



Identity Sources

Identity sources are servers and databases that define user accounts. You can use this information in a variety of ways, such as providing the user identity associated with an IP address, or authenticating remote access VPN connections or access to the device manager.

The following topics explain how to define the identity sources. You would then use these objects when you configure the services that require an identity source.

- [About Identity Sources, on page 1](#)
- [Active Directory \(AD\) Identity Realms, on page 3](#)
- [RADIUS Servers and Groups, on page 9](#)
- [Identity Services Engine \(ISE\), on page 12](#)
- [SAML Servers, on page 16](#)
- [Local Users, on page 18](#)

About Identity Sources

Identity sources are the AAA servers and databases that define user accounts for the people in your organization. You can use this information in a variety of ways, such as providing the user identity associated with an IP address, or authenticating remote access VPN connections or access to the device manager.

Use the **Objects > Identity Sources** page to create and manage your sources. You would then use these objects when you configure the services that require an identity source

Following are the supported identity sources and their uses:

Active Directory (AD) Identity Realm

Active Directory provides user account and authentication information. See [Active Directory \(AD\) Identity Realms, on page 3](#).

You can use this source for the following purposes:

- Remote Access VPN, as a primary identity source. You can use AD in conjunction with a RADIUS server.
- Identity policy, for active authentication and as the user identity source used with passive authentication.

AD (Active Directory) Realm Sequence

An AD realm sequence is an ordered list of AD realm objects. Realm sequences are useful if you manage more than one AD domain in your network. See [Configuring an AD Realm Sequence, on page 7](#).

You can use this source for the following purposes:

- Identity policy, as the user identity source used with passive authentication. The order of realms in the sequence determines how the system determines user identity in the rare cases there is a conflict.

Cisco Identity Services Engine (ISE) or Cisco Identity Services Engine Passive Identity Connector (ISE PIC)

If you are using ISE, you can integrate the threat defense device with your ISE deployment. See [Identity Services Engine \(ISE\), on page 12](#).

You can use this source for the following purposes:

- Identity policy, as a passive identity source to collect user identity from ISE.

RADIUS Server, RADIUS Server Group

If you are using RADIUS servers, you can also use them with the device manager. You must define each server as a separate object, then put them in server groups (where the servers in a given group are copies of each other). You assign the server group to features, you do not assign individual servers. See [RADIUS Servers and Groups, on page 9](#).

You can use this source for the following purposes:

- Remote Access VPN, as an identity source for authentication, and for authorization and accounting. You can use AD in conjunction with a RADIUS server.
- Identity policy, as a passive identity source to collect user identity from remote access VPN logins.
- External authentication for the device manager or the threat defense CLI management users. You can support multiple management users with different authorization levels. These users can log into the system for device configuration and monitoring purposes.

SAML Server

Security Assertion Markup Language 2.0 (SAML 2.0) is an open standard for exchanging authentication and authorization data between parties, specifically an Identity Provider (IdP) and Service Provider (SP).

You can use this source for the following purposes:

- Remote access VPN, as a single sign-on (SSO) authentication source.
- External authentication for the device manager users. You can support multiple management users with different authorization levels. These users can log into the system for device configuration and monitoring purposes.

LocalIdentitySource

This is the local user database, which includes users that you have defined in the device manager. Select **Objects > Users** to manage the user accounts in this database. See [Local Users, on page 18](#).



Note The local identity source database does not include users you configure in the CLI for CLI access (using the **configure user add** command). CLI users are completely separate from those you create in the device manager.

You can use this source for the following purposes:

- Remote Access VPN, as a primary or fallback identity source.
- Identity policy, as a passive identity source to collect user identity from remote access VPN logins.

Active Directory (AD) Identity Realms

Microsoft Active Directory (AD) defines user accounts. You can create an AD identity realm for an Active Directory domain. The following topics explain how to define an AD identity realm.

Supported Directory Servers

You can use Microsoft Active Directory (AD) on Windows Server 2012, 2016, and 2019.

Note the following about your server configuration:

- If you want to perform user control on user groups or on users within groups, you must configure user groups on the directory server. The system cannot perform user group control if the server organizes the users in basic object hierarchy.
- The directory server must use the field names listed in the following table in order for the system to retrieve user metadata from the servers for that field.

Metadata	Active Directory Field
LDAP user name	samaccountname
first name	givenname
last name	sn
email address	mail userprincipalname (if mail has no value)
department	department distinguishedname (if department has no value)
telephone number	telephonenumber

Limitations on Number of Users

Device Manager can download information on up to 50,000 users from the directory server.

If your directory server includes more than 50,000 user accounts, you will not see all possible names when selecting users in an access rule or when viewing user-based dashboard information. You can write rules on only those names that were downloaded.

The limit also applies to the names associated with groups. If a group has more than 50,000 members, only the 50,000 names that were downloaded can be matched against the group membership.

Determining the Directory Base DN

When you configure directory properties, you need to specify the common base distinguished name (DN) for users and groups. The base is defined in your directory server, and differs from network to network. You must enter the correct bases for identity policies to work. If the base is wrong, the system cannot determine user or group names, and thus identity-based policies will be inoperable.



Tip To get the correct bases, consult the administrator who is responsible for the directory servers.

For active directory, you can determine the correct bases by logging into the Active Directory server as domain administrator, and using the **dsquery** command at a command prompt as follows to determine the bases:

User search base

Enter the **dsquery user** command with a known username (partial or complete) to determine the base distinguished name. For example, the following command uses the partial name “John*” to return information for all users that start with “John.”

```
C:\Users\Administrator>dsquery user -name "John*"
"CN=John Doe,CN=Users,DC=csc-lab,DC=example,DC=com"
```

The base DN would be “DC=csc-lