

Network Discovery Policies

The following topics describe how to create, configure, and manage network discovery policies:

- Overview: Network Discovery Policies, on page 1
- Requirements and Prerequisites for Network Discovery Policies, on page 2
- Network Discovery Customization, on page 2
- Network Discovery Rules, on page 3
- Configuring Advanced Network Discovery Options, on page 12
- Troubleshooting Your Network Discovery Strategy, on page 20

Overview: Network Discovery Policies

The network discovery policy on the Cloud-Delivered Firewall Management Center controls how the system collects data on your organization's network assets and which network segments and ports are monitored.

You can configure network discovery policies only for Secure Firewall Threat Defense devices that send events to a Network Analytics Manager. (The Network Analytics Manager is an on-premises Secure Firewall Management Center that is configured to provide event analytics only.)

Discovery rules within the policy specify which networks and ports the system monitors to generate discovery data based on network data in traffic, and the zones to which the policy is deployed. Within a rule, you can configure whether hosts, applications, and non-authoritative users are discovered. You can create rules to exclude networks and zones from discovery. You can configure discovery of data from NetFlow exporters and restrict the protocols for traffic where user data is discovered on your network.

The network discovery policy has a single default rule in place, configured to discover applications from all observed traffic. The rule does not exclude any networks, zones, or ports, host and user discovery is not configured, and the rule is not configured to monitor a NetFlow exporter. This policy is deployed by default to any managed devices when they are registered to the Cloud-Delivered Firewall Management Center. To begin collecting host or user data, you must add or modify discovery rules and re-deploy the policy to a device.

If you want to adjust the scope of network discovery, you can create additional discovery rules and modify or remove the default rule.

Remember that the access control policy for each managed device defines the traffic that you permit for that device and, therefore, the traffic you can monitor with network discovery. If you block certain traffic using access control, the system cannot examine that traffic for host, user, or application activity. For example, if an access control policy blocks access to social networking applications, the system cannot provide any discovery data on those applications.

If you enable traffic-based user detection in your discovery rules, you can detect non-authoritative users through user login activity in traffic over a set of application protocols. You can disable discovery in particular protocols across all rules if needed. Disabling some protocols can help avoid reaching the user limit associated with your Cloud-Delivered Firewall Management Center model, reserving available user count for users from the other protocols.

Advanced network discovery settings allow you to manage what data is logged, how discovery data is stored, what indications of compromise (IOC) rules are active, what vulnerability mappings are used for impact assessment, and what happens when sources offer conflicting discovery data. You can also add sources for host input and NetFlow exporters to monitor.

Requirements and Prerequisites for Network Discovery Policies

Model support

Any.

Supported domains

Leaf

User roles

- Admin
- Discovery Admin

Network Discovery Customization

The information about your network traffic collected by the system is most valuable to you when the system can correlate this information to identify the hosts on your network that are most vulnerable and most important.

As an example, if you have several devices on your network running a customized version of SuSE Linux, the system cannot identify that operating system and so cannot map vulnerabilities to the hosts. However, knowing that the system has a list of vulnerabilities for SuSE Linux, you may want to create a custom fingerprint for one of the hosts that can then be used to identify the other hosts running the same operating system. You can include a mapping of the vulnerability list for SuSE Linux in the fingerprint to associate that list with each host that matches the fingerprint.

The system also allows you to input host data from third-party systems directly into the network map, using the host input feature. However, third-party operating system or application data does not automatically map to vulnerability information. If you want to see vulnerabilities and perform impact correlation for hosts using third-party operating system, server, and application protocol data, you must map the vendor and version information from the third-party system to the vendor and version listed in the vulnerability database (VDB). You also may want to maintain the host input data on an ongoing basis. Note that even if you map application data to system vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

If the system cannot identify application protocols running on hosts on your network, you can create user-defined application protocol detectors that allow the system to identify the applications based on a port

or a pattern. You can also import, activate, and deactivate certain application detectors to further customize the application detection capability.

You can also replace detection of operating system and application data using scan results from the Nmap active scanner or augment the vulnerability lists with third-party vulnerabilities. The system may reconcile data from multiple sources to determine the identity for an application.

Configuring the Network Discovery Policy

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Configure the following components of your policy:

- Discovery rules — See Configuring Network Discovery Rules, on page 4.
- Traffic-based detection for users — See Configuring Traffic-Based User Detection, on page 11.
- Advanced network discovery options — See Configuring Advanced Network Discovery Options, on page 12.
- Custom operating system definitions (fingerprints) — See Creating a Custom Fingerprint for Clients and Creating a Custom Fingerprint for Servers.

