



Introduction to Network Discovery and Identity

The following topics provide an introduction to network discovery and identity policies and data:

- [About Detection of Host, Application, and User Data, on page 1](#)
- [Host and Application Detection Fundamentals, on page 2](#)
- [About User Identity, on page 9](#)
- [Firepower System Host and User Limits, on page 18](#)

About Detection of Host, Application, and User Data

The Firepower System uses *network discovery* and *identity* policies to collect host, application, and user data for traffic on your network. You can use certain types of discovery and identity data to build a comprehensive map of your network assets, perform forensic analysis, behavioral profiling, access control, and mitigate and respond to the vulnerabilities and exploits to which your organization is susceptible.

Host and Application Data

Host and application data is collected by host identity sources and application detectors according to the settings in your network discovery policy. Managed devices observe traffic on the network segments you specify.

For more information, see [Host and Application Detection Fundamentals, on page 2](#).

User Data

User data is collected by user identity sources according to the settings in your network discovery and identity policies. You can use the data for user awareness and user control.

For more information, see [About User Identity, on page 9](#).

Logging discovery and identity data allows you to take advantage of many features in the Firepower System, including:

- Viewing the network map, which is a detailed representation of your network assets and topology that you can view by grouping hosts and network devices, host attributes, application protocols, or vulnerabilities.
- Performing application and user control; that is, writing access control rules using application, realm, user, user group, and ISE attribute conditions.
- Viewing host profiles, which are complete views of all the information available for your detected hosts.

- Viewing dashboards, which (among other capabilities) can provide you with an at-a-glance view of your network assets and user activity.
- Viewing detailed information on the discovery events and user activity logged by the system.
- Associating hosts and any servers or clients they are running with the exploits to which they are susceptible. This enables you to identify and mitigate vulnerabilities, evaluate the impact that intrusion events have on your network, and tune intrusion rule states so that they provide maximum protection for your network assets
- Alerting you by email, SNMP trap, or syslog when the system generates either an intrusion event with a specific impact flag, or a specific type of discovery event
- Monitoring your organization's compliance with an allow list of allowed operating systems, clients, application protocols, and protocols
- Creating correlation policies with rules that trigger and generate correlation events when the system generates discovery events or detects user activity
- Logging and using NetFlow connections, if applicable.

Related Topics

[Host Identity Sources](#)

[Application Detection](#)

[User Identity Sources](#)

Host and Application Detection Fundamentals

You can configure your network discovery policy to perform host and application detection.

For more information, see [Overview: Host Data Collection](#) and [Overview: Application Detection](#).

Passive Detection of Operating System and Host Data

Passive detection is the system's default method of populating the network map by analyzing network traffic (and any exported NetFlow data). Passive detection provides contextual information about your network assets, such as operating systems and running applications.

If traffic from a monitored host does not offer conclusive evidence of the host's operating system, the network map displays the most likely operating system. For example, a NAT device may appear to be running several operating systems because of the hosts "behind" the NAT device. To make this most-likely determination, the system uses a confidence value it assigns to each detected operating system, and the amount of corroborating data among detected operating systems.



Note The system does not consider reported "unknown" applications and operating systems in its determination.

If passive detection inaccurately identifies your network assets, consider the placement of your managed devices. You can also augment the system's passive detection capabilities with custom operating-system fingerprints and custom application detectors. Or, you can use *active detection*, which is not based on traffic

analysis, but instead allows you to directly update the network map using scan results or other information sources.

Active Detection of Operating System and Host Data

Active detection adds host information collected by active sources to network maps. For example, you can use the Nmap scanner to actively scan the hosts that you target on your network. Nmap discovers operating systems and applications on hosts.

In addition, the host input feature allows you to actively add *host input data* to network maps. There are two different categories of host input data:

- *user input data*—Data added through the Firepower System user interface. You can modify a host's operating system or application identity through this interface.
- *host import input data*—Data imported using a command line utility.

The system retains one identity for each active source. When you run an Nmap scan instance, for example, the results of the previous scan are replaced with the new scan results. However, if you run an Nmap scan and then replace those results with data from a client whose results are imported through the command line, the system retains both the identities from the Nmap results and the identities from the import client. The system then uses the priorities set in the network discovery policy to determine which active identity to use as the current identity.

