

Identity Overview

The following topics discuss user identity:

- User identity, on page 1
- Recommendation: identity realm limit, on page 14
- Cloud-Delivered Firewall Management Center host and user limits, on page 15
- User limits for microsoft azure active directory realms, on page 17
- Maximum mappings per device, on page 17

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.

User identity information can help you 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.

Identity terminology

Identity terminology is a set of concepts that

- identifies users on networks using identity sources
- enables user awareness from authoritative and non-authoritative sources, and
- manages user access using identity policies.

Key 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, TS Agent, or Cisco ISE). 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. Access control rules specify user and group conditions and traffic actions. 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 source**: a trusted server that validates user logins (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 repository. ISE/ISE-PIC, the TS Agent, Microsoft Active Directory, and Microsoft Azure Active Directory, and the Passive Identity Agent are passive authentication user repositories supported by the system.

    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. Active authentication generally uses the same user repositories as passive authentication (the exceptions being ISE/ISE-PIC, and TS Agent, and the Passive Identity Agent, which are passive only).
- **Non-authoritative identity source**: an unknown or untrusted server that validates 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

Learn how to select and configure user identity sources for the firewall to enable user awareness and access control policies. This overview covers supported identity sources, authentication types, and key considerations for optimizing user identity management.

The 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...
Captive portal	OpenLDAP Microsoft Active Directory Microsoft Azure Active Directory	Authoritative	Active	Yes	Captive portal identity source Create a Microsoft Azure AD (SAML) realm for active authentication (captive portal)
Passive authentication	OpenLDAP Microsoft Active Directory	Non-authoritative	Active	Yes	Create an LDAP realm or an Active Directory realm and realm directory
Passive authentication	Microsoft Azure Active Directory	Passive	Passive	Yes	Configure Microsoft Entra ID for passive authentication
Passive authentication with the TS Agent	Microsoft Windows Terminal Server	Authoritative	Passive	Yes	Terminal services agent identity source
Passive authentication with pxGrid Cloud identity source	Cisco ISE	Non-authoritative	Passive	Yes	pxGrid Cloud identity source
Remote Access VPN	OpenLDAP or Microsoft Active Directory	Authoritative	Active	Yes	Remote access VPN identity sources
	RADIUS	Authoritative	Active	No, awareness only	
ISE/ISE-PIC	Microsoft Active Directory	Authoritative	Passive	Yes	The ISE/ISE-PIC identity source
Traffic-based detection (Configured in the network discovery policy.)	—	Non-authoritative	—	No, awareness only	Traffic-based detection identity sources

Consider the information 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 Security Cloud Control's users database and the user activity database. You can configure Security Cloud Control-server user downloads to automatically and regularly download new user data to your databases.

After you configure identity rules using the chosen identity source, you must associate the identity policy 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 Associate prefilter, decryption, and identity policies with an access control.

For general information about user identity, see User identity, on page 1.

Best practices for user identity

General recommendations for identity policy setup

Review the following recommendations before you set up identity policies.

- Know user limits
- Health monitor
- Use latest version of ISE/ISE-PIC, two types of remediation
- Captive portal requires routed interface, several individual tasks

Microsoft active directory and LDAP

Create one realm per LDAP server or Active Directory domain. The system supports Active Directory, LDAP, and other user repositories for user awareness and control. The association between an Active Directory or LDAP repository and the Security Cloud Control is referred to as a *realm*. 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.

Passive Identity Agent identity source

The passive identity agent software can monitor multiple Microsoft Active Directory servers and domain controllers and send user name and IP address information to the Security Cloud Control. For more information, see The passive identity agent identity source.

pxGrid Cloud Identity Source

The Cisco Identity Services Engine (Cisco ISE) pxGrid Cloud identity source enables you to use subscription and user data from Cisco ISE in Cloud-Delivered Firewall Management Center access control rules. Also,

the identity source uses constantly changing dynamic objects from Cisco ISE in access control policies in the Cloud-Delivered Firewall Management Center.

For more information, see pxGrid Cloud identity source.

Proxy sequence

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 Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

Although you can use one managed device as a proxy sequence, we strongly recommend you set up two or more so that, in the event one managed device cannot communicate with Active Directory or ISE/ISE-PIC, another managed device can take over.