Network Discovery Rules

Network discovery rules allow you to tailor the information discovered for your network map to include only the specific data you want. Rules in your network discovery policy are evaluated sequentially. You can create rules with overlapping monitoring criteria, but doing so may affect your system performance.

When you exclude a host or a network from monitoring, the host or network does not appear in the network map and no events are reported for it. However, when the host discovery rules for the local IP are disabled, the detection engine instances are impacted by a higher processing load, as it builds data from each flow afresh rather than using the existing host data.

We recommend that you exclude load balancers (or specific ports on load balancers) and NAT devices from monitoring. These devices may create excessive and misleading events, filling the database and overloading the Cloud-Delivered Firewall Management Center. For example, a monitored NAT device might exhibit multiple updates of its operating system in a short period of time. If you know the IP addresses of your load balancers and NAT devices, you can exclude them from monitoring.

Tip The system can identify many load balancers and NAT devices by examining your network traffic.

In addition, if you need to create a custom server fingerprint, you should temporarily exclude from monitoring the IP address that you are using to communicate with the host you are fingerprinting. Otherwise, the network map and discovery event views will be cluttered with inaccurate information about the host represented by that IP address. After you create the fingerprint, you can configure your policy to monitor that IP address again.

Cisco also recommends that you **not** monitor the same network segment with NetFlow exporters and managed devices. Although ideally you should configure your network discovery policy with non-overlapping rules, the system does drop duplicate connection logs generated by managed devices. However, you **cannot** drop duplicate connection logs for connections detected by both a managed device and a NetFlow exporter.

Configuring Network Discovery Rules

You can configure discovery rules to tailor the discovery of host and application data to your needs.

Tip In most cases, we recommend restricting discovery to the addresses in RFC 1918.

Before you begin

- Make sure you are logging connections for the traffic where you want to discover network data.
- If you want to collect exported NetFlow records, add a NetFlow Exporter as described in Adding NetFlow Exporters to a Network Discovery Policy, on page 17.
- If you will want to view discovery performance graphs, you must enable hosts, users, and applications in your discovery rule. Note that this may impact system performance.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Set the **Action** for the rule as described in Actions and Discovered Assets, on page 4.

Step 4 Set optional discovery parameters:

- Restrict the rule action to specific networks; see Restricting the Monitored Network, on page 6.
- Restrict the rule action to traffic in specific zones; see Configuring Zones in Network Discovery Rules, on page 9.
- Exclude ports from monitoring; see Excluding Ports in Network Discovery Rules, on page 8.
- Configure the rule for NetFlow data discovery; see Configuring Rules for NetFlow Data Discovery, on page 6.

Step 5 Click **Save**.

What to do next

- Deploy configuration changes.

Actions and Discovered Assets

When you configure a discovery rule, you must select an action for the rule. The effect of that action depends on whether you are using the rule to discover data from a managed device or from a NetFlow exporter.

The following table describes what assets are discovered by rules with the specified action settings in those two scenarios.

Table 1: Discovery Rule Actions

Action	Option	Managed Device	NetFlow Exporter
Exclude	--	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.	Excludes the specified network from monitoring. If the source or destination host for a connection is excluded from discovery, the connection is recorded but discovery events are not created for excluded hosts.
Discover	Hosts	Adds hosts to the network map based on discovery events. (Optional, unless user discovery is enabled, then required.)	Adds hosts to the network map and logs connections based on NetFlow records. (Required)
Discover	Applications	Adds applications to the network map based on application detectors. Note that you cannot discover hosts or users in a rule without also discovering applications. (Required)	Adds application protocols to the network map based on NetFlow records and the port-application protocol correlation in `/etc/sf/services`. (Optional)
Discover	Users	Adds users to the users table and logs user activity based on traffic-based detection on the user protocols configured in the network discovery policy. (Optional)	n/a
Log NetFlow Connections	--	n/a	Logs NetFlow connections only. Does not discover hosts or applications.

If you want the rule to monitor managed device traffic, application logging is required. If you want the rule to monitor users, host logging is required. If you want the rule to monitor exported NetFlow records, you cannot configure it to log users, and logging applications is optional.

Note The system detects connections in exported NetFlow records based on the **Action** settings in the network discovery policy. The system detects connections in managed device traffic based on access control policy settings.

