



User Identity Overview

The following topics discuss user identity:

- [About User Identity](#), on page 1
- [Host and User Limits](#), on page 14

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 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.

For more information about identity sources, see [About User Identity Sources](#), on page 2.

Related Topics

- [Identity Terminology](#), on page 2
- [About User Identity Sources](#), on page 2
- [Identity Deployments](#), on page 5
- [How to Set Up an Identity Policy](#), on page 10

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 2](#).

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 with Cisco ISE.
- *Active authentications* occur when a user authenticates through preconfigured managed devices. Captive portal is another name for active authentication. Remote Access VPN is another authentication method supported by the system. Active authentication generally uses the same user repositories as passive authentication (the exceptions being ISE/ISE-PIC, TS Agent, and the Passive Identity Agent, which are passive only).

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

User Identity Source	Server Requirements	Login Type	Authentication Type	User Control	For more, see...
Captive portal	OpenLDAP Microsoft Active Directory	Authoritative	Active	Yes	The Captive Portal Identity Source
Remote Access VPN	OpenLDAP or Microsoft Active Directory	Authoritative	Active	Yes	The Remote Access VPN Identity Source
	RADIUS	Authoritative	Active	No, awareness only	
Traffic-based detection (Configured in the network discovery policy.)	—	Non-authoritative	—	No, awareness only	The Traffic-Based Detection Identity Source

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 management center's users database and the user activity database. You can configure 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 [Associating Other Policies with Access Control](#).

For general information about user identity, see [About User Identity, on page 1](#).

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
- 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

Active Directory, LDAP, and realms

The system supports either Active Directory or LDAP for user awareness and control. The association between an Active Directory or LDAP repository and the management center 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 an LDAP Realm or an Active Directory 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 management center health monitor provides valuable information about the status of various management center functions, including:

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

For more information about health modules, see *Health Modules* in the [Cisco Secure Firewall Management Center Administration Guide](#).

To set up policies to monitor health modules, see *Creating Health Policies* in the [Cisco Secure Firewall Management Center Administration Guide](#).

Device-specific user limits

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

User limits are discussed in [User Limits for Microsoft Active Directory, on page 15](#).

If you use the ISE/ISE-PIC identity source, you can optionally limit the subnets the management center 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 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](#).

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 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.

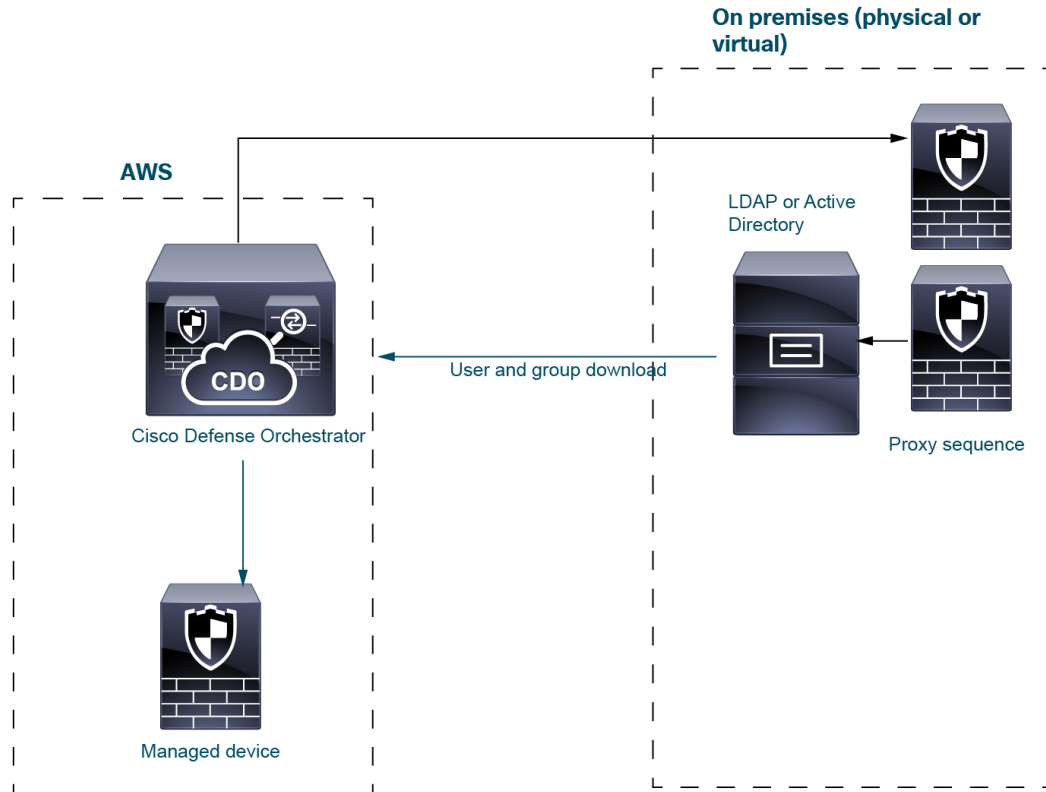
The group to which the user belongs is associated with the user as soon as the user is seen by the management center.