Health monitor

The Security Cloud Control health monitor provides valuable information about the status of various Security Cloud Control functions, including user/realm mismatches, Snort memory usage, and 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

- User limits are discussed in User limits for microsoft active directory.
- If you use the ISE/ISE-PIC identity source, you can optionally limit the subnets the Cloud-Delivered Firewall 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 Best practices for ISE/ISE-PIC integration.
- To set up the ISE/ISE-PIC identity source, see ISE/ISE-PIC configurations for user control.

Captive portal information

You can use captive portal active authentication with any of the following: LDAP, Microsoft AD, or Microsoft Azure AD (SAML).

- For LDAP, see Captive portal identity source.
- For Microsoft AD, see Captive portal identity source.
- For Microsoft Azure AD (SAML), see Create a Microsoft Azure AD (SAML) realm.

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 Security Cloud Control.

- 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 Best practices for TS agent configuration.

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 Associate prefilter, decryption, and identity policies with an access control.

Identity deployments

The Secure Firewall Management Center uses various identity sources, such as TS Agent, ISE/ISE-PIC, pxGrid Cloud, passive identity agent, and captive portal, to collect and manage user and group data for access control and user awareness.

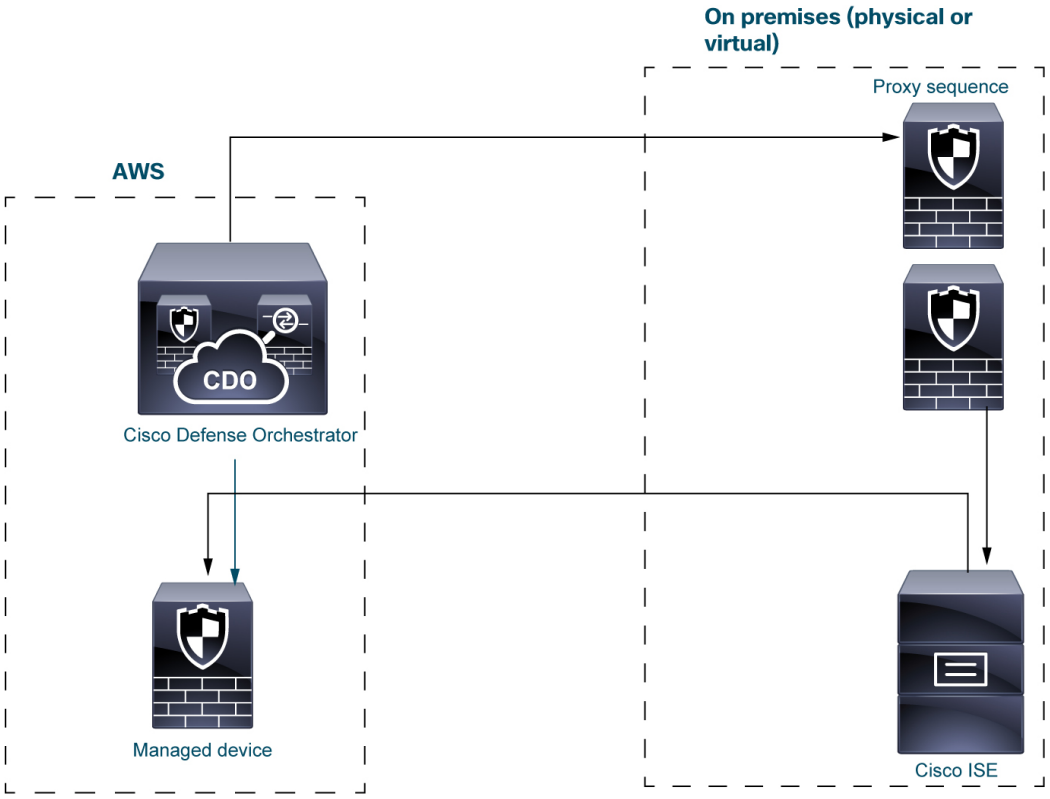
User data collection and storage

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 Security Cloud Control 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 Security Cloud Control.

Here are examples of identity deployments.

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



In the preceding figure, Security Cloud Control 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 Security Cloud Control 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 Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.)