Monitored Networks

A discovery rule causes discovery of monitored assets only in traffic to and from hosts in the specified networks. For a discovery rule, discovery occurs for connections that have at least one IP address within the networks specified, with events generated only for IP addresses within the networks to monitor. The default discovery rule discovers applications from all observed traffic (`0.0.0.0/0` for all IPv4 traffic, and `::/0` for all IPv6 traffic).

If you configure a rule to handle NetFlow discovery and log only connections data, the system also logs connections to and from IP addresses in the specified networks. Note that network discovery rules provide the only way to log NetFlow network connections.

You can also use network object or object groups to specify the networks to monitor.

Restricting the Monitored Network

Every discovery rule must include at least one network.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Click **Networks**, if it is not already open.

Step 4 (Optional) Add network objects to the Available Networks list as described in Creating Network Objects During Discovery Rule Configuration, on page 7.

If you modify a network object used in the network discovery policy, the changes do not take effect for discovery until you deploy the configuration changes.

Step 5 Specify a network:

- Choose a network from the **Available Networks** list. If the network does not immediately appear on the list, click **Reload** ( ).
- Enter the IP address into the text box below the Available Networks label.

Step 6 Click **Add**.

Step 7 Click **Save**.

What to do next

- Deploy configuration changes.

Configuring Rules for NetFlow Data Discovery

The system can use data from NetFlow exporters to generate connection and discovery events, and to add host and application data to the network map.

If you choose a NetFlow exporter in a discovery rule, the rule is limited to discovery of NetFlow data for the specified networks. Choose the NetFlow device to monitor before you configure other aspects of rule behavior, as the available rule actions change when you choose a NetFlow device. You cannot configure port exclusions for monitoring NetFlow exporters.

Before you begin

- Add NetFlow-enabled devices to the network discovery policy; see Adding NetFlow Exporters to a Network Discovery Policy, on page 17.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Choose **NetFlow Device**.

Step 4 From the **NetFlow Device** drop-down list, choose the IP address of the NetFlow exporter to be monitored.

Step 5 Specify the type of NetFlow data you want the system managed device to collect:

- Connection only — Choose `Log NetFlow Connections` from the **Action** drop-down list.
- Host, Application, and Connection — Choose `Discover` from the **Action** drop-down list. The system automatically checks the **Hosts** check box and enables collection of connection data. Optionally, you can check the **Application** check box to collect application data.

Step 6 Click **Save**.

What to do next

- Deploy configuration changes.

Creating Network Objects During Discovery Rule Configuration

You can add new network objects to the list of available networks that appears in a discovery rule by adding them to the list of reusable network objects and groups.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 In **Networks**, click **Add Rule**.

Step 3 Click **Add** (+) next to **Available Networks**.

Step 4 Create a network object as described in Creating Network Objects.

Step 5 Finish adding the network discovery rule as described in Configuring Network Discovery Rules, on page 4.

Port Exclusions

Just as you can exclude hosts from monitoring, you can exclude specific ports from monitoring. For example:

- Load balancers can report multiple applications on the same port in a short period of time. You can configure your network discovery rules so that they exclude that port from monitoring, such as excluding port 80 on a load balancer that handles a web farm.
- Your organization may use a custom client that uses a specific range of ports. If the traffic from this client generates excessive and misleading events, you can exclude those ports from monitoring. Similarly, you may decide that you do not want to monitor DNS traffic. In that case, you could configure your rules so that your discovery policy does not monitor port 53.

When adding ports to exclude, you can decide whether to use a reusable port object from the Available Ports list, add ports directly to the source or destination exclusion lists, or create a new reusable port and then move it into the exclusion lists.

Note You cannot exclude ports in rules handling NetFlow data discovery.

Excluding Ports in Network Discovery Rules

You cannot exclude ports in rules handling NetFlow data discovery.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Click **Port Exclusions**.

Step 4 Optionally, add port objects to the Available Ports list as described in Creating Port Objects During Discovery Rule Configuration, on page 8.

Step 5 Exclude specific source ports from monitoring, using either of the following methods:

- Choose a port or ports from the **Available Ports** list and click **Add to Source**.
- To exclude traffic from a specific source port without adding a port object, under the **Selected Source Ports** list, choose a **Protocol**, enter a **Port** number (a value from 1 to 65535), and click **Add**.

