workflow**).
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [Host Attribute Data Fields, on page 22](#).
 - Assign a host attribute to specific hosts; see [Setting Host Attributes for Selected Hosts, on page 23](#).
-

Host Attribute Data Fields

Note that the host attributes table does not display hosts identified only by MAC addresses.

Descriptions of the fields that can be viewed and searched in the host attributes table follow.

IP Address

The IP addresses associated with a host.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Host Criticality

The user-assigned importance of a host to your enterprise. You can use the host criticality in correlation rules and policies to tailor policy violations and their responses to the importance of a host involved in an event. You can assign a host criticality of low, medium, high, or none.

Notes

Information about the host that you want other analysts to view.

Any user-defined host attribute, including those for compliance allow lists

The value of the user-defined host attribute. The host attributes table contains a field for each user-defined host attribute.

Domain

The domain associated with the host. This field is only present if you have ever configured the management center for multitenancy.

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.

Related Topics

[Event Searches](#)

Setting Host Attributes for Selected Hosts

You can configure predefined and user-defined host attributes from a host workflow.

Procedure

-
- Step 1** In a host workflow, check the check boxes next to the hosts to which you want to add a host attribute.
- Tip** Use the sort and search features to isolate the hosts to which you want to assign particular attributes.
- Step 2** At the bottom of the page, click **Set Attributes**.
- Step 3** Optionally, set the host criticality for the hosts you selected. You can choose **None**, **Low**, **Medium**, or **High**.
- Step 4** Optionally, add notes to the host profiles of the hosts you selected in the text box.
- Step 5** Optionally, set any user-defined host attributes you have configured.
- Step 6** Click **Save**.
-

Indications of Compromise Data

The system correlates various types of data (intrusion events, Security Intelligence, connection events, and file or malware events) to determine whether a host on your monitored network is likely to be compromised by malicious means. Certain combinations and frequencies of event data trigger indications of compromise (IOC) tags on affected hosts. The IP addresses of these hosts appear in event views with a **Red Compromised Host icon**.

When a host is identified as potentially compromised, the user associated with that compromise is also tagged. These users appear in event views with a **Red User icon**.

If a file containing malware is seen again within 300 seconds of being tagged as an IOC, another IOC is not generated. If the same file is seen more than 300 seconds later, a new IOC will be generated.

To configure the system to tag events as indications of compromise, see *Enabling Indications of Compromise Rules* in the [Cisco Secure Firewall Management Center Device Configuration Guide](#).

Related Topics

[Editing Server Identities](#)

View and Work with Indications of Compromise Data

You can use the management center to view tables showing Indications of Compromise (IOC). Manipulate the event view depending on the information you are looking for.

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.

The page you see depends on the workflow you use. The predefined IOC workflows terminate in a profile view, which contains a host or user profile for every host or user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Before you begin

- For your system to detect and tag indications of compromise (IOC), you must activate the IOC feature in the network discovery policy and enable at least one IOC rule. See *Enabling Indications of Compromise Rules* in the [Cisco Secure Firewall Management Center Device Configuration Guide](#).
- Users must be identified in an active Identity policy.

Procedure

Step 1

Determine which location in the web interface presents information that meets your needs.

You can use the following locations to view or work with Indication of Compromise data:

- Event Viewer (under the Analysis menu) — Connection, Security Intelligence, intrusion, malware, and IOC discovery event views indicate whether an event triggered an IOC. Note that malware events generated by Secure Endpoint that trigger IOC rules have the event type `AMP IOC` and appear with an event subtype that specifies the compromise.
- Dashboard — In the dashboard, Threats of the Summary Dashboard displays, by default, IOC tags by host and by user. The Custom Analysis widget offers presets based on IOC data.
- Context Explorer — The Indications of Compromise section of the Context Explorer displays graphs of hosts by IOC category and IOC categories by host.
- Network Map page — The Indications of Compromise under Analysis > Hosts > Network Map groups potentially compromised hosts on your network by type of compromise and IP address.
- Network File Trajectory details page — The details pages for files listed under Analysis > Files > Network File Trajectory let you track indications of compromise on your network.
- Host Indications of Compromise page — The Host Indications of Compromise page under the Analysis > Hosts menu lists monitored hosts, grouped by IOC tag. Use the workflows on this page to drill down into your data.
- User Indications of Compromise page — The User Indications of Compromise page under the Analysis > Users menu lists users associated with potential IOC events, grouped by IOC tag. Use the workflows on this page to drill down into your data.
- Host Profile page — The host profile for a potentially compromised host displays all IOC tags associated with that host, and lets you resolve IOC tags and configure IOC rule states.

- **User Profile page** — The user profile for a user associated with a potential IOC event displays all IOC tags associated with that user, and lets you resolve IOC tags and configure IOC rule states. (The user profile is labeled "User Identity" in the management center web interface.)

Step 2 If applicable, do one of the following and use the rest of the steps in this procedure:

Option	Description
To research IOCs on hosts:	<ul style="list-style-type: none"> • If you are using the predefined workflow, choose Analysis > Hosts > Indications of Compromise. • If you are using a custom workflow that does not include the Host IOC table view, click (switch workflow), then choose Host Indications of Compromise.
To research IOCs associated with users:	<ul style="list-style-type: none"> • If you are using the predefined workflow, choose Analysis > Users > Indications of Compromise. • If you are using a custom workflow that does not include the User IOC table view, click (switch workflow), then choose User Indications of Compromise.

Step 3 You have the following options:

- Use a different workflow, including a custom workflow, by clicking (**switch workflow**).
- Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
- Learn more about the contents of the columns in the table; see [Indications of Compromise Data Fields, on page 25](#).
- On a Host Indications of Compromise page: View the host profile for a compromised host by clicking **Compromised Host** in the **IP Address** column.
- On a User Indications of Compromise page: View the user profile associated with a compromise by clicking **Red User** in the **User** column.
- Mark IOC events resolved so they no longer appear in the list. To do so, check the check boxes next to the IOC events you want to modify, then click **Mark Resolved**.
- View details of events that triggered the IOC by clicking **View** (🔍) in the **First Seen** or **Last Seen** columns.
- See more options: Right-click a value in the table.

Indications of Compromise Data Fields

The following are the fields in Host or User IOC (indication of compromise) tables. Not every IOC-related table includes all fields.

IP Address (When viewing Host IOC data)

The IP address associated with the host that triggered the IOC.

User (When viewing User IOC data)

The username, realm, and authentication source of the user associated with the event that triggered the IOC.

Category

Brief description of the type of compromise indicated, such as `Malware Executed` or `Impact 1 Attack`.

Event Type

Identifier associated with a specific IOC, referring to the event that triggered it.

Description

Description of the impact on the potentially compromised host, such as `This host may be under remote control` or `Malware has been executed on this host`.

First Seen/Last Seen

The first/most recent date and time that events triggering the IOC occurred.

Domain

The domain of the host that triggered the IOC. This field is only present if you have ever configured the management center for multitenancy.

Related Topics

[Event Searches](#)

Editing Indication of Compromise Rule States for a Single Host or User


When enabled in a network discovery policy, indication of compromise rules apply to all hosts in the monitored network and to authoritative users that are associated with IOC events on that network. You can disable a rule for an individual host or user to avoid unhelpful IOC tags (for example, you may not want to see IOC tags for a DNS server.) If a rule is disabled in the applicable network discovery policy, it cannot be enabled for a specific host or user. Disabling a rule for a particular host does not affect tagging for the user involved in the same event, and vice-versa.

Procedure

-
- Step 1** Navigate to the **Indications of Compromise** section of a host or user profile.
 - Step 2** Click **Edit Rule States**.
 - Step 3** In the **Enabled** column for a rule, click the slider to enable or disable it.
 - Step 4** Click **Save**.
-

Viewing Source Events for Indication of Compromise Tags

You can use the Indications of Compromise section of the host profile and the user profile to navigate quickly to the events that triggered the IOC tags. Analyzing these events can give you the information you need to determine what, and whether, action is required to address threats of compromise.

Clicking **View** () next to the timestamp of an IOC tag navigates to the table view of events for the relevant event type, constrained to show only the event that triggered the IOC tag.

Only the first instance of a User IOC is displayed in the management center. Subsequent instances are caught by the DNS Server."

Procedure


- Step 1** In a host or user profile, navigate to the **Indications of Compromise** section.
- Step 2** Click **View** () in the **First Seen** or **Last Seen** column for the IOC tag you want to investigate.
-

Resolving Indication of Compromise Tags

After you have analyzed and addressed the threats indicated by an indication of compromise (IOC) tag, or if you determine that an IOC tag represents a false positive, you can mark an event resolved. Marking an event resolved removes it from the host profile and the user profile; when all active IOC tags on a profile are resolved, the **Compromised Host** or a user is associated with an indication of compromise **Red User icon** no longer appears. You can still view the IOC-triggering events for the resolved IOC.

If the events that triggered the IOC tag recur, the tag is set again unless you have disabled the IOC rule for the host or user.

Procedure

- Step 1** In a host or user profile, navigate to the **Indications of Compromise** section.
- Step 2** You have two choices:
- To mark an individual IOC tag resolved, click **Delete** () to the right of the tag you want to resolve.
 - To mark all IOC tags on the profile resolved, click **Mark All Resolved**.
-

Server Data

The system collects information about all servers running on hosts on monitored network segments. This information includes:

- the name of the server
- the application and network protocols used by the server
- the vendor and version of the server
- the IP address associated with the host running a server
- the port on which the server communicates

When the system detects a server, it generates a discovery event unless the associated host has already reached its maximum number of servers. You can use the management center web interface to view, search, and delete server events.

You can also base correlation rules on server events. For example, you could trigger a correlation rule when the system detects a chat server, such as ired, running on one of your hosts.

The system can add hosts to the network map from exported NetFlow records, but the available information for these hosts is limited; see [Differences between NetFlow and Managed Device Data](#).

Viewing Server Data

You can use the management center to view a table of detected servers. Then, you can manipulate the event view depending on the information you are looking for.

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.

The page you see when you access servers differs depending on the workflow you use. All the predefined workflows terminate in a host view, which contains a host profile for every host that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1

Access the server data:

- If you are using the predefined workflow, choose **Analysis > Hosts > Servers**.
- If you are using a custom workflow that does not include the table view of servers, click (**switch workflow**), then choose **Servers**.

Step 2

You have the following options:

- Use a different workflow, including a custom workflow, by clicking (**switch workflow**).
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [Server Data Fields, on page 28](#).
 - Edit server identities by checking the check boxes next to the events for servers you want to edit, then clicking **Set Server Identity**.
 - Right-click an item in the table to see options. (Not every column offers options.)
-

Server Data Fields

Descriptions of the fields that can be viewed and searched in the servers table follow below.

Last Used

The date and time the server was last used on the network or the date and time that the server was originally updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects a server information update.

IP Address

The IP address associated with the host running the server.

Port

The port where the server is running.

Protocol

The network or transport protocol used by the server.

Application Protocol

One of the following:

- the name of the application protocol for the server
- `pending` if the system cannot positively or negatively identify the server for one of several reasons
- `unknown` if the system cannot identify the server based on known server fingerprints or if the server was added through host input and did not include the application protocol

Category, Tags, Risk, or Business Relevance for Application Protocols

The categories, tags, risk level, and business relevance assigned to the application protocol. These filters can be used to focus on a specific set of data.

Vendor

One of the following:

- the server vendor as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its vendor based on known server fingerprints, or if the server was added to the network map using NetFlow data

Version

One of the following:

- the server version as identified by the system, Nmap or another active source, or that you specified using the host input feature
- blank, if the system cannot identify its version based on known server fingerprints, or if the server was added to the network map using NetFlow data

Web Application

The web application based on the payload content detected by the system in the HTTP traffic. Note that if the system detects an application protocol of `HTTP` but cannot detect a specific web application, the system supplies a generic web browsing designation.

Category, Tags, Risk, or Business Relevance for Web Applications

The categories, tags, risk level, and business relevance assigned to the web application. These filters can be used to focus on a specific set of data.

Hits

The number of times the server was accessed. For servers added using the host input feature, this value is always 0.

Source Type

One of the following values:

- User: user_name
- Application: app_name
- Scanner: scanner_type (Nmap or scanner added through network discovery configuration)
- Firepower, Firepower Port Match, Or Firepower Pattern Match for servers detected by the system
- NetFlow for servers added using NetFlow data

Domain

The domain of the host running the server. This field is only present if you have ever configured the management center for multitenancy.

Device

Either the managed device that detected the traffic or the device that processed NetFlow or host input data.

Current User

The user identity (username) of the currently logged in user on the host.

When a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Count

The number of events that match the information that appears in each row. This field appears only after you apply a constraint that creates two or more identical rows.

Related Topics

[Event Searches](#)

Application and Application Details Data

When a monitored host connects to another host, the system can, in many cases, determine what application was used. The system detects the use of many email, instant messaging, peer-to-peer, web applications, as well as other types of applications.

For each detected application, the system logs the IP address that used the application, the product, the version, and the number of times its use was detected. You can use the web interface to view, search, and delete application events. You can also update application data on a host or hosts using the host input feature.

If you know which applications are running on which hosts, you can use that knowledge to create host profile qualifications, which constrain the data you collect while building a traffic profile, and also can limit the conditions under which you want to trigger a correlation rule. You can also base correlation rules on the detection of application. For example, if you want your employees to use a specific mail client, you could trigger a correlation rule when the system detects a different mail client running on one of your hosts.

You can obtain the latest information about application detectors by carefully reading both the release notes for each system update and advisories for each VDB update.

To collect and store application data for analysis, make sure that you enable application detection in your network discovery policy.

Viewing Application Data

You can use the management center to view a table of detected applications. Then, you can manipulate the event view depending on the information you are looking for.

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.

The page you see when you access applications differs depending on the workflow you use. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

- Step 1** Access the application data:
- If you are using the predefined workflow, choose **Analysis > Hosts > Application Details**.
 - If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then choose **Clients**.
- Step 2** You have the following options:
- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [Application Data Fields, on page 31](#).
 - Open the Application Detail View for a specific application by clicking **Application Detail View** next to a client, application protocol, or web application.
 - View data in sources external to your system, by right-clicking 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, see [Event Investigation Using Web-Based Resources](#).
 - Gather intelligence about an event by right-clicking an event value in the table and choosing 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, see [Event Investigation Using Web-Based Resources](#).
-

Application Data Fields

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

Descriptions of the fields that can be viewed and searched in the applications table follow.

Application

The name of the detected application.

IP Address

The IP address associated with the host using the application.

Type

The type of application:

Application Protocols

Represents communications between hosts.

Client Applications

Represents software running on a host.

Web Applications

Represents the content or requested URL for HTTP traffic.

Category

A general classification for the application that describes its most essential function. Each application belongs to at least one category.

Tag

Additional information about the application. Applications can have any number of tags, including none.

Risk

How likely the application is to be used for purposes that might be against your organization's security policy. An application's risk can range from Very Low to Very High.

Of Application Protocol Risk, Client Risk, and Web Application Risk, the highest of the three detected, when available, in the traffic that triggered the intrusion event.

Business Relevance

The likelihood that the application is used within the context of your organization's business operations, as opposed to recreationally. An application's business relevance can range from Very Low to Very High.

Of Application Protocol Business Relevance, Client Business Relevance, and Web Application Business Relevance, the lowest of the three detected, when available, in the traffic that triggered the intrusion event.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the

current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

Domain

The domain of the host using the application. This field is only present if you have ever configured the management center for multitenancy.

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.

Related Topics

[Event Searches](#)

Viewing Application Detail Data

You can use the management center to view a table of detected application details. Then, you can manipulate the event view depending on the information you are looking for.

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.

The page you see when you access application details differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

- Step 1** Access the application details data:
- If you are using the predefined workflow, choose **Analysis > Hosts > Application Details**.
 - If you are using a custom workflow that does not include the table view of application details, click **(switch workflow)**, then select **Clients**.
- Step 2** You have the following options:
- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [Application Detail Data Fields, on page 34](#).
 - Open the Application Detail View for a specific application by clicking **Application Detail View** next to a client.
 - View data in available sources external to your system, by right-clicking 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, see [Event Investigation Using Web-Based Resources](#)
 - Gather intelligence about an event by right-clicking an event value in the table and choosing 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, see [Event Investigation Using Web-Based Resources](#).
-

Application Detail Data Fields

When the system detects traffic for a known client, application protocol, or web application, it logs information about the application and the host running it.

Descriptions of the fields that can be viewed and searched in the application details table follow.

Last Used

The time that the application was last used or the time that the application data was updated using the host input feature. The Last Used value is updated at least as often as the update interval you configured in the network discovery policy, as well as when the system detects an application information update.

IP Address

The IP address associated with the host using the application.

Client

The name of the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Version

The version of the application.

Category, Tags, Risk, or Business Relevance for Clients, Application Protocols, and Web Applications

The categories, tags, risk level, and business relevance assigned to the application. These filters can be used to focus on a specific set of data.

Application Protocol

The application protocol used by the application. Note that if the system detected an application protocol but could not detect a specific client, `client` is appended to the application protocol name to provide a generic name.

Web Application

The web application based on the payload content or URL detected by the system in the HTTP traffic. Note that if the system detects an application protocol of HTTP but cannot detect a specific web application, the system supplies a generic web browsing designation here.

Hits

The number of times the system detected the application in use. For applications added using the host input feature, this value is always 0.

Domain

The domain of the host using the application. This field is only present if you have ever configured the management center for multitenancy.

Device

The device that generated the discovery event containing the application detail.

Current User

The user identity (username) of the currently logged in user on the host.

Note that when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user. In addition, when a non-authoritative user is the current user on a host, that user still cannot be used for user control.

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.

Related Topics

[Event Searches](#)

Vulnerability Data

The system includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network. The operating systems, servers, and clients running on your hosts have different sets of associated vulnerabilities.

You can use the management center to:

- Track and review the vulnerabilities for each host.
- Deactivate vulnerabilities for a host after you patch the host or otherwise judge it immune to a vulnerability.

Vulnerabilities for vendorless and versionless servers are not mapped unless the applications protocols used by the servers are mapped in the management center configuration. Vulnerabilities for vendorless and versionless clients cannot be mapped.

Related Topics

[Mapping Vulnerabilities for Servers](#)

Vulnerability Data Fields

Except as noted, these fields appear on all pages under **Analysis > Hosts > Vulnerabilities**.

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.

CVE ID

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

To view details about this vulnerability in the National Vulnerability Database (NVD), right-click the CVE ID and choose **View description in NVD**.

Date Published

The date the vulnerability was published.

Description

A brief description of the vulnerability, from the National Vulnerability Database (NVD).

For the complete description, right-click the CVE ID and choose **View description in NVD** to view details in the National Vulnerability Database (NVD).

Impact

See "Vulnerability Impact" (below.)

Impact Qualification

This field is available only on the Vulnerability Details page.

Use the drop-down list to enable or disable a vulnerability. The management center ignores disabled vulnerabilities in its impact correlations.

The setting you specify here determines how the vulnerability is treated on a system-wide basis and is not limited to the host profile where you select the value.

Remote

Indicates whether the vulnerability is remotely exploitable (TRUE/FALSE).

Severity

The base score and Common Vulnerability Scoring System score (CVSS) from the National Vulnerability Database (NVD).

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

SVID

The vulnerability identification number that the system uses to track vulnerabilities.

To view details for this vulnerability, click **View** (👁).

Vulnerability Impact/Impact

The severity of the vulnerability on a scale of 0 to 10, with 10 being the most severe.

Related Topics

[Event Searches](#)

Vulnerability Deactivation

Deactivating a vulnerability prevents the system from using that vulnerability to evaluate intrusion impact correlations. You can deactivate a vulnerability after you patch the hosts on your network or otherwise judge them immune. Note that if the system discovers a new host that is affected by that vulnerability, the vulnerability is considered valid (and is not automatically deactivated) for that host.

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network. You can deactivate vulnerabilities within the vulnerabilities workflow only on:

- the second page of the default vulnerabilities workflow, **Vulnerabilities on the Network**, which shows only the vulnerabilities that apply to the hosts on your network
- a page in a vulnerabilities workflow, custom or predefined, that you constrained based on IP address using a search.

You can deactivate a vulnerability for a single host using the network map, using the host's host profile, or by constraining the vulnerabilities workflow based on the IP addresses of the host or hosts for which you want to deactivate vulnerabilities. For hosts with multiple associated IP addresses, this function applies only to the single, selected IP address of that host.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices if the vulnerability is activated in the ancestor domain.

Related Topics

[Deactivating Vulnerabilities for Individual Hosts](#)

[Deactivating Individual Vulnerabilities](#)

[Deactivating Multiple Vulnerabilities](#), on page 38

Viewing Vulnerability Data

You can use the management center to view a table of vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

The page you see when you access vulnerabilities differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of vulnerabilities. The table view contains a row for each vulnerability in the database, regardless of whether any of your detected hosts exhibit the vulnerabilities. The second page of the predefined workflow contains a row for each vulnerability (that you have not deactivated) that applies to detected hosts on your network. The predefined workflow terminates in a vulnerability detail view, which contains a detailed description for every vulnerability that meets your constraints.



Tip If you want to see the vulnerabilities that apply to a single host or set of hosts, perform a search for vulnerabilities, specifying an IP address or range of IP addresses for the hosts.

You can also create a custom workflow that displays only the information that matches your specific needs.

The table of vulnerabilities is not restricted by domain in a multidomain deployment.

Procedure

- Step 1** Access the table of vulnerabilities:
- If you are using the predefined vulnerabilities workflow, choose **Analysis > Hosts > Vulnerabilities**.
 - If you are using a custom workflow that does not include the table view of vulnerabilities, click (**switch workflow**), then choose **Vulnerabilities**.
- Step 2** You have the following options:
- Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Deactivate vulnerabilities so they are no longer used for intrusion impact correlation for currently vulnerable hosts; see [Deactivating Multiple Vulnerabilities, on page 38](#).
 - View the details for a vulnerability by clicking **View** (👁) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page. See options for viewing additional details at [Viewing Vulnerability Details, on page 38](#).
 - View the full text of a vulnerability title by right-clicking the title and choosing **Show Full Text**.
-

Viewing Vulnerability Details

Procedure

You can view vulnerability details in any of the following ways:

- Choose **Analysis > Hosts > Vulnerabilities**, and click **View** (👁) next to the SVID.
 - Choose **Analysis > Hosts > Third-Party Vulnerabilities** and click **View** (👁) next to the SVID.
 - Choose **Analysis > Hosts > Network Map**, and click **Vulnerabilities**.
 - View the profile of a host affected by the vulnerability (**Analysis > Hosts > Network Map**, click **Hosts**, then drill down and click the host you are investigating), and expand the **Vulnerabilities** section of the profile.
 - In any table under **Analysis > Hosts > Vulnerabilities**, right-click the value in the **CVE ID** column and choose **View description in NVD** to view that CVE on the NVD (National Vulnerabilities Database) web site.
-

Deactivating Multiple Vulnerabilities

Deactivating a vulnerability within a vulnerabilities workflow that is **not** constrained by IP addresses deactivates the vulnerability for *all* detected hosts on your network.

In a multidomain deployment, deactivating a vulnerability in an ancestor domain deactivates it in all descendant domains. Leaf domains can activate or deactivate a vulnerability for their devices so long as the vulnerability is activated in the ancestor domain.

Procedure

- Step 1** Access the table of vulnerabilities:
- If you are using the predefined vulnerabilities workflow, choose **Analysis > Hosts > Vulnerabilities**.
 - If you are using a custom workflow that does not include the table view of vulnerabilities, click (**switch workflow**), then choose **Vulnerabilities**.
- Step 2** Click **Vulnerabilities on the Network**.
- Step 3** Check the check boxes next to vulnerabilities you want to deactivate.
- Step 4** Click **Review** at the bottom of the page.
-

Related Topics

- [Deactivating Vulnerabilities for Individual Hosts](#)
- [Deactivating Individual Vulnerabilities](#)

Third-Party Vulnerability Data

The system includes its own vulnerability tracking database which is used, in conjunction with the system's fingerprinting capability, to identify the vulnerabilities associated with the hosts on your network.