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

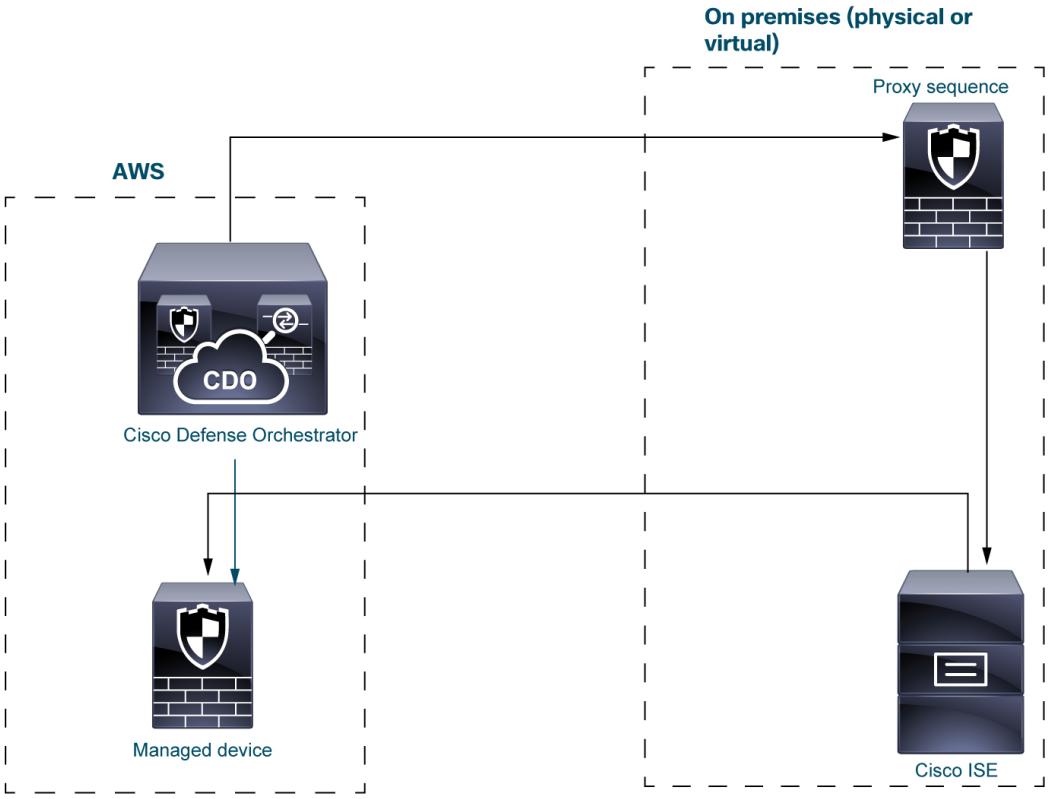
For more information about setting up a system like this, see Set up an identity policy, on page 9.

ISE/ISE-PIC identity source

When you deploy the ISE/ISE-PIC identity source, Security Cloud Control contacts the proxy sequence if Security Cloud Control 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 this figure.

For more information about ISE/ISE-PIC, see The ISE/ISE-PIC identity source.



Passive identity agent source

Software that monitors one or several Microsoft AD domain controllers, sending user name and IP address information to the Security Cloud Control.

For more information, see The passive identity agent identity source.