Step 6 Exclude specific destination ports from monitoring, using either of the following methods:

- Choose a port or ports from the **Available Ports** list and click **Add to Destination**.
- To exclude traffic from a specific destination port without adding a port object, under the **Selected Destination Ports** list, choose a **Protocol**, enter a **Port** number, and click **Add**.

Step 7 Click **Save** to save the changes you made.

What to do next

- Deploy configuration changes.

Creating Port Objects During Discovery Rule Configuration

You can add new port objects to the list of available ports that appears in a discovery rule by adding them to the list of reusable port objects and groups that can be used anywhere in the system.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 In Networks, click **Add Rule**.

Step 3 Click **Port Exclusions**.

Step 4 To add a port to the Available Ports list, click **Add** (+).

Step 5 Enter **Name**.

Step 6 In the **Protocol** field, specify the protocol of the traffic you want to exclude.

Step 7 In the **Port** field, enter the ports you want to exclude from monitoring.

You can specify a single port, a range of ports using the dash (-), or a comma-separated list of ports and port ranges. Allowed port values are from 1 to 65535.

Step 8 Click **Save**.

Step 9 If the port does not immediately appear on the list, click **Refresh**.

What to do next

- Deploy configuration changes.

Zones in Network Discovery Rules

To improve performance, discovery rules can be configured so that the zones in the rule include the sensing interfaces on your managed devices that are physically connected to the networks-to-monitor in the rule.

Unfortunately, you may not always be kept informed of network configuration changes. A network administrator may modify a network configuration through routing or host changes without informing you, which may make it challenging to stay on top of proper network discovery policy configurations. If you do not know how the sensing interfaces on your managed devices are physically connected to your network, leave the zone configuration as the default. This default causes the system to deploy the discovery rule to all zones in your deployment. (If no zones are excluded, the system deploy the discovery policy to all zones.)

Configuring Zones in Network Discovery Rules

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Add Rule**.

Step 3 Click **Zones**.

Step 4 Choose a zone or zones from the **Available Zones** list.

Step 5 Click **Save** to save the changes you made.

What to do next

- Deploy configuration changes.

The Traffic-Based Detection Identity Source

Traffic-based detection is the only non-authoritative identity source supported by the system. When configured, managed devices detect LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), FTP, HTTP, MDNS, and SMTP logins on the networks you specify. The data gained from traffic-based detection can be used only for user awareness. Unlike authoritative identity sources, you configure traffic-based detection in your network discovery policy as described in Configuring Traffic-Based User Detection, on page 11.

Note the following limitations:

- Traffic-based detection interprets only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.
- Traffic-based detection detects AIM logins using the OSCAR protocol only. They cannot detect AIM logins using TOC2.
- Traffic-based detection cannot restrict SMTP logging. This is because users are not added to the database based on SMTP logins; although the system detects SMTP logins, the logins are not recorded unless there is already a user with a matching email address in the database.

Traffic-based detection also records failed login attempts. A failed login attempt does not add a new user to the list of users in the database. The user activity type for detected failed login activity detected by traffic-based detection is **Failed User Login**.

Note The system cannot distinguish between failed and successful HTTP logins. To see HTTP user information, you must enable **Capture Failed Login Attempts** in the traffic-based detection configuration.

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior for more information.

Traffic-Based Detection Data

When a device detects a login using traffic-based detection, it sends the following information to the Cloud-Delivered Firewall Management Center to be logged as user activity:

- The user name identified in the login.
- The time of the login.
- The IP address involved in the login, which can be the IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for HTTP, MDNS, FTP, SMTP and Oracle logins), or the session originator (for SIP logins).
- The user's email address (for POP3, IMAP, and SMTP logins).
- The name of the device that detected the login.

If the user was previously detected, the Cloud-Delivered Firewall Management Center updates that user's login history. Note that the Cloud-Delivered Firewall Management Center can use the email addresses in POP3 and IMAP logins to correlate with LDAP users. This means that, for example, if the Cloud-Delivered Firewall Management Center detects a new IMAP login, and the email address in the IMAP login matches that for an existing LDAP user, the IMAP login does not create a new user; rather, it updates the LDAP user's history.

If the user was previously undetected, the Cloud-Delivered Firewall Management Center adds the user to the users database. Unique AIM, SIP, and Oracle logins always create new user records, because there is no data

in those login events that the Cloud-Delivered Firewall Management Center can correlate with other login types.

The Cloud-Delivered Firewall Management Center does **not** log user activity or user identities in the following cases:

- If you configured the network discovery policy to ignore that login type
- If a managed device detects an SMTP login, but the users database does not contain a previously detected LDAP, POP3, or IMAP user with a matching email address

The user data is added to the users table.