Note that user input is considered one source, even if it comes from different users. As an example, if UserA sets the operating system through the host profile, and then UserB changes that definition through the host profile, the definition set by UserB is retained, and the definition set by UserA is discarded. In addition, note that user input overrides all other active sources and is used as the current identity if it exists.


Current Identities for Applications and Operating Systems

The *current identity* for an application or an operating system on a host is the identity that the system finds most likely to be correct.

The system uses the current identity for an operating system or application for the following purposes:

- to assign vulnerabilities to a host
- for impact assessment
- when evaluating correlation rules written against operating system identifications, host profile qualifications, and compliance allow lists
- for display in the Hosts and Servers table views in workflows
- for display in the host profile
- to calculate the operating system and application statistics on the Discovery Statistics page

The system uses source priorities to determine which active identity should be used as the current identity for an application or operating system.



For example, if a user sets the operating system to Windows 2003 Server on a host, Windows 2003 Server is the current identity. Attacks which target Windows 2003 Server vulnerabilities on that host are given a higher impact, and the vulnerabilities listed for that host in the host profile include Windows 2003 Server vulnerabilities.

The database may retain information from several sources for the operating system or for a particular application on a host.

The system treats an operating system or application identity as the current identity when the source for the data has the highest source priority. Possible sources have the following priority order:

1. user
2. scanner and application (set in the network discovery policy)
3. managed devices
4. NetFlow records

A new higher priority application identity will not override a current application identity if it has less detail than the current identity.

In addition, when an identity conflict occurs, the resolution of the conflict depends on settings in the network discovery policy or on your manual resolution.

Current User Identities

When the system detects multiple logins to the same host by different users, the system assumes that only one user is logged into any given host at a time, and that the current user of a host is the last authoritative user login. If only non-authoritative user logins have been logged into the host, the last non-authoritative user login is considered the current user. If multiple users are logged in through remote sessions, the last user reported by the server is the user reported to the Firepower Management Center.

When the system detects multiple logins to the same host by the same user, the system records the first time that a user logs into a specific host and disregards subsequent logins. If an individual user is the only person who logs into a specific host, the only login that the system records is the original login.


If another user logs into that host, however, the system records the new login. Then, if the original user logs in again, his or her new login is recorded.

Application and Operating System Identity Conflicts

An *identity conflict* occurs when the system reports a new passive identity that conflicts with the current active identity and previously reported passive identities. For example, the previous passive identity for an operating system is reported as Windows 2000, then an active identity of Windows XP becomes current. Next, the system detects a new passive identity of Ubuntu Linux 8.04.1. The Windows XP and the Ubuntu Linux identities are in conflict.

When an identity conflict exists for the identity of the host's operating system or one of the applications on the host, the system lists both conflicting identities as current and uses both for impact assessment until the conflict is resolved.

A user with Administrator privileges can resolve identity conflicts automatically by choosing to always use the passive identity or always use the active identity. Unless you disable automatic resolution of identity conflicts, identity conflicts are always automatically resolved.



A user with Administrator privileges can also configure the system to generate an event when an identity conflict occurs. That user can then set up a correlation policy with a correlation rule that uses an Nmap scan as a correlation response. When an event occurs, Nmap scans the host to obtain updated host operating system and application data.

Netflow Data in the Firepower System

NetFlow is a Cisco IOS application that provides statistics on packets flowing through a router. It is available on Cisco networking devices and can also be embedded in Juniper, FreeBSD, and OpenBSD devices.

When NetFlow is enabled on a network device, a database on the device (the NetFlow cache) stores records of the flows that pass through the router. A flow, called a *connection* in the Firepower System, is a sequence

of packets that represents a session between a source and destination host, using specific ports, protocol, and application protocol. The network device can be configured to export this NetFlow data. In this documentation, network devices configured in this way are called *NetFlow exporters*.

Firepower System managed devices can be configured to collect records from NetFlow exporters, generate unidirectional end-of-connection events based on the data in those records, and finally send those events to the Firepower Management Center to be logged in the connection event database. You can also configure the network discovery policy to add host and application protocol information to the database based on the information in NetFlow connections.

You can use this discovery and connection data to supplement the data gathered directly by your managed devices. This is especially useful if you have NetFlow exporters monitoring networks that your managed devices cannot monitor.