pxGrid cloud identity source

pxGrid Cloud Identity Source

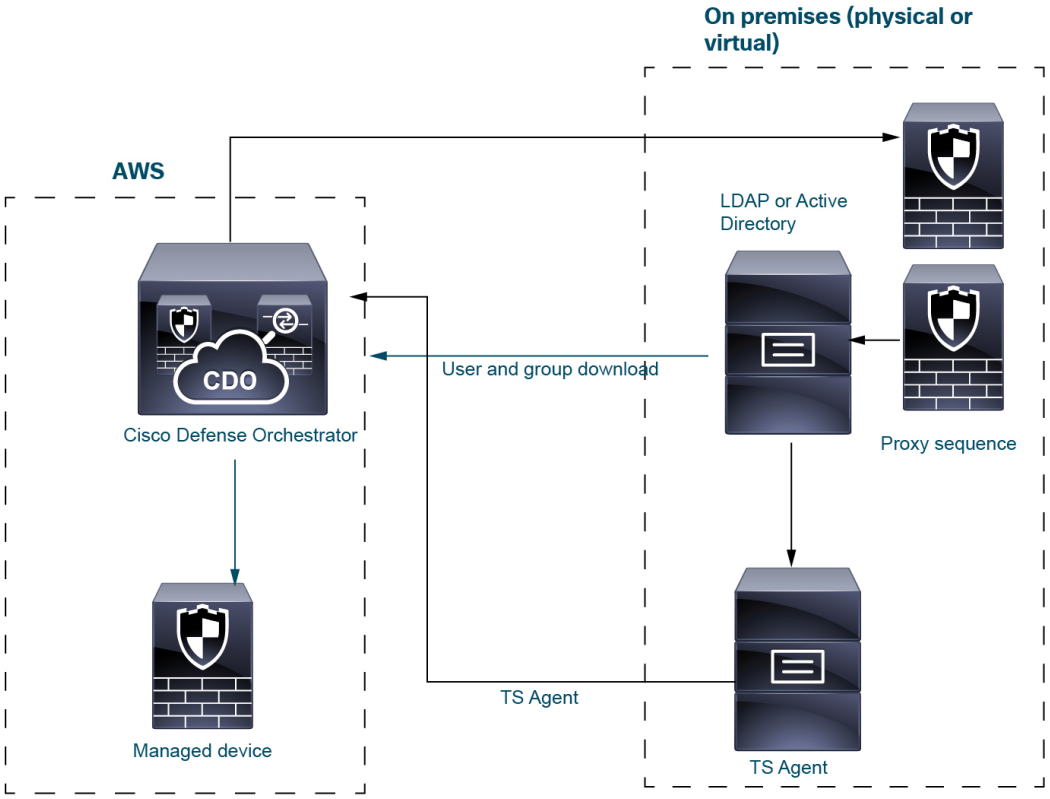
The Cisco Identity Services Engine (Cisco ISE) pxGrid Cloud identity source enables you to use subscription and user data from Cisco ISE in Cloud-Delivered Firewall Management Center access control rules. Also, the identity source uses constantly changing dynamic objects from Cisco ISE in access control policies in the Cloud-Delivered Firewall Management Center.

For more information, see pxGrid Cloud identity source.

TS Agent identity source

The Terminal Services (TS) Agent software runs on a Microsoft Server and sends Security Cloud Control 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 Security Cloud Control.

For more information about the TS Agent identity source, see Terminal services agent identity source.



Captive portal

You can use captive portal active authentication with any of these

- LDAP

  For more information, see Captive portal identity source.

- Microsoft AD

  For more information, see Captive portal identity source.

- Microsoft Azure AD (SAML)

  For more information, see Create a Microsoft Azure AD (SAML) realm.

Set up an identity policy

Set up an identity policy to control user access and filter network traffic using various user identity sources such as TS Agent, ISE/ISE-PIC, Microsoft Azure AD (SAML), captive portal.

An identity policy is not required if all these conditions are true:

- You use the ISE or 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.

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.

Procedure

	Command or Action	Purpose
Step 1	Choose your identity source.	• To use Microsoft Active Directory or LDAP, continue with the next step. • To use a Microsoft Azure AD (SAML) realm, see Create a Microsoft Azure AD (SAML) realm. • To use the Cisco ISE identity source, see Ways to configure the Cisco Identity Services Engine (Cisco ISE) identity source. • To use the passive identity agent identity source, see The passive identity agent identity source. • To use the pxGrid Cloud identity source, see pxGrid Cloud identity source.
Step 2	(Optional.) Create a proxy sequence.	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 Security Cloud Control cannot communicate with your Active Directory or ISE/ISE-PIC server. (For example, Security Cloud Control might be in a public cloud but Active Directory or ISE/ISE-PIC might be in a private cloud.) Although you can use one managed device as a proxy sequence, we strongly recommend you set up two or more so that, in the event one managed device cannot communicate with Active Directory or ISE/ISE-PIC, another managed device can take over. See Create a proxy sequence.
Step 3	(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 Cloud-Delivered Firewall Management Center realms and directories can be used in identity policies.	Creating a realm, realm directory, and proxy sequence is optional if any of these 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.

	Command or Action	Purpose
		• A proxy sequence is required only if you use Security Cloud Control (Security Cloud Control) and it cannot directly communicate with Active Directory or ISE/ISE-PIC. The *realm* is a trusted user and group store, typically a Microsoft Active Directory repository. The Cloud-Delivered Firewall Management Center downloads users and groups at intervals you specify. You can include or exclude users and groups from being downloaded. 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. A *directory* 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 not user, group, realm, Endp
Step 4	Synchronize users and groups from the realm.	To be able to control users and groups, you must synchronize them with the Cloud-Delivered Firewall 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.

	Command or Action	Purpose
		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 5	(Optional.) Create a realm sequence.	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.
Step 6	Create a method to retrieve user and group data (the *identity source*).	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 these: • Configure the captive portal for user control • Configure Remote Access VPN for AD-authenticated users
Step 7	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 se not user, group, realm, Endpoint only to filter network traffic.
Step 8	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 9	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 Associate prefilter, decryption, and identity policies with an access control.

What to do next

1. 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.

2. 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. .

User activity database

The user activity database on the Cloud-Delivered Firewall Management Center is a logging system that contains records of user activity on your network detected or reported by all of your configured identity sources.

User activity logging events

The system logs events in these 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 Cloud-Delivered Firewall 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 Cloud-Delivered Firewall Management Center.

You can view user activity detected by the system using the Cloud-Delivered Firewall Management Center. (**Events & Logs** > **Users** > **User Activity**.)

Users database

A users database on the Cloud-Delivered Firewall Management Center is a system component that

- contains a record for each user detected or reported by all configured identity sources
- stores data obtained from authoritative sources for user control, and
- prioritizes user data based on identity source type when user limits are reached.

User database management and limitations

For more information about the supported non-authoritative and authoritative identity sources, refer to About user identity sources, on page 2.

The total number of users the Cloud-Delivered Firewall Management Center can store depends on the Cloud-Delivered Firewall Management Center model. After the user limit is reached, the system prioritizes previously-undetected user data based on its identity source, as described in this list:

- Non-authoritative identity source: 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.
- Authoritative identity source: 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 Cloud-Delivered Firewall Management Center. These excluded user names remain in the database, but are not associated with IP addresses.

If you have Cloud-Delivered Firewall 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 Cloud-Delivered Firewall Management Center. The unidentified users are logged as Unknown users on the Cloud-Delivered Firewall 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 Cloud-Delivered Firewall 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 Cloud-Delivered Firewall Management Center.

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

- A user on the Cloud-Delivered Firewall 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.

Recommendation: identity realm limit

We limit you to a maximum of 25 realms regardless of Firewall Threat Defense model; the same limit applies to the Cloud-Delivered Firewall Management Center.

The number of *users* you can download to a system is limited according to the information provided in:

Cloud-Delivered Firewall Management Center user limit, on page 16

Cloud-Delivered Firewall Management Center host and user limits

Cloud-Delivered Firewall Management Center host and user limits are restrictions that define the maximum number of hosts and users that can be managed or supported by the system.

User and host limit details

These topics describe the user and host limits for the Cloud-Delivered Firewall Management Center.

Cloud-Delivered Firewall Management Center host limit

Cloud-Delivered Firewall Management Center host limit describes the storage capacity that

- allows the system to store a maximum of 600,000 hosts in its host database
- automatically adds a host to the network map when activity is detected for an IP address in your monitored network, and
- controls host visibility and contextual data access based on network map inclusion.

Host limit recommendations

The system can store up to 600,000 hosts but we recommend limits based on the number of devices managed.

Number of devices managed by Security Cloud Control	Recommended number of hosts
1-50	100,000
51-300	300,000
301-1000	600,000

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

When you reach the host limit, the network discovery policy controls what happens when you detect a new host. 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.

Cloud-Delivered Firewall Management Center user limit

A user is added to the Cloud-Delivered Firewall 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).

A Cloud-Delivered Firewall Management Center can store a maximum of 600,000 users in its host database but we recommend these limits.