Traffic-Based Detection Strategies

You can restrict the protocols where user activity is discovered to reduce the total number of detected users so you can focus on users likely to provide the most complete user information. Limiting protocol detection helps minimize user name clutter and preserve storage space on your Cloud-Delivered Firewall Management Center.

Consider the following when selecting traffic-based detection protocols:

- Obtaining user names through protocols such as AIM, POP3, and IMAP may introduce user names not relevant to your organization due to network access from contractors, visitors, and other guests.
- AIM, Oracle, and SIP logins may create extraneous user records. This occurs because these login types are not associated with any of the user metadata that the system obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your managed devices detect. Therefore, the Cloud-Delivered Firewall Management Center cannot correlate these users with other types of users.

Configuring Traffic-Based User Detection

When you enable traffic-based user detection in a network discovery rule, host discovery is automatically enabled. For more information about traffic-based detection, see The Traffic-Based Detection Identity Source, on page 9.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Users**.

Step 3 Click **Edit** ( ).

Step 4 Check the check boxes for protocols where you want to detect logins or clear check boxes for protocols where you do not want to detect logins, and choose whether you want to **Capture Failed Login Attempts**.

Step 5 Click **Save**.

What to do next

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery policy restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See Snort Restart Traffic Behavior for more information.

- Configure network discovery rules to discover users as described in Configuring Network Discovery Rules, on page 4.
- Deploy configuration changes.

Configuring Advanced Network Discovery Options

The Advanced of the network discovery policy allows you to configure policy-wide settings for what events are detected, how long discovery data is retained and how often it is updated, what vulnerability mappings are used for impact correlation, and how operating system and server identity conflicts are resolved. In addition, you can add host input sources and NetFlow exporters to allow import of data from other sources.

Note Database event limits for discovery and user activity events are set in system configuration.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) or **Add** ( ) next to the setting you want to modify:

- Data Storage Settings — Update the settings as described in Configuring Network Discovery Data Storage, on page 19.
- Event Logging Settings — Update the settings as described in Configuring Network Discovery Event Logging, on page 19.
- General Settings — Update the settings as described in Configuring Network Discovery General Settings, on page 13.
- Identity Conflict Settings — Update the settings as described in Configuring Network Discovery Identity Conflict Resolution, on page 14.
- Indications of Compromise Settings — Update the settings as described in Enabling Indications of Compromise Rules, on page 16.
- NetFlow Exporters — Update the settings as described in Adding NetFlow Exporters to a Network Discovery Policy, on page 17.
- OS and Server Identity Sources — Update the settings as described in Adding Network Discovery OS and Server Identity Sources, on page 19.

- Vulnerabilities to use for Impact Assessment — Update the settings as described in Enabling Network Discovery Vulnerability Impact Assessment, on page 15.

Step 4 Click **Save**.

What to do next

- Deploy configuration changes.

Network Discovery General Settings

The general settings control how often the system updates network maps and whether server banners are captured during discovery.

Capture Banners

Select this check box if you want the system to store header information from network traffic that advertises server vendors and versions ("banners"). This information can provide additional context to the information gathered. You can access server banners collected for hosts by accessing server details.

Update Interval

The interval at which the system updates information (such as when any of a host's IP addresses was last seen, when an application was used, or the number of hits for an application). The default setting is 3600 seconds (1 hour).

Note that setting a lower interval for update timeouts provides more accurate information in the host display, but generates more network events.

Configuring Network Discovery General Settings

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **General Settings**.

Step 4 Update the settings as described in Network Discovery General Settings, on page 13.

Step 5 Click **Save** to save the general settings.

What to do next

- Deploy configuration changes.

Network Discovery Identity Conflict Settings

The system determines which operating system and applications are running on a host by matching fingerprints for operating systems and servers against patterns in traffic. To provide the most reliable operating system and server identity information, the system collates fingerprint information from several sources.

The system uses all passive data to derive operating system identities and assign a confidence value.

By default, unless there is an identity conflict, identity data added by a scanner or third-party application overrides identity data detected by the system. You can use the Identity Sources settings to rank scanner and third-party application fingerprint sources by priority. The system retains one identity for each source, but only data from the highest priority third-party application or scanner source is used as the current identity. Note, however, that user input data overrides scanner and third-party application data regardless of priority.

An identity conflict occurs when the system detects an identity that conflicts with an existing identity that came from either the active scanner or third-party application sources listed in the Identity Sources settings or from a system user. By default, identity conflicts are not automatically resolved and you must resolve them through the host profile or by rescanning the host or re-adding new identity data to override the passive identity. However, you can set your system to automatically resolve the conflict by keeping either the passive identity or the active identity.

Generate Identity Conflict Event

Specifies whether the system generates an event when an identity conflict occurs.

Automatically Resolve Conflicts

From the **Automatically Resolve Conflicts** drop-down list, choose one of the following:

- **Disabled** if you want to force manual conflict resolution of identity conflicts
- **Identity** if you want to the system to use the passive fingerprint when an identity conflict occurs
- **Keep Active** if you want the system to use the current identity from the highest priority active source when an identity conflict occurs

Configuring Network Discovery Identity Conflict Resolution

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **Identity Conflict Settings**.

Step 4 Update the settings in the Edit Identity Conflict Settings pop-up window as described in Network Discovery Identity Conflict Settings, on page 14.

Step 5 Click **Save** to save the identity conflict settings.

What to do next

- Deploy configuration changes.

Network Discovery Vulnerability Impact Assessment Options

You can configure how the system performs impact correlation with intrusion events. Your choices are as follows:

- Check the **Use Network Discovery Vulnerability Mappings** check box if you want to use system-based vulnerability information to perform impact correlation.
- Check the **Use Third-Party Vulnerability Mappings** check box if you want to use third-party vulnerability references to perform impact correlation. For more information, see the *Firepower System Host Input API Guide*.

You can check either or both of the check boxes. If the system generates an intrusion event and the host involved in the event has servers or an operating system with vulnerabilities in the selected vulnerability mapping sets, the intrusion event is marked with the Vulnerable (level 1: red) impact icon. For any servers which do not have vendor or version information, note that you need to enable vulnerability mapping in the Cloud-Delivered Firewall Management Center configuration.

If you clear both check boxes, intrusion events will **never** be marked with the Vulnerable (level 1: red) impact icon.

Related Topics

Mapping Third-Party Vulnerabilities

Enabling Network Discovery Vulnerability Impact Assessment

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **Vulnerabilities to use for Impact Assessment**.

Step 4 Update the settings in the Edit Vulnerability Settings pop-up window as described in Network Discovery Vulnerability Impact Assessment Options, on page 15.

Step 5 Click **Save** to save the vulnerability settings.

What to do next

- Deploy configuration changes.

Indications of Compromise

The system uses IOC rules in the network discovery policy to identify a host as likely to be compromised by malicious means. When a host meets the conditions specified in these system-provided rules, the system tags

it with an *indication of compromise* (IOC). The related rules are known as *IOC rules*. Each IOC rule corresponds to one type of IOC tag. The *IOC tags* specify the nature of the likely compromise.

The Cloud-Delivered Firewall Management Center can tag the host and user involved when one of the following things occurs:

- The system correlates data gathered about your monitored network and its traffic, using intrusion, connection, Security Intelligence, and file or malware events, and determines that a potential IOC has occurred.
- The Cloud-Delivered Firewall Management Center can import IOC data from your Secure Endpoint deployments via the AMP cloud. Because this data examines activity on a host itself—such as actions taken by or on individual programs—it can provide insights into possible threats that network-only data cannot. For your convenience, the Cloud-Delivered Firewall Management Center automatically obtains any new IOC tags that Cisco develops from the AMP cloud.

To configure this feature, see Enabling Indications of Compromise Rules, on page 16.

You can also write correlation rules against host IOC data and compliance allow lists that account for IOC-tagged hosts.

To investigate and work with tagged IOCs, see *Indications of Compromise Data* and its subtopics in *Cisco Secure Firewall Management Center Administration Guide*.

Enabling Indications of Compromise Rules

For your system to detect and tag indications of compromise (IOC), you must first activate at least one IOC rule in your network discovery policy. Each IOC rule corresponds to one type of IOC tag, and all IOC rules are predefined by Cisco; you cannot create original rules. You can enable any or all rules, depending on the needs of your network and organization. For example, if hosts using software such as Microsoft Excel never appear on your monitored network, you may decide not to enable the IOC tags that pertain to Excel-based threats.

Before you begin

Because IOC rules trigger based on data provided by other components of the system and by Secure Endpoint, those components must be correctly licensed and configured for IOC rules to set IOC tags. Enable the system features associated with the IOC rules you will enable, such as intrusion detection and prevention (IPS) and Advanced Malware Protection (AMP). If an IOC rule's associated feature is not enabled, no relevant data is collected and the rule cannot trigger.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **Indications of Compromise Settings**.

Step 4 To toggle the entire IOC feature off or on, click the slider next to **Enable IOC**.

Step 5 To globally enable or disable individual IOC rules, click the slider in the rule's **Enabled** column.

Step 6 Click **Save** to save your IOC rule settings.

What to do next

- Deploy configuration changes.

Adding NetFlow Exporters to a Network Discovery Policy

Use this procedure to add NetFlow exporters. Note that you cannot delete an exporter if it is currently in use in a disocvery rule.

Before you begin

- Review prerequisites: Requirements for using NetFlow data (principle)
- Configure NetFlow exporters: NetFlow data

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Add** (+) next to **NetFlow Devices**.

Step 4 Enter the **IP Address** of the exporter.

Step 5 Click **Save**.

What to do next

- Configure a network discovery rule to monitor NetFlow traffic: Configuring Network Discovery Rules, on page 4
- Deploy configuration changes.

Network discovery data storage settings

Discovery data storage settings include the host limit and timeout settings.

When host limit reached

The number of hosts a Cloud-Delivered Firewall Management Center can monitor, and therefore store in the network map, depends on its model. The **When Host Limit Reached** option controls what happens when you detect a new host after you reach the host limit. You can:

Drop hosts

The system drops the host that has remained inactive for the longest time, then adds the new host. This is the default setting.

Don't insert new hosts

The system does not track any newly discovered hosts. Only after the host count falls below the limit does the system resume tracking new hosts. This can occur if an administrator increases the domain's host limit, manually deletes hosts from the network map, or the system identifies hosts as timed out due to inactivity.

Table 2: Reaching the host limit with multitenancy

Setting	Domain Host Limit Set?	Domain Host Limit Reached	Ancestor Domain Host Limit Reached
Drop hosts	yes	Drops oldest host in the constrained domain.	Drops the oldest host among all descendant leaf domains configured to drop hosts. If no host can be dropped, does not add the host.
	no	n/a	Drops the oldest host among all descendant leaf domains configured to drop hosts and that share the general pool.
Don't insert new hosts	yes or no	Does not add the host.	Does not add the host.

Host timeout

The amount of time that passes, in minutes, before the system drops a host from the network map due to inactivity. The default setting is 10080 minutes (one week). Individual host IP and MAC addresses can time out individually, but a host does not disappear from the network map unless all its associated addresses time out.

To avoid premature timeout of hosts, make sure that the host timeout value is longer than the update interval in the network discovery policy general settings.

Server timeout

The amount of time that passes, in minutes, before the system drops a server from the network map due to inactivity. The default setting is seven days (10,080 minutes).

To avoid premature timeout of servers, make sure that the service timeout value is longer than the update interval in the network discovery policy general settings.

Client application timeout

The amount of time that passes, in minutes, before the system drops a client from the network map due to inactivity. The default setting is 10080 minutes (one week).

To ensure optimal performance, set the client timeout value to exceed the update interval in the network discovery policy general settings.

Related Topics

Host limits

Configuring Network Discovery Data Storage

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **Network Discovery Data Storage Settings**.

Step 4 Update the settings in the Data Storage Settings dialog as described in Network discovery data storage settings, on page 17.

Step 5 Click **Save** to save the data storage settings.

What to do next

- Deploy configuration changes.

Configuring Network Discovery Event Logging

The Event Logging Settings control whether discovery and host input events are logged. If you do not log an event, you cannot retrieve it in event views or use it to trigger correlation rules.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** (✎) next to **Event Logging Settings**.

Step 4 Check or clear the check boxes next to the discovery and host input event types you want to log in the database .

Step 5 Click **Save** to save the event logging settings.

What to do next

- Deploy configuration changes.

Adding Network Discovery OS and Server Identity Sources

In Advanced of the network discovery policy, you can add new active sources or change the priority or timeout settings for existing sources.

Adding a scanner to this page does not add the full integration capabilities that exist for the Nmap scanners, but does allow integration of imported third-party application or scan results.

If you import data from a third-party application or scanner, make sure that you map vulnerabilities from the source to the vulnerabilities detected in your network.

Procedure

Step 1 Choose **Policies** > **Network Discovery**.