Requirements for Using NetFlow Data

Before you configure the Firepower System to analyze NetFlow data, you must enable the NetFlow feature on the routers or other NetFlow-enabled network devices you plan to use, and configure the devices to broadcast NetFlow data to a destination network where the sensing interface of a managed device is connected.

The Firepower System can parse both NetFlow version 5 and NetFlow version 9 records. NetFlow exporters **must** use one of those versions if you want to export the data to the Firepower System. In addition, the system requires that specific fields be present in the exported NetFlow templates and records. If your NetFlow exporters are using version 9, which you can customize, you **must** make sure that the exported templates and records contain the following fields, in any order:

- IN_BYTES (1)
- IN_PKTS (2)
- PROTOCOL (4)
- TCP_FLAGS (6)
- L4_SRC_PORT (7)
- IPV4_SRC_ADDR (8)
- L4_DST_PORT (11)
- IPV4_DST_ADDR (12)
- LAST_SWITCHED (21)
- FIRST_SWITCHED (22)
- IPV6_SRC_ADDR (27)
- IPV6_DST_ADDR (28)

Because the Firepower System uses managed devices to analyze NetFlow data, your deployment must include at least one managed device that can monitor NetFlow exporters. At least one sensing interface on that managed device must be connected to a network where it can collect the exported NetFlow data. Because the sensing interfaces on managed devices do not usually have IP addresses, the system does not support the direct collection of NetFlow records.

Note that the Sampled NetFlow feature available on some network devices collects NetFlow statistics on only a subset of packets that pass through the devices. Although enabling this feature can improve CPU utilization on the network device, it may affect the NetFlow data you are collecting for analysis by the Firepower System.

Differences between NetFlow and Managed Device Data

The traffic represented by NetFlow data is not directly analysed. Instead, the exported NetFlow records are converted into connection logs and host and application protocol data.

As a result, there are several differences between converted NetFlow data and the discovery and connection data gathered directly by your managed devices. You should keep these differences in mind when performing analysis that requires:

- Statistics on the number of detected connections
- Operating system and other host-related information (including vulnerabilities)
- Application data, including client information, web application information, and vendor and version server information
- Knowing which host in a connection is the initiator and which is the responder

Network Discovery Policy versus Access Control Policy

You configure NetFlow data collection, including connection logging, using rules in the network discovery policy. Contrast this with connection logging for connections detected by managed devices, which you configure per access control rule.

Types of Connection Events

Because NetFlow data collection is linked to networks rather than access control rules, you do not have granular control over which NetFlow connections the system logs.

NetFlow data cannot generate Security Intelligence events.

NetFlow-based connection events can be stored in the connection event database only; you cannot send them to the system log or an SNMP trap server.

Number of Connection Events Generated Per Monitored Session

For connections detected directly by managed devices, you can configure the access control rule to log a bidirectional connection event at the beginning or end of a connection, or both.

In contrast, because exported NetFlow records contain unidirectional connection data, the system generates at least two connection events for each NetFlow record it processes. This also means that a summary's connection count is incremented by two for every connection based on NetFlow data, providing an inflated count of the number of connections that are actually occurring on your network.

Because the NetFlow exporter outputs records at a fixed interval even if a connection is still ongoing, long-running sessions can result in multiple exported records, each of which generates a connection event. For example, if the NetFlow exporter exports every five minutes, and a particular connection lasts twelve minutes, the system generates six connection events for that session:

- One pair of events for the first five minutes
- One pair for the second five minutes

- A final pair when the connection is terminated

Host and Operating System Data

Hosts added to the network map from NetFlow data do not have operating system, NetBIOS, or host type (host vs network device) information. You can, however, manually set a host's operating system identity using the host input feature.

Application Data

For connections detected directly by managed devices, the system can identify application protocols, clients, and web applications by examining the packets in the connection.

When the NetFlow records are processed, the system uses a port correlation in `/etc/sf/services` to extrapolate application protocol identity. However, there is no vendor or version information for those application protocols, nor do connection logs contain information on client or web applications used in the session. You can, however, manually provide this information using the host input feature.

Note that a simple port correlation means that application protocols running on non-standard ports may be unidentified or misidentified. Additionally, if no correlation exists, the system marks the application protocol as `unknown` in connection logs.

Vulnerability Mappings

The system cannot map vulnerabilities to hosts monitored by NetFlow exporters, unless you use the host input feature to manually set either a host's operating system identity or an application protocol identity. Note that because there is no client information in NetFlow connections, you cannot associate client vulnerabilities with hosts created from NetFlow data.

Initiator and Responder Information in Connections

For connections detected directly by managed devices, the system can identify which host is the initiator, or source, and which is the responder, or destination. However, NetFlow data does not contain initiator or responder information.

When the system processes NetFlow records, it uses an algorithm to determine this information based on the ports each host is using, and whether those ports are well-known:

- If both or neither port being used is a well-known port, the system considers the host using the lower-number port to be the responder.
- If only one of the hosts is using a well-known port, the system considers that host to be the responder.

For this purpose, a well-known port is any port that is either numbered from 1 to 1023, or that contains application protocol information in `/etc/sf/services` on the managed device.

In addition, for connections detected directly by managed devices, the system records two byte counts in the corresponding connection event:

- The **Initiator Bytes** field records bytes sent.
- The **Responder Bytes** field records bytes received.

Connection events based on unidirectional NetFlow records contain only one byte count, which the system assigns to either **Initiator Bytes** or **Responder Bytes**, depending on the port-based algorithm. The system

sets the other field to 0. Note that if you are viewing connection summaries (aggregated connection data) of NetFlow records, both fields may be populated.

NetFlow-only Connection Event Fields

A small number of fields are present only in connection events generated from NetFlow records; see [Information Available in Connection Event Fields](#).

Related Topics

[Information Available in Connection Event Fields](#)

About User Identity

User identity information can help you to identify the source of policy breaches, attacks, or network vulnerabilities, and trace them to specific users. For example, you could determine:

- Who owns the host targeted by an intrusion event that has a Vulnerable (level 1: red) impact level.
- Who initiated an internal attack or portscan.
- Who is attempting unauthorized access to a specified host.
- Who is consuming an unreasonable amount of bandwidth.
- Who has not applied critical operating system updates.
- Who is using instant messaging software or peer-to-peer file-sharing applications in violation of company policy.
- Who is associated with each indication of compromise on your network.

Armed with this information, you can use other features of the Firepower System to mitigate risk, perform access control, and take action to protect others from disruption. These capabilities also significantly improve audit controls and enhance regulatory compliance.

After you configure user identity sources to gather user data, you can perform user awareness and user control.

Video [YouTube videos for configuring identity](#).

Related Topics

[Identity Terminology](#), on page 9

[Identity Deployments](#), on page 13

[About User Identity Sources](#), on page 10

[How to Set Up an Identity Policy](#), on page 14

Identity Terminology

This topic discusses common terminology for user identity and user control.

User awareness

Identifying users on your network using *identity sources* (such as or TS Agent). User awareness enables you to identify users from both *authoritative* (such as Active Directory) and *non-authoritative*

(application-based) sources. To use Active Directory as an identity source, you must configure a realm and directory. For more information, see [About User Identity Sources, on page 10](#).

User control

Configuring an *identity policy* that you associate with an *access control policy*. (The identity policy is then referred to as an access control *subpolicy*.) The identity policy specifies the identity source and, optionally, users and groups belonging to that source.

By associating the identity policy with an access control policy, you determine whether to monitor, trust, block, or allow users or user activity in traffic on your network. For more information, see [Access Control Policies](#).

Authoritative identity sources

A trusted server validated the user login (for example, Active Directory). You can use the data obtained from authoritative logins to perform user awareness and user control. Authoritative user logins are obtained from passive and active authentications:

- *Passive authentications* occur when a user authenticates through an external source. ISE/ISE-PIC and the TS Agent are the passive authentication methods supported by the Firepower System.
- *Active authentications* occur when a user authenticates through preconfigured managed devices. Captive portal and Remote Access VPN are the active authentication methods supported by the Firepower System.

Non-authoritative identity sources

An unknown or untrusted server validated the user login. Traffic-based detection is the only non-authoritative identity source supported by the Firepower System. You can use the data obtained from non-authoritative logins to perform user awareness.

About User Identity Sources

The following table provides a brief overview of the user identity sources supported by the Firepower System. Each identity source provides a store of users for user awareness. These users can then be controlled with identity and access control policies.

User Identity Source	Policy	Server Requirements	Type	Authentication Type	User Awareness?	User Control?	For more information, see...
ISE/ISE-PIC	Identity	Microsoft Active Directory	Authoritative logins	Passive	Yes	Yes	The ISE/ISE-PIC Identity Source
TS Agent	Identity	Microsoft Windows Terminal Server	Authoritative logins	Passive	Yes	Yes	The Terminal Services (TS) Agent Identity Source

User Identity Source	Policy	Server Requirements	Type	Authentication Type	User Awareness?	User Control?	For more information, see...
Captive portal	Identity	Microsoft Active Directory	Authoritative logins	Active	Yes	Yes	The Captive Portal Identity Source
Remote Access VPN	Identity	OpenLDAP or Microsoft Active Directory	Authoritative logins	Active	Yes	Yes	The Remote Access VPN Identity Source
	Identity	RADIUS	Authoritative logins	Active	Yes	No	
Traffic-based detection	Network discovery	n/a	Non-authoritative logins	n/a	Yes	No	The Traffic-Based Detection Identity Source



Note The Cisco Firepower User Agent is not supported in and cannot be enabled in FMC version 6.7.

Consider the following when selecting identity sources to deploy:

- You must use traffic-based detection for non-LDAP user logins.
- You must use traffic-based detection or captive portal to record failed login or authentication activity. A failed login or authentication attempt does not add a new user to the list of users in the database.
- The captive portal identity source requires a managed device with a routed interface. You *cannot* use an inline (also referred to as tap mode) interface with captive portal.

Data from those identity sources is stored in the Firepower Management Center's users database and the user activity database. You can configure Firepower Management Center-server user downloads to automatically and regularly download new user data to your databases.

After you configure identity rules using the desired identity source, you must associate each rule with an access control policy and deploy the policy to managed devices for the policy to have any effect. For more information about access control policies and deployment, see [User, Realm, and ISE Attribute Conditions \(User Control\)](#).

For general information about user identity in the Firepower System, see [About User Identity, on page 9](#).

Video icon [YouTube videos for configuring identity sources](#).

Best Practices for User Identity

We recommend you review the following information before you set up identity policies.

- Know user limits

- Create one realm per AD domain, something about trust
- Health monitor
- Use latest version of ISE/ISE-PIC, two types of remediation
- User agent support drops in 6.7
- Captive portal requires routed interface, several individual tasks
- See TS Agent troubleshooting

Active Directory, LDAP, and realms

The Firepower System supports either Active Directory or LDAP for user awareness and control. The association between an Active Directory or LDAP repository and the FMC is referred to as a *realm*. You should create one realm per LDAP server or Active Directory domain. For details about which versions are supported, see [Supported Servers for Realms](#).

The only user identity source supported by LDAP is captive portal. To use other identity sources (with the exception of ISE/ISE-PIC), you must use Active Directory.

For Active Directory only:

- Create one *directory* per domain controller.

For details, see [Create a Realm and Realm Directory](#)

- Users and groups in trust relationships between two domains are supported provided you add all Active Directory domains and domain controllers as realms and directories, respectively.

For more information, see [Realms and Trusted Domains](#).

Health monitor

The FMC health monitor provides valuable information about the status of various FMC functions, including:

- User/realm mismatches
- Short memory usage
- ISE connection status

For more information about health modules, see [Health Modules](#).

To set up policies to monitor health modules, see [Creating Health Policies](#).

Device-specific user limits

Every physical or virtual FMC device has limits to the number of users that can be downloaded. When the user limit is reached, the FMC can run out of memory and can function unreliably as a result.

User limits are discussed in [Firepower System User Limit, on page 19](#).

If you use the ISE/ISE-PIC identity source, you can optionally limit the subnets the FMC monitors to reduce memory usage using identity mapping filters as discussed in [Create an Identity Policy](#).

Use the latest version of ISE/ISE-PIC

If you expect to use the ISE/ISE-PIC identity source, we strongly recommend you always use the latest version to make sure you get the latest features and bug fixes.

pxGrid 2.0 (which is used by version 2.6 patch 6 or later; or 2.7 patch 2 or later) also changes the remediation used by ISE/ISE-PIC from Endpoint Protection Service (EPS) to Adaptive Network Control (ANC). If you upgrade ISE/ISE-PIC, you must migrate your mediation policies from EPS to ANC.

More information about using ISE/ISE-PIC can be found in [ISE/ISE-PIC Guidelines and Limitations](#).

To set up the ISE/ISE-PIC identity source, see [How to Configure ISE/ISE-PIC for User Control](#).

Captive portal information

Captive portal is the only user identity source for which you can use either LDAP or Active Directory. In addition, your managed devices must be configured to use a routed interface.

Additional guidelines can be found in [Captive Portal Guidelines and Limitations](#).

Setting up captive portal requires performing several independent tasks. For more information, see [How to Configure the Captive Portal for User Control](#).

TS Agent information

The TS Agent user identity source is required to identify user sessions on a Windows Terminal Server. The TS Agent software must be installed on the Terminal Server machine as discussed in the *Cisco Terminal Services (TS) Agent Guide*. In addition, you must synchronize the time on your TS Agent server with the time on the Firepower Management Center.

TS Agent data is visible in the Users, User Activity, and Connection Event tables and can be used for user awareness and user control.

For more information, see [TS Agent Guidelines](#).

Associate the identity policy with an access control policy

After you configure your realm, directory, and user identity source, you must set up identity rules in an identity policy. To make the policy effective, you must associate the identity policy with an access control policy.

For more information about creating an identity policy, see [Create an Identity Policy](#).

For more information about creating identity rules, see [Create an Identity Rule](#).

To associate an identity policy with an access control policy, see [Associating Other Policies with Access Control](#).

User agent deprecation and end of support by FMC


For more information, see [End-of-Life and End-of-Support for the Cisco Firepower User Agent](#).

Identity Deployments

When the system detects user data from a user login, from any identity source, the user from the login is checked against the list of users in the Firepower Management Center user database. If the login user matches an existing user, the data from the login is assigned to the user. Logins that do not match existing users cause a new user to be created, unless the login is in SMTP traffic. Non-matching logins in SMTP traffic are discarded.

The group to which the user belongs is associated with the user as soon as the user is seen by the Firepower Management Center.

The following diagram illustrates how the Firepower System collects and stores user data:



How to Set Up an Identity Policy

This topic provides a high-level overview of setting up an identity policy using any available user identity source: TS Agent, ISE/ISE-PIC, captive portal, or Remote Access VPN.

Procedure

	Command or Action	Purpose
Step 1	(Optional.) Create a realm and directory, one realm for every domain in the Active Directory trust and one directory for every domain controller. Only users and groups that have corresponding FMC realms and directories can be used in identity policies..	Creating a realm and realm directory is optional if any of the following are true: • You use SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions. • You are using an identity policy only to filter network traffic. The <i>realm</i> is a trusted user and group store, typically a Microsoft Active Directory domain. You must create a realm only for a domain that contains users you wish to use in user control. The Firepower Management Center downloads users and groups at intervals you specify. You can include or exclude users and groups from being downloaded.

	Command or Action	Purpose
		See Create a Realm and Realm Directory. For details about the options to create a realm, see Realm Fields. A <i>directory</i> is an Active Directory domain controller that organizes information about a computer network's users and network shares. An Active Directory controller provides Directory Services for the realm. Active Directory distributes user and group objects across multiple domain controllers, which are peers that propagate local changes between each other by the use of Directory Services. For more information, see the Active Directory technical specification glossary on MSDN. You can specify more than one directory for a realm, in which case each domain controller is queried in the order listed on the realm's Directory tab page to match user and group credentials for user control. Note Configuring a realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions.
Step 2	Synchronize users and groups from the realm.	To be able to control users and groups, you must synchronize them with the Firepower Management Center. You can synchronize them with users and groups whenever you want or you can configure the system to synchronize them with them at a specified interval. When you synchronize users and groups, you can specify exceptions; for example, you can exclude the Engineering group from all user control for that realm, or you can exclude the user joe.smith from user controls that apply to the Engineering group. See Synchronize Users and Groups
Step 3	(Optional.) Create a realm sequence.	A realm sequence is an ordered list of realms that, when used in an identity policy, causes the Firepower System to search the realms in the specified order to find users to match the rule. See Create a Realm Sequence.
Step 4	Create a method to retrieve user and group data (the <i>identity source</i>).	Set up an identity source with its unique configuration to be able to control users and