Number of devices managed by Cisco Security Cloud	Recommended number of users
1-50	100,000
51-300	300,000
301-1000	600,000

Note The Secure Firewall 200 series can download a *cumulative total* of 10,000 user IPs, SXP/SGT mappings, endpoint profiles, and dynamic objects. After a total of 10,000 have been downloaded, the Secure Firewall 200 stops downloading any objects until some previously downloaded objects have been removed. (For example, if users log out, the memory is free for other objects.) This limit applies regardless of how many devices Cisco Security Cloud manages.

Because of its limited capacity, we recommend you *do not* use a Secure Firewall 220 as a proxy sequence.

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

The Cloud-Delivered Firewall Management Center can store 600,000 sessions in its user database.

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

This reference provides user limit information for Microsoft Azure Active Directory realms, showing how different device models handle user monitoring and login capacity.

Microsoft azure active directory user limits

About user limits

Your Cloud-Delivered Firewall Management Center model determines how many individual users you can monitor.

Note the following:

- The maximum number of *downloaded* users depends on your Cloud-Delivered Firewall 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 Firewall Threat Defense devices. If you have devices with different user concurrent user session limits, the Firewall Threat Defense with the smallest limit reports health warnings when its memory reaches the configured limit. (For example, if your Cloud-Delivered Firewall 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 these tables.

Table 1: Maximum Downloaded Users by Cloud-Delivered Firewall Management Center Model[1]

Cloud-Delivered Firewall Management Center Model	Number of Dynamic Attributes Connector Connectors	Maximum Downloaded Users
Firewall Management Center Virtual (any supported hypervisor)	10	50,000
Firewall Management Center Virtual 300 (any supported hypervisor)	20	150,000

[1]—Cloud-Delivered Firewall 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.

Maximum mappings per device

The identity module is optimized to accommodate a higher number of mappings. The mappings include users, shared users, Security Group Exchange Protocol (SXP), and Attribute-Based Policy (ABP) IP bindings. The

maximum number of mappings depends on the managed device model. A single user can have multiple sessions from different IP addresses. An IP address can have multiple mappings.

This table shows the maximum number of mappings and the experimental enhanced values for each Firewall Threat Defense device.

Table 2: Maximum Mapping Limits by Firewall Threat Defense

Firewall Threat Defense Model	Maximum Mappings	Enhanced Maximum Mappings (Beta)*
Firewall Threat Defense Virtual [1]: • 4 vCPU / 8 GB RAM • 8 vCPU / 16 GB RAM • 12 vCPU / 24 GB RAM • 16 vCPU / 32 GB RAM	50,000	128,000
Firewall Threat Defense Virtual [1]: • 32 vCPU / 64 GB RAM • 64 vCPU / 128 GB RAM	50,000	300000
Secure Firewall 220	10,000	10,000
Firepower 1010, 1120, 1140, 1150 [1]	50,000	128,000
Firepower 1210, 1220, 1230, 1240, 1250 [1]	50,000	128,000
Firepower 2110, 2120, 2130 [1]	50,000	NA
Firepower 2140	150,000	NA
Secure Firewall 3105, 3110, 3120, 3130, 3140 [1]	50,000	300,000
Firepower 4112, 4115, 4125, 4145, 4215	150,000	300,000
Firepower 9300 [2]	225,000	300,000
Secure Firewall 4225, 4245	300,000	600,000
Secure Firewall 6160, 6170	300,000	600,000

The Secure Firewall 200 series devices can download a *cumulative total* of 10,000 user IPs, SXP/SGT mappings, endpoint profiles, and dynamic objects. After reaching the total of 10,000, the device stops downloading new objects until previously downloaded objects are removed. For example, when users log out, the memory frees space for other objects.

[1] These devices support a maximum of 64,000 mappings when you restrict to using only Active Directory user mappings.

[2] This device supports a maximum of 300,000 mappings when you restrict to using only Active Directory user mappings.

Maximum unique user groups and dynamic objects in a policy (Beta)*

An access control policy supports the following number of unique user groups and dynamic objects, across all device models.

- Maximum number of AD and Azure AD groups: 4,000
- Maximum number of dynamic objects: 8,000

These values are experimental. Typically, an IP address can be associated with 32 unique dynamic objects. A user can normally be associated with 32 groups.



Note *This is a Beta feature. Enabling enhanced maximum mappings may increase memory utilization on the firewall in larger scale deployments. Carefully evaluate and monitor the use of this feature in your environment.