You can augment the system's vulnerability data with imported network map data from third-party applications. To do so, your organization must be able to write scripts or create command line import files to import the data. For more information, see the *Firepower System Host Input API Guide*.

To include imported data in impact correlations, you must map third-party vulnerability information to the operating system and application definitions in the database. You cannot map third-party vulnerability information to client definitions.

Viewing Third-Party Vulnerability Data

After you use the host input feature to import third-party vulnerability data, you can use the management center to view a table of third-party vulnerabilities. Then, you can manipulate the event view depending on the information you are looking for.

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.

The page you see when you access third-party vulnerabilities differs depending on the workflow you use. There are two predefined workflows. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

- Step 1** Access the third-party vulnerabilities data:
- If you are using the predefined workflow, choose **Analysis > Hosts > Third-Party Vulnerabilities**.
 - If you are using a custom workflow that does not include the table view of third-party vulnerabilities, click (**switch workflow**), then choose **Vulnerabilities by Source** or **Vulnerabilities by IP Address**.

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking (**switch workflow**).
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [Third-Party Vulnerability Data Fields, on page 40](#).
 - View the vulnerability details for a third-party vulnerability by clicking **View** (👁) in the SVID column. Alternatively, constrain on the vulnerability ID and drill down to the vulnerability details page.
-

Third-Party Vulnerability Data Fields

Descriptions of the fields that can be viewed and searched in the third-party vulnerabilities table follow.

Vulnerability Source

The source of the third-party vulnerabilities, for example, QualysGuard or NeXpose.

Vulnerability ID

The ID number associated with the vulnerability for its source.

IP Address

The IP address associated with the host affected by the vulnerability.

Port

A port number, if the vulnerability is associated with a server running on a specific port.

Bugtraq ID

The identification number associated with the vulnerability in the Bugtraq database. (<http://www.securityfocus.com/bid/>)

CVE ID

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

SVID

The legacy vulnerability identification number that the system uses to track vulnerabilities

Click **View** (👁) to access the vulnerability details for the SVID.

Snort ID

The identification number associated with the vulnerability in the Snort ID (SID) database. That is, if an intrusion rule can detect network traffic that exploits a particular vulnerability, that vulnerability is associated with the intrusion rule's SID.

Note that a vulnerability can be associated with more than one SID (or no SIDs at all). If a vulnerability is associated with more than one SID, the vulnerabilities table includes a row for each SID.

Title

The title of the vulnerability.

Description

A brief description of the vulnerability.

Domain

The domain of the host with the vulnerability. This field is only present if you have ever configured the management center for multitenancy.

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.

Related Topics

[Event Searches](#)

Active Sessions, Users, and User Activity Data

Identity sources collect active session data, user data, and user activity data. The data is displayed in individual user-related workflows:

- **Active Sessions** — this workflow displays all current user sessions on your network. A single user running several simultaneous active sessions would occupy several rows in this table. For more information about the types of user data displayed in this workflow, see [Active Sessions Data, on page 48](#).
- **Users** — this workflow displays all users seen on your network. A single user occupies a single row in this table. For more information about the types of user data displayed in this workflow, see [User Data, on page 49](#).
- **User Activity** — this workflow displays all user activity seen on your network. A single user with more than one instance of user activity would occupy several rows in this table. For more information about the types of user activity displayed in this workflow, see [User Activity Data, on page 52](#).

For more information about the user identity sources that populate these workflows, see the [Cisco Secure Firewall Management Center Device Configuration Guide](#).

User-Related Fields

User-related data is displayed in the active sessions, users, and user activity tables.