	Command or Action	Purpose
		groups using data stored in the realm. Identity sources include TS Agent, captive portal, or Remote VPN. See one of the following: • How to Configure the Captive Portal for User Control • Configure ISE/ISE-PIC for User Control • Configure RA VPN for User Control
Step 5	Create an identity policy.	An identity policy contains one or more identity rules, optionally organized in categories. See Create an Identity Policy. Note Configuring a realm or realm sequence is optional if you plan to configure SGT ISE attribute conditions but not user, group, realm, Endpoint Location, or Endpoint Profile conditions; or if you use your identity policy only to filter network traffic.
Step 6	Create one or more identity rules.	Identity rules enable you to specify a number of matching criteria, including the type of authentication, network zones, networks or geolocation, realms, realm sequences, and so on. See Create an Identity Rule.
Step 7	Associate your identity policy with an access control policy.	An access control policy filters and optionally inspects traffic. An identity policy must be associated with an access control policy to have any effect. See Associating Other Policies with Access Control.
Step 8	Deploy the access control policy to at least one managed device.	To use your policy to control user activity, the policy must be deployed to the managed devices to which clients connect. See Deploy Configuration Changes.
Step 9	Monitor user activity.	View a list of active sessions collected by user identity sources or a list of user information collected by user identity sources. See Using Workflows. An identity policy is not required if all of the following are true: • You use the ISE/ISE-PIC identity source. • You do not use users or groups in access control policies.

	Command or Action	Purpose
		You use Security Group Tags (SGT) in access control policies. For more information, see ISE SGT vs Custom SGT Rule Conditions.

Related Topics

[Configuring Traffic-Based User Detection](#)

The User Activity Database

The user activity database on the Firepower Management Center contains records of user activity on your network detected or reported by all of your configured identity sources. The system logs events in the following circumstances:

- When it detects individual logins or logoffs.
- When it detects a new user.
- When a system administrator manually delete a user.
- When the system detects a user that is not in the database, but cannot add the user because you have reached your user limit.
- When you resolve an indication of compromise associated with a user, or enable or disable indication of compromise rules for a user.

**Note**

If the TS Agent monitors the same users as another passive authentication identity source (such as the ISE/ISE-PIC), the Firepower Management Center prioritizes the TS Agent data. If the TS Agent and another passive source report identical activity from the same IP address, only the TS Agent data is logged to the Firepower Management Center.

You can view user activity detected by the system using the Firepower Management Center web interface. (**Analysis > Users > User Activity**).

The Users Database

The users database on the Firepower Management Center contains a record for each user detected or reported by all of your configured identity sources. You can use data obtained from an authoritative source for user control.

See [About User Identity Sources, on page 10](#) for more information about the supported non-authoritative and authoritative identity sources.

The total number of users the Firepower Management Center can store depends on the Firepower Management Center model, as described in [Firepower System User Limit, on page 19](#). After the user limit is reached, the system prioritizes previously-undetected user data based on its identity source, as follows:

- If the new user is from a non-authoritative identity source, the system does not add the user to the database. To allow new users to be added, you must delete users manually or with a database purge.

- If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period and adds the new user to the database.

If an identity source is configured to exclude specific user names, user activity data for those user names are not reported to the Firepower Management Center. These excluded user names remain in the database, but are not associated with IP addresses. For more information about the type of data stored by the system, see [User Data](#).

If you have Firepower Management Center high availability configured and the primary fails, no logins reported by a captive portal, ISE/ISE-PIC, TS Agent, or Remote Access VPN device can be identified during failover downtime, even if the users were previously seen and downloaded to the Firepower Management Center. The unidentified users are logged as Unknown users on the Firepower Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.



Note If the TS Agent monitors the same users as another passive authentication identity source (ISE/ISE-PIC), the Firepower Management Center prioritizes the TS Agent data. If the TS Agent and another passive source report identical activity from the same IP address, only the TS Agent data is logged to the Firepower Management Center.

When the system detects a new user session, the user session data remains in the users database until one of the following occurs:

- A user on the Firepower Management Center manually deletes the user session.
- An identity source reports the logoff of that user session.
- 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.

Firepower System Host and User Limits

Your Firepower Management Center model determines how many individual hosts you can monitor with your deployment, as well as how many users you can monitor and use to perform user control.