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



Sample identity deployments

The sample deployments discussed in this section are based on the system shown in the following figure.



In the preceding figure, CDO and one managed device are deployed to AWS and the other devices are located on premises. These devices can be either physical or virtual; they just need to be able to communicate with each other.

The two on-premises managed devices are intended to be used as a proxy sequence. You must add those devices to CDO as well.

A *proxy sequence* is one or more managed devices that can be used to communicate with an LDAP, Active Directory, or ISE/ISE-PIC server. It is necessary only if Cisco Defense Orchestrator (CDO) cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, CDO might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

LDAP or Active Directory are needed only for TS Agent and captive portal, as the following paragraphs explain.

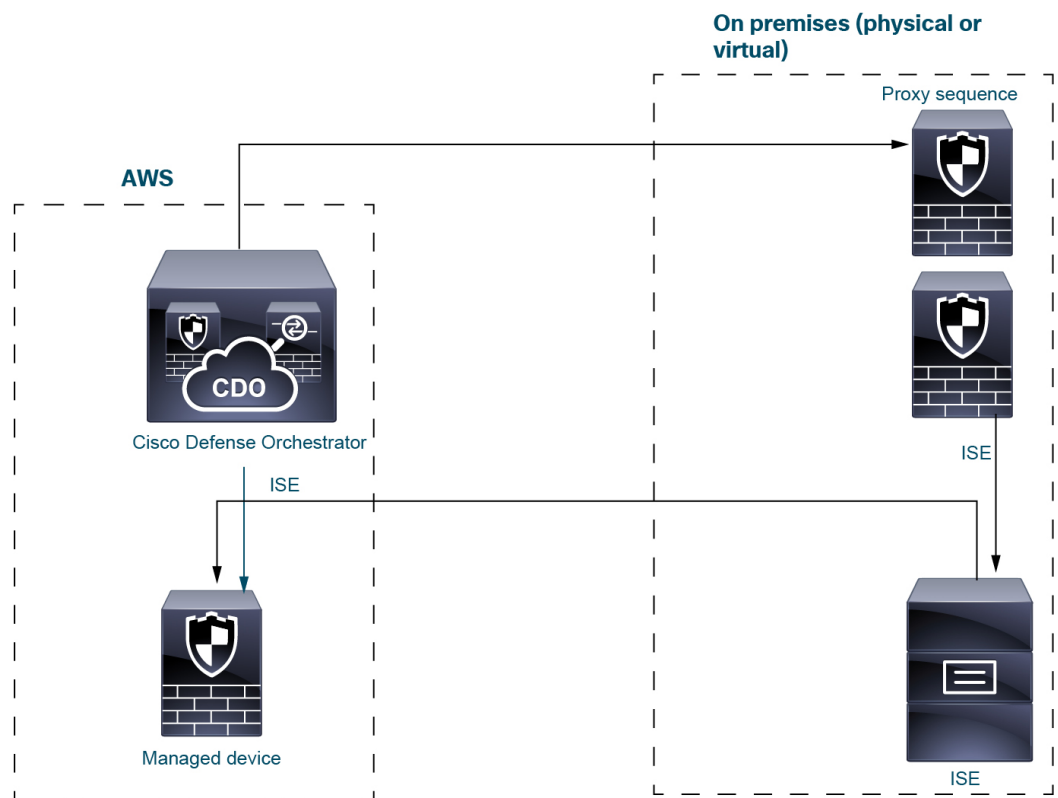
For more information about setting up a system like this, see [How to Set Up an Identity Policy, on page 10](#).

ISE/ISE-PIC identity source

When you deploy the ISE/ISE-PIC identity source, CDO contacts the proxy sequence if CDO cannot contact the ISE/ISE-PIC server directly. Users, groups, and subscriptions are sent from the ISE/ISE-PIC server to the managed device in AWS.

You can optionally have an LDAP server in an ISE/ISE-PIC deployment but because it's optional, it isn't shown in the following figure.

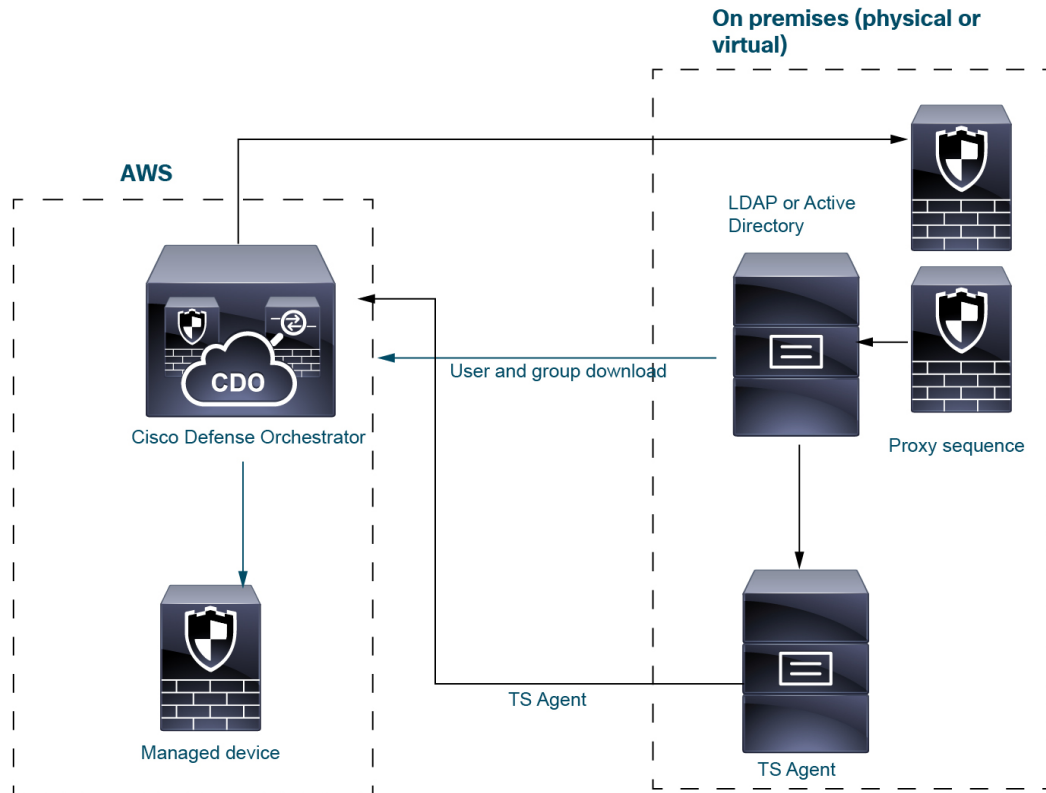
For more information about ISE/ISE-PIC, see [The ISE/ISE-PIC Identity Source](#).



TS Agent identity source

The Terminal Services (TS) Agent software runs on a Microsoft Server and sends CDO user information based on the port range the users log in to the server with. TS Agent gets user identity information from LDAP or Active Directory and sends it to CDO.

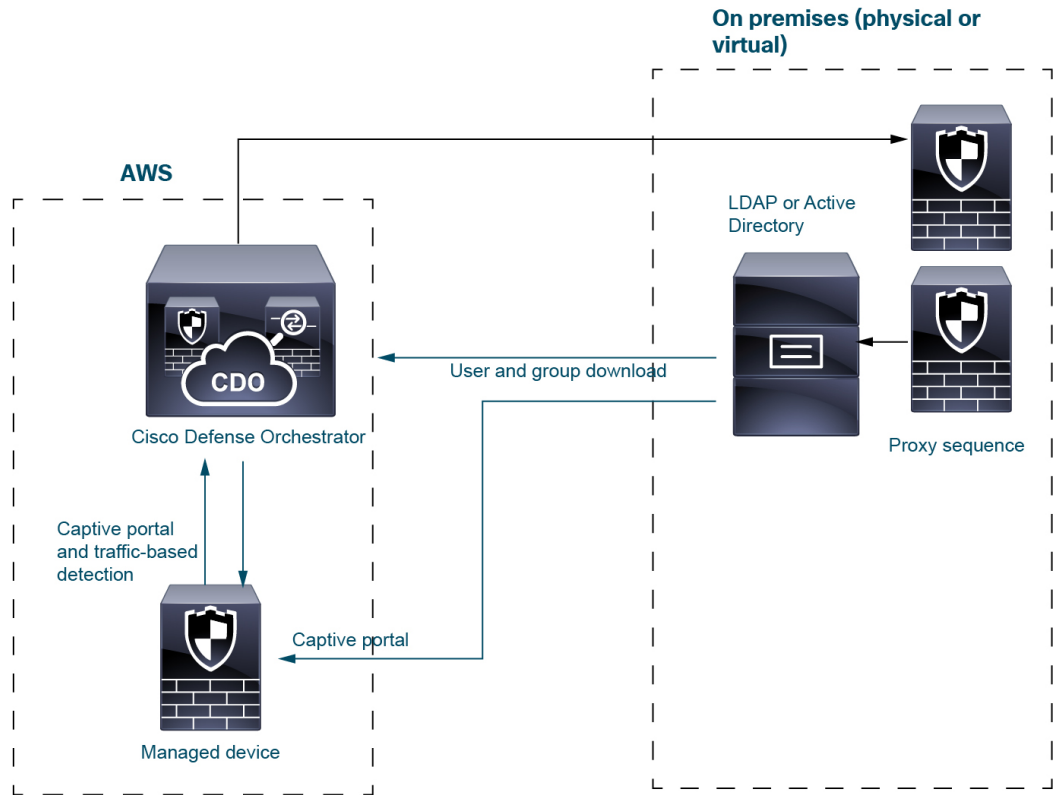
For more information about the TS Agent identity source, see [The Terminal Services \(TS\) Agent Identity Source](#).



Captive portal identity source

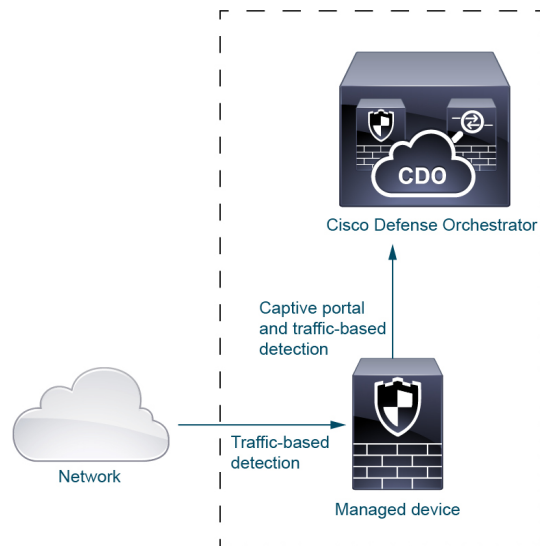
Captive portal is the only identity source that supports LDAP in addition to Active Directory. The captive portal identity source is triggered when a user tries to access network resources using the Managed device in AWS, using either an IP address or host name. Captive portal gets user information from LDAP or Active Directory using the proxy sequence and sends user information to CDO.

For more information about the captive portal identity source, see [The Captive Portal Identity Source](#).



Traffic-based detection

Traffic-based detection is designed only to detect applications on the network and therefore has no need for a user repository like Active Directory or for a proxy sequence. For more information about it, see [About Detection of Host, Application, and 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, Microsoft Azure AD (SAML), captive portal, or Remote Access VPN.

Procedure

	Command or Action	Purpose
Step 1	<p>(Optional.) Create a realm and directory, one realm for every domain in the forest that contain users you want to use in user control. Also create one directory for every domain controller. Only users and groups that have corresponding management center realms and directories can be used in identity policies..</p>	<p>Creating a realm, realm directory is optional if any of the following are true:</p> <ul style="list-style-type: none"> • 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. <p>The <i>realm</i> is a trusted user and group store, typically a Microsoft Active Directory repository. The management center downloads users and groups at intervals you specify. You can include or exclude users and groups from being downloaded.</p> <p>See Create an LDAP Realm or an Active Directory Realm and Realm Directory. For details about the options to create a realm, see Realm Fields.</p> <p>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.</p> <p>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.</p> <p>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.</p>

	Command or Action	Purpose
Step 2	Synchronize users and groups from the realm.	<p>To be able to control users and groups, you must synchronize them with the management center. You can synchronize them with users and groups whenever you want or you can configure the system to synchronize them at a specified interval.</p> <p>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.</p> <p>See Synchronize Users and Groups</p>
Step 3	(Optional.) Create a realm sequence.	<p>A realm sequence is an ordered list of realms that, when used in an identity policy, causes the system to search the realms in the specified order to find users to match the rule. See Create a Realm Sequence.</p>
Step 4	Create a method to retrieve user and group data (the <i>identity source</i>).	<p>Set up an identity source with its unique configuration to be able to control users and groups using data stored in the realm. Identity sources include TS Agent, captive portal, or Remote VPN. See one of the following:</p> <ul style="list-style-type: none"> • How to Configure the Captive Portal for User Control • Ways to Configure the Identity Services Engine Identity Source • Configure RA VPN for User Control
Step 5	Create an identity policy.	<p>An identity policy contains one or more identity rules, optionally organized in categories. See Create an Identity Policy.</p> <p>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.</p>
Step 6	Create one or more identity rules.	<p>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.</p>

	Command or Action	Purpose
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 <i>Using Workflows</i> in the Cisco Secure Firewall Management Center Administration Guide . An identity policy is not required if all of the following are true: <ul style="list-style-type: none"> • You use the ISE/ISE-PIC identity source. • You do not use users or groups in access control policies. • 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 Secure Firewall 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 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 management center.

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

The Users Database

The users database on the Secure Firewall 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 2](#) for more information about the supported non-authoritative and authoritative identity sources.

The total number of users the Secure Firewall Management Center can store depends on the Secure Firewall Management Center model, as described in [User Limits for Microsoft Active Directory, on page 15](#). 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 Secure Firewall 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* in the [Cisco Secure Firewall Management Center Administration Guide](#).