Table 1: Active Sessions, Users, and User Activity Field Descriptions

Field	Description	Active Sessions Table	Users Table	User Activity Table
Active Session Count	The number of active sessions associated with the user.	No	Yes	No
Authentication Type	The type of authentication: No Authentication, Passive Authentication, Active Authentication, Guest Authentication, Failed Authentication, or VPN Authentication. For more information about the supported identity sources for each Authentication Type, see Cisco Secure Firewall Management Center Device Configuration Guide .	Yes	No	Yes
Available for Policy	A value of Yes means the user was retrieved from the user store (for example, Active Directory). A value of No means the management center received a report of a login for that user but the user is not in the user store. One way this can happen is if a user in an excluded group logs in to the user store. You can exclude groups from being downloaded when you configure a realm. Users not available for policy are recorded in the management center but are not sent to managed devices.	No	Yes	No
Count	Note The Count field is displayed only after you apply a constraint that creates two or more identical rows. Depending on the table, the number of sessions, users, or activity events that match the information that appears in a particular row.	Yes	Yes	Yes
Current IP	The IP address associated with the host that the user is logged into. This field is blank in the Users table if there are no active sessions for a user.	Yes	Yes	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
Department	<p>The user's department, as obtained by a realm. If there is no department explicitly associated with the user on your servers, the department is listed as whatever default group the server assigns. For example, on Active Directory, this is <code>Users (ad)</code>. This field is blank if:</p> <ul style="list-style-type: none"> You have not configured a realm. The management center cannot correlate the user in the management center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). 	Yes	Yes	No
Description	More information, if available, about the session, user, or user activity.	No	No	Yes
Device	<p>For user activity detected by traffic-based detection or an active authentication identity source, the name of the device that identified the user.</p> <p>For other types of user activity, the managing management center.</p> <p>Note If you have configured your VPN in a high-availability deployment, the device name displayed against active VPN sessions can be the primary or secondary device that identified the user session.</p>	Yes	No	Yes
Discovery Application	<p>The application or protocol used to detect the user.</p> <ul style="list-style-type: none"> For user activity detected by traffic-based detection, one of the following: ldap, pop3, imap, oracle, sip, http, ftp, mdns, or aim. <p>Note Users are not added to the database based on SMTP logins.</p> <ul style="list-style-type: none"> For all other user activity: ldap. 	Yes	Yes	Yes
Current IP Domain/Domain	<p>In the Active Sessions table, the multitenancy domain where the user activity was detected.</p> <p>In the Users table, the multitenancy domain associated with the user's realm.</p> <p>In the User Activity table, the multitenancy domain where the user activity was detected.</p> <p>This field is only present if you have ever configured the management center for multitenancy.</p>	Yes	Yes	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
E-Mail	The user's email address. This field is blank if: <ul style="list-style-type: none"> The user was added to the database via an AIM login. The user was added to the database via an LDAP login and there is no email address associated with the user on your LDAP servers. 	Yes	Yes	No
End Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the end value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	Yes	No	Yes
Endpoint Location	The IP address of the network device that used ISE to authenticate the user, as identified by ISE. If you do not configure ISE, this field is blank.	No	No	Yes
Endpoint Profile	The user's endpoint device type, as identified by Cisco ISE. If you do not configure ISE, this field is blank.	No	No	Yes
Event	The user activity event type.	No	No	Yes
First Name	The user's first name, as obtained by a realm. This field is blank if: <ul style="list-style-type: none"> You have not configured a realm. The management center cannot correlate the user in the management center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). There is no first name associated with the user on your servers. 	Yes	Yes	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
IP Address	<p>For User Login user activity, the IP address or internal IP address involved in the login:</p> <ul style="list-style-type: none"> • LDAP, POP3, IMAP, FTP, HTTP, MDNS, and AIM logins — the address of the user's host • SMTP and Oracle logins — the address of the server • SIP logins — the address of the session originator <p>An associated IP address does not mean the user is the current user for that IP address; when a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user.</p> <p>For other types of user activity, this field is blank.</p>	No	No	Yes
Last Name	<p>The user's last name, as obtained by a realm. This field is blank if:</p> <ul style="list-style-type: none"> • You have not configured a realm. • The management center cannot correlate the user in the management center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no last name associated with the user on your servers. 	Yes	Yes	No
Last Seen	<p>The date and time that a session was last initiated (or user data was updated) for the user.</p>	Yes	Yes	No
Login Time	<p>The date and time that the session was initiated for the user.</p>	Yes	No	No
Phone	<p>The user's telephone number, as obtained by a realm. This field is blank if:</p> <ul style="list-style-type: none"> • You have not configured a realm. • The management center cannot correlate the user in the management center database with an LDAP record (for example, for users added to the database via an AIM, Oracle, or SIP login). • There is no telephone number associated with the user on your servers. 	Yes	Yes	No

Field	Description	Active Sessions Table	Users Table	User Activity Table
Realm	The identity realm associated with the user. Note Active sessions for Azure AD realm users appear only in the Active Sessions new UI layout and not in the legacy UI.	Yes	Yes	Yes
Security Group Tag	The Security Group Tag (SGT) attribute applied by Cisco TrustSec as the packet entered a trusted TrustSec network. If you do not configure ISE, this field is blank.	No	No	Yes
Session Duration	The duration of the user session, calculated from the Login Time and the current time.	Yes	No	No
Start Port	If the user was reported by the TS Agent and their session is currently active, this field identifies the start value for the port range assigned to the user. This field is blank if the user's TS Agent session is inactive or if the user was reported by another identity source.	Yes	No	Yes
Time	The time that the system detected the user activity.	No	No	Yes
User	At minimum, this field displays the user's realm and username. For example, Lobby\jsmith, where Lobby is the realm and jsmith is the username. If a realm downloads additional user data from an LDAP server and the system associates it with a user, this field also displays the user's first name, last name, and type. For example, John Smith (Lobby\jsmith, LDAP), where John Smith is the user's name and LDAP is the type. Note Because traffic-based detection can record unsuccessful AIM logins, the management center may store invalid AIM users (for example, if a user misspelled his or her username).	Yes	Yes	No
Username	The username associated with the user.	Yes	Yes	Yes
VPN Bytes In	For Remote Access VPN-reported user activity, the total number of bytes received from the remote peer or client by the threat defense. Note You can view the total number of bytes received once the user's VPN session is terminated. For ongoing VPN sessions, this is not a dynamic counter. For other types of user activity, this field is blank.	No	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
VPN Bytes Out	<p>For Remote Access VPN-reported user activity, the total number of bytes transmitted to the remote peer or client by the threat defense.</p> <p>Note You can view the total number of bytes transmitted once the user's VPN session is terminated. For ongoing VPN sessions, this is not a dynamic counter.</p> <p>For other types of user activity, this field is blank.</p>	No	No	Yes
VPN Client Application	<p>For Remote Access VPN-reported user activity, the remote user's AnyConnect VPN module of Cisco Secure Client application.</p> <p>For other types of user activity, this field is blank.</p>	Yes	No	Yes
VPN Client Country	<p>For Remote Access VPN-reported user activity, the country name as reported by the Secure Client VPN.</p> <p>For other types of user activity, this field is blank.</p>	No	No	Yes
VPN Client OS	<p>For Remote Access VPN-reported user activity, the remote user's endpoint operating system as reported by the Secure Client VPN.</p> <p>For other types of user activity, this field is blank.</p>	Yes	No	Yes
VPN Client Public IP	<p>For Remote Access VPN-reported user activity, the publicly routable IP address of the Secure Client VPN device.</p> <p>For other types of user activity, this field is blank.</p>	Yes	No	Yes
VPN Connection Duration	<p>For Remote Access VPN-reported user activity, the total time (HH:MM:SS) that the session was active.</p> <p>For other types of user activity, this field is blank.</p>	No	No	Yes
VPN Connection Profile	<p>For Remote Access VPN-reported user activity, the name of the connection profile (tunnel group) used by the VPN session. Connection profiles are part of a Remote Access VPN Policy.</p> <p>For other types of user activity, this field is blank.</p>	Yes	No	Yes