Related Topics

[Purging Data from the FMC Database](#)

Firepower System Host Limit

The system adds a host to the network map when it detects activity associated with an IP address in your monitored network (as defined in your network discovery policy). The number of hosts a Firepower Management Center can monitor, and therefore store in the network map, depends on its model.

Table 1: Host Limits by Firepower Management Center Model

FMC Model	Hosts
MC1000	50,000
MC1600	50,000

FMC Model	Hosts
MC2500	150,000
MC2600	150,000
MC4500	600,000
MC4600	600,000
virtual	50,000

You cannot view contextual data for hosts not in the network map. However, you can perform access control. For example, you can perform application control on traffic to and from a host not in the network map, even though you cannot use a compliance allow list to monitor the host's network compliance.



Note The system counts MAC-only hosts separately from hosts identified by both IP addresses and MAC addresses. All IP addresses associated with a host are counted together as one host.

Reaching the Host Limit and Deleting Hosts

The network discovery policy controls what happens when you detect a new host after you reach the host limit; you can drop the new host, or replace the host that has been inactive for the longest time. You can also set the period after which the system removes a host from the network map due to inactivity. Although you can manually delete a host, an entire subnet, or all of your hosts from the network map, if the system detects activity associated with a deleted host, it re-adds the host.

In a multidomain deployment, each leaf domain has its own network discovery policy. Therefore, each leaf domain governs its own behavior when the system discovers a new host.

Related Topics

[Domain Properties](#)

[Network Discovery Data Storage Settings](#)

Firepower System User Limit

Your Firepower Management Center model determines how many individual users you can monitor. The user is added to the Firepower Management Center user database when:

- The user is downloaded from a realm.
- A captive portal or RA-VPN user logs in.
- A user is detected from any identity source (for example, TS Agent).

Only authoritative users are available for user control with access control policies.

Note the following:

- The maximum number of *downloaded* users depends on your FMC model.
- The maximum number of *concurrent* user sessions (that is, logins) depends on your FTD model. A single user can have multiple sessions from different unique IP addresses.



Note The Firepower System downloads all user sessions to all FTD devices. If you have devices with different user concurrent user session limits, the FTD with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your FMC manages an FTD 4110 and a 4120, the 4110 reports health warnings when the number of concurrent user sessions approaches its maximum of 64,000.)

Table 2: Maximum Concurrent User Login Limits by Firepower Threat Defense Model

FTD Model	Maximum Concurrent User Logins
FTDv 5, 10, 20, 30, 50 (any supported hypervisor)	64,000
ASA 5508-X, 5516-X	64,000
Firepower 1010, 1120, 1140, 1150 Firepower 2110, 2120, 2130, 4110	64,000
Firepower 2140, 4112, 4115, 4120, 4125	150,000
Firepower 4140, 4145, 4150, 9300	300,000
ASA FirePOWER Services Module	2,000

Table 3: Maximum Downloaded Users by Firepower Management Center Model¹

FMC Model	Maximum Downloaded Users
FMC1000	50,000
FMC1600	50,000
FMC2500	150,000
FMC2600	150,000
FMC4500	600,000
FMC4600	600,000
FMCv (any supported hypervisor)	50,000
FMCv 300 (any supported hypervisor)	150,000

¹—FMC models are subject to end of life and end of sale. For more information, see [End-Of-Life and End-Of-Sale Notices](#).

When the system detects a new, previously-undetected user after the limit has been reached, it prioritizes user data based on their identity source:

- If the new user is from a non-authoritative source, the system does not add the non-authoritative user to the database. To allow new users to be added, you must delete users manually or purge the database.
- If the new user is from an authoritative identity source, the system deletes the non-authoritative user who has remained inactive for the longest period of and adds the new authoritative user to the database.

If there are only authoritative users, the system deletes the authoritative user who has remained inactive for the longest period of time and adds the new user to the database.

Troubleshooting information can be found in [Troubleshoot User Control](#).



Tip Note that if you are using traffic-based detection, you can restrict user logging by protocol to help minimize username clutter and preserve space in the database. For example, you could prevent the system from adding users discovered in AIM, POP3, and IMAP traffic because you know it is traffic from specific contractors or visitors you do not want to monitor.