If you have 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 management center. The unidentified users are logged as Unknown users on the 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 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 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 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.

Host and User Limits

Your Secure Firewall 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.

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 Secure Firewall Management Center can monitor, and therefore store in the network map, depends on its model.

Table 1: Host Limits by Secure Firewall Management Center Model

Management Center Model	Hosts
MC1000	50,000
MC1600	50,000
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.

Related Topics[Network Discovery Data Storage Settings](#)

User Limits for Microsoft Active Directory

About user limits

Your management center model determines how many individual users you can monitor. The user is added to the 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 management center model.
- The maximum number of *concurrent* user sessions (that is, logins) depends on your managed device model. A single user can have multiple sessions from different unique IP addresses.



Note The system downloads all user sessions to all threat defense devices. If you have devices with different user concurrent user session limits, the threat defense with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your management center manages a Firepower 2110 and a 4125, the 2110 reports health warnings when the number of concurrent user sessions approaches its maximum of 64,000.)

User Limits for Microsoft Active Directory

Table 2: Maximum Concurrent User Login Limits by Threat Defense

Threat Defense Model	Maximum Concurrent User Logins per Realm
Threat Defense Virtual 5, 10, 20, 30, 50 (any supported hypervisor)	64,000
Firepower 1010, 1120, 1140, 1150 Secure Firewall 3105, 3110, 3120	64,000
Secure Firewall 3130, 3140 Firepower 4112, 4115, 4125	150,000
Firepower 4145 Firepower 9300	300,000

Threat Defense Model	Maximum Concurrent User Logins per Realm
Secure Firewall 4215	300,000
Secure Firewall 4225, 4245	315,000

User limits are applied per Microsoft Active Directory realm. That is, if you attempt to download more than the maximum users in a *single* realm, the download stops after that many users and a health alert is displayed. If, however, you attempt to download more than the maximum number of users spread across *different* realms, the download succeeds (unless any one realm has more than 150,000 users, in which case the download fails for that realm).

Table 3: Maximum Downloaded Users by Management Center Model

Management Center Model	Maximum Downloaded Users
FMC 1000	50,000
FMC 1600	50,000
FMC 1700	50,000
FMC 2500	150,000
FMC 2600	150,000
FMC 2700	150,000
FMC 4500	600,000
FMC 4600	600,000
FMC 4700	600,000
Management Center Virtual (any supported hypervisor)	50,000
Management Center Virtual 300 (any supported hypervisor)	150,000

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.

User Limits for Microsoft Azure Active Directory Realms

Microsoft Azure Active Directory User Limits

About user limits

Your management center model determines how many individual users you can monitor.

Note the following:

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



Note The system downloads all user sessions to all threat defense devices. If you have devices with different user concurrent user session limits, the threat defense with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your management center manages a Firepower 2110 and a 4125, the 2110 reports health warnings when the number of concurrent user sessions approaches its maximum of 64,000.)

Refer to the following tables.

Table 4: Maximum Downloaded Users by Management Center Model¹

Management Center Model	Number of Cisco Secure Dynamic Attributes Connector Connectors	Maximum Downloaded Users
FMC1600, FMC1700	10	50,000
FMC2600, FMC2700	20	150,000
FMC4600, FMC4700	30	600,000
Management Center Virtual (any supported hypervisor)	10	50,000
Management Center Virtual 300 (any supported hypervisor)	20	150,000

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

Table 5: Maximum Concurrent User Login Limits by Threat Defense

Threat Defense Model	Maximum Concurrent User Logins
Threat Defense Virtual 5, 10, 20, 30, 50 (any supported hypervisor)	50,000
Firepower 1010, 1120, 1140, 1150 Secure Firewall 3110, 3120, 3130, 3140	50,000
Secure Firewall 3130, 3140	150,000
Firepower 4140, 4145, 4150 Firepower 9300 Secure Firewall 4215	225,000
Secure Firewall 4225, 4245	300,000