Field	Description	Active Sessions Table	Users Table	User Activity Table
VPN Group Policy	For Remote Access VPN-reported user activity, the name of the group policy assigned to the client when the VPN session is established; either the statically-assigned group policy associated with the VPN Connection Profile, or the dynamically-assigned group policy if RADIUS is used for authentication. If assigned by the RADIUS server, this group policy overrides the static policy configured for the VPN Connection Profile. Group policies configure common attributes for groups of users in Remote Access VPN policies. For other types of user activity, this field is blank.	Yes	No	Yes
VPN Session Type	For Remote Access VPN-reported user activity, the type of session: LAN-to-LAN or Remote. For other types of user activity, this field is blank.	Yes	No	Yes

Active Sessions Data

The **Analysis > Users > Active Sessions** workflow displays select information about current user sessions. When a user on your network runs several sessions simultaneously, the system can uniquely identify the sessions if:

- they have unique **IP Address** values.
- they have unique **Start Port** and **End Port** values, as provided by the Cisco Terminal Services (TS) Agent.
- they have unique **Current IP Domain** values.
- they were authenticated by different identity sources.
- they were associated with different identity realms.

For more information about the user and user activity data stored by the system, see [User Data, on page 49](#) and [User Activity Data, on page 52](#).

For information about general user-related event troubleshooting and Remote Access VPN Troubleshooting, see the *Troubleshoot Realms and User Downloads* and *VPN Troubleshooting* in the [Cisco Secure Firewall Management Center Device Configuration Guide](#).

Viewing Active Session Data

You can view a table of active sessions, and then manipulate the event view depending on the information you are looking for.

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.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

- Step 1** Access the users data:
- If you are using the predefined workflow, click **Analysis > Users > Active Sessions**.
 - If you are using a custom workflow that does not include the table view of active sessions, click (**switch workflow**), then choose **Active Sessions**.
- Step 2** You have the following options:
- Use a different workflow, including a custom workflow, by clicking (**switch workflow**).
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [Active Sessions Data, on page 48](#) and [User-Related Fields, on page 41](#).

User Data

When an identity source reports a user login for a user who is not already in the database, the user is added to the database, unless you have specifically restricted that login type.

The system updates the users database when one of the following occurs:

- A user on the management center manually deletes a non-authoritative user from the Users table.
- An identity source reports a logoff by that user.
- A realm ends the user session as specified by the realm's **User Session Timeout: Authenticated Users**, **User Session Timeout: Failed Authentication Users**, or **User Session Timeout: Guest Users** setting.



Note If you have ISE/ISE-PIC configured, you may see host data in the users table. Because host detection by ISE/ISE-PIC is not fully supported, you cannot perform user control using ISE-reported host data.

The type of user login that the system detected determines what information is stored about the new user.

Identity Source	Login Type	User Data Stored
ISE/ISE-PIC	Active Directory LDAP RADIUS RSA	<ul style="list-style-type: none"> • username • current IP address • Security Group Tag (SGT) — not supported with ISE-PIC • endpoint profile/device type — not supported with ISE-PIC • endpoint location/location IP — not supported with ISE-PIC • type (LDAP)

Identity Source	Login Type	User Data Stored
TS Agent	Active Directory	<ul style="list-style-type: none"> • username • current IP address • start port • end port • type (LDAP)
captivity portal	Active Directory LDAP	<ul style="list-style-type: none"> • username • current IP address • type (LDAP)
traffic-based detection	LDAP AIM Oracle SIP HTTP FTP MDNS	<ul style="list-style-type: none"> • username • current IP address • type (AD)
	POP3 IMAP	<ul style="list-style-type: none"> • username • current IP address • email address • type (pop3 or imap)



Note No data about Microsoft Azure Active Directory users is displayed in this table.

If you configure a realm to automatically download users, the management center queries the servers based on the interval you specified. It may take five to ten minutes for the management center database to update with user metadata after the system detects a new user login. The management center obtains the following information and metadata about each user:

- username
- first and last names
- email address
- department
- telephone number

- current IP address
- Security Group Tag (SGT), if available
- endpoint profile, if available
- endpoint location, if available
- start port, if available
- end port, if available

The number of users the management center can store in its database depends on your management center model. When a non-authoritative user login to a host is detected, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user login is detected for that host, only another authoritative user login changes the current user.

Note that traffic-based detection of AIM, Oracle, and SIP logins create duplicate user records because they are not associated with any of the user metadata that the system obtains from LDAP servers. To prevent overuse of user count because of duplicate user records from these protocols, configure traffic-based detection to ignore those protocols.

You can search, view, and delete users from the database; you can also purge all users from the database.

For information about general user-related event troubleshooting, see [Cisco Secure Firewall Management Center Device Configuration Guide](#).

Viewing User Data

You can view a table of users, and then manipulate the event view depending on the information you are looking for.

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.

The page you see when you access users differs depending on the workflow you use. You can use the predefined workflow, which includes a table view of users that lists all detected users, and terminates in a user details page. The user details page provides information on every user that meets your constraints.

Procedure

- Step 1** Access the users data:
- If you are using the predefined workflow, choose **Analysis > Users > Users**.
 - If you are using a custom workflow that does not include the table view of users, click **(switch workflow)**, then choose **Users**.
- Step 2** You have the following options:
- Use a different workflow, including a custom workflow, by clicking **(switch workflow)**.
 - Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
 - Learn more about the contents of the columns in the table; see [User-Related Fields, on page 41](#).
-

User Activity Data

The system generates events that communicate the details of user activity on your network. When the system detects user activity, the user activity data is logged to the database. You can view, search, and delete user activity; you can also purge all user activity from the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

The system also correlates user activity with other types of events. For example, intrusion events can tell you the users who were logged into the source and destination hosts at the time of the event. This correlation can tell you who was logged into the host that was targeted by an attack, or who initiated an internal attack or portscan.

You can also use user activity in correlation rules. Based on the type of user activity as well as other criteria that you specify, you can build correlation rules that, when used in a correlation policy, launch remediations and alert responses when network traffic meets your criteria.



Note If you have ISE/ISE-PIC configured, you may see host data in the users table. Because host detection by ISE/ISE-PIC is not fully supported, you cannot perform user control using ISE-reported host data.