Step 2 Click **Advanced**.

Step 3 Click **Edit** ( ) next to **OS and Server Identity Sources**.

Step 4 To add a new source, click **Add Source**.

Step 5 Enter a **Name**.

Step 6 Choose the input source **Type** from the drop-down list:

- Choose **Scanner** if you plan to import scan results using the AddScanResult function.
- Choose **Application** if you do not plan to import scan results.

Step 7 To indicate the duration of time that should elapse between the addition of an identity to the network map by this source and the deletion of that identity, choose **Hours**, **Days**, or **Weeks** from the **Timeout** drop-down list and enter the appropriate duration.

Step 8 Optionally:

- To promote a source and cause the operating system and application identities to be used in favor of sources below it in the list, choose the source and click the up arrow.
- To demote a source and cause the operating system and application identities to be used only if there are no identities provided by sources above it in the list, choose the source and click the down arrow.
- To delete a source, click **Delete** ( ) next to the source.

Step 9 Click **Save** to save the identity source settings.

What to do next

- Deploy configuration changes.

Related Topics

Mapping Third-Party Vulnerabilities

Troubleshooting Your Network Discovery Strategy

Before you make any changes to the system's default detection capabilities, you should analyze what hosts are not being identified correctly and why, so you can decide what solution to implement.

Are Your Managed Devices Correctly Placed?

If network devices such as load balancers, proxy servers, or NAT devices reside between the managed device and the unidentified or misidentified host, place a managed device closer to the misidentified host rather than using custom fingerprinting. Cisco does not recommend using custom fingerprinting in this scenario.

Do Unidentified Operating Systems Have a Unique TCP Stack?

If the system misidentifies a host, you should investigate why the host is misidentified to help you decide between creating and activating a custom fingerprint or substituting Nmap or host input data for discovery data.



Caution If you encounter misidentified hosts, contact your support representative before creating custom fingerprints.

If a host is running an operating system that is not detected by the system by default and does not share identifying TCP stack characteristics with existing detected operating systems, you should create a custom fingerprint.

For example, if you have a customized version of Linux with a unique TCP stack that the system cannot identify, you would benefit from creating a custom fingerprint, which allows the system to identify the host and continuing monitoring it, rather than using scan results or third-party data, which require you to actively update the data yourself on an ongoing basis.

Note that many open source Linux distributions use the same kernel, and as such, the system identifies them using the Linux kernel name. If you create a custom fingerprint for a Red Hat Linux system, you may see other operating systems (such as Debian Linux, Mandrake Linux, Knoppix, and so on) identified as Red Hat Linux, because the same fingerprint matches multiple Linux distributions.

You should not use a fingerprint in every situation. For example, a modification may have been made to a host's TCP stack so that it resembles or is identical to another operating system. For example, an Apple Mac OS X host is altered, making its fingerprint identical to a Linux 2.4 host, causing the system to identify it as Linux 2.4 instead of Mac OS X. If you create a custom fingerprint for the Mac OS X host, it may cause all legitimate Linux 2.4 hosts to be erroneously identified as Mac OS X hosts. In this case, if Nmap correctly identifies the host, you could schedule regular Nmap scans for that host.

If you import data from a third-party system using host input, you must map the vendor, product, and version strings that the third party uses to describe servers and application protocols to the Cisco definitions for those products. Note that even if you map application data to system vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.

The system may reconcile data from multiple sources to determine the current identity for an operating system or application.

For Nmap data, you can schedule regular Nmap scans. For host input data, you can regularly run the Perl script for the import or the command line utility. However, note that active scan data and host input data may not be updated with the frequency of discovery data.

Can the System Identify All Applications?

If a host is correctly identified by the system but has unidentified applications, you can create a user-defined detector to provide the system with port and pattern matching information to help identify the application.

Have You Applied Patches that Fix Vulnerabilities?

If the system correctly identifies a host but does not reflect applied fixes, you can use the host input feature to import patch information. When you import patch information, you must map the fix name to a fix in the database.

Do You Want to Track Third-Party Vulnerabilities?

If you have vulnerability information from a third-party system that you want to use for impact correlation, you can map the third-party vulnerability identifiers for servers and application protocols to vulnerability identifiers in the Cisco database and then import the vulnerabilities using the host input feature. For more information on using the host input feature, see the *Firepower System Host Input API Guide*. Note that even if you map application data to system vendor and version definitions, imported third-party vulnerabilities are not used for impact assessment for clients or web applications.