Descriptions of the four types of user activity data follow.

New User Identity

This type of event is generated when the system detects a login by an unknown user that is not in the database.

The system logs a user activity event when a user is seen on your network for the first time. Subsequent appearances by that user do not log new user activity events. However, if the user's IP address changes, the system logs a new user activity event.

User Login

This type of event is generated when any of the following occur:

- Captive portal performs a successful or failed user authentication.
- Traffic-based detection detects a successful or failed user login.



Note SMTP logins detected by traffic-based detection are not recorded unless there is already a user with a matching email address in the database.

When a non-authoritative user logs into a host, that login is recorded in the user and host history. If no authoritative user is associated with the host, a non-authoritative user can be the current user for the host. However, after an authoritative user logs into the host, only a login by another authoritative user changes the current user.

If you are using captive portal or traffic-based detection, note the following about failed user login and failed user authentication data:

- Failed logins reported by traffic-based detection (LDAP, IMAP, FTP, and POP3 traffic) are displayed in the table view of user activity, but not in the table view of users. If a known user failed to log in, the system identifies them by their username. If an unknown user failed to log in, the system uses **Failed Authentication** as their username.
- Failed authentications reported by captive portal are displayed in both the table view of user activity and the table view of users. If a known user failed to authenticate, the system identifies them by their username. If an unknown user failed to authenticate, the system identifies them by the username they entered.

Delete User Identity

This type of event is generated when you manually delete a user from the database.

User Identity Dropped: User Limit Reached

This type of event is generated when the system detects a user that is not in the database, but cannot add the user because you have reached the maximum number of users in the database as determined by your management center model.

After you reach the user limit, in most cases the system stops adding new users to the database. To add new users, you must either manually delete old or inactive users from the database, or purge all users from the database.

However, the system favors authoritative users. If you have reached the limit and the system detects a login for a previously undetected authoritative user, the system deletes the non-authoritative user who has remained inactive for the longest time, and replaces it with the new authoritative user.

User Indications of Compromise Events

The following user IOC changes are logged in the user activity database:

- When indications of compromise are resolved.
- When indication of compromise rules are enabled or disabled for users.

For information about general user-related event troubleshooting, see the [Cisco Secure Firewall Management Center Device Configuration Guide](#).

Viewing User Activity Data

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.

You can view a table of user activity, and then manipulate the event view depending on the information you are looking for. The page you see when you access user activity differs depending on the workflow you use. You can use the predefined workflow, which includes the table view of user activity and terminates in a user details page, which contains user details for every user that meets your constraints. You can also create a custom workflow that displays only the information that matches your specific needs.

Procedure

Step 1

Access the user activity data:

- If you are using the predefined workflow, choose **Analysis > Users > User Activity**.

- If you are using a custom workflow that does not include the table view of user activity, click (**switch workflow**), then choose **User Activity**.

Tip If no events appear, you may need to adjust the time range; see [Changing the Time Window](#).

Step 2 You have the following options:

- Use a different workflow, including a custom workflow, by clicking (**switch workflow**).
- Perform basic workflow actions; see [Using Discovery and Identity Workflows, on page 6](#).
- Learn more about the contents of the columns in the table; see [User-Related Fields, on page 41](#).

User Profile and Host History

You can learn more about a specific user by viewing the User pop-up window. The page that appears, called the "User Profile" in this document, is titled "User Identity" in the web interface.

You can display the window from:

- any event view that associates user data with other kinds of events
- the table view of active sessions
- the table view of users

User information also appears in the terminating page for users workflows.

The user data you see is the same as you would see in the table view of users.

Indications of Compromise Section

For information about this section, see:

- *Indications of Compromise* in the [Cisco Secure Firewall Management Center Device Configuration Guide](#)
- [Indications of Compromise Data Fields, on page 25](#)
- [Editing Indication of Compromise Rule States for a Single Host or User, on page 26](#)
- [Resolving Indication of Compromise Tags, on page 27](#)
- [Viewing Source Events for Indication of Compromise Tags, on page 26](#)

Host History Section

The host history provides a graphic representation of the last twenty-four hours of the user's activity. A list of IP addresses of the hosts that the user logged into and logged off of approximates login and logout times with bar graphs. A typical user might log on to and off of multiple hosts in the course of a day. For example, periodic automated logins to a mail server would display as multiple short sessions, while longer logins (such as during working hours) display longer sessions.

If you use traffic-based detection or captive portal to capture failed logins, the host history also includes hosts where the user failed to log in.

The data used to generate the host history is stored in the user history database, which by default stores 10 million user login events. If you do not see any data in the host history for a particular user, either that user is inactive, or you may need to increase the database limit.

Related Topics

[User Data Fields](#)

Viewing User Details and Host History

Procedure

You have two options:

- In any event view that lists users, click user that appears next to a user identity **User icon**, or, for users associated with an indication of compromise, **Red User icon**.
 - In any users workflow, click the Users terminating page.
-

History for Working with Discovery Events

Table 2:

Feature	Minimum Management Center	Minimum Threat Defense	Details
Vulnerabilities pages changes	6.7	Any	<p>Bugtraq and its vulnerability data are no longer available. The following changes have been made:</p> <ul style="list-style-type: none"> • Most vulnerability data now comes from the National Vulnerability Database (NVD). • Obsolete and redundant fields have been removed. • A new CVE ID column has been added to table views, and a new Severity field has been added to tables and details pages. • You can now right-click the CVE ID in tables to view details about that vulnerability in the NVD. • The Vulnerability Impact column in tables has been renamed to Impact. (No change to the field name in Detail views.) • When viewing vulnerabilities in host profiles under Analysis > Hosts > Network Map > Hosts, details for vulnerabilities (excluding third-party vulnerabilities) use the new set of fields. • The Bugtraq option has been removed from the Vulnerabilities options on the Analysis > Hosts > Network Map > Vulnerabilities page. <p>Modified screens:</p> <ul style="list-style-type: none"> • All pages under Analysis > Hosts > Vulnerabilities • Hosts and Vulnerabilities tabs on Analysis > Hosts > Network Map pages <p>Supported Platforms: management center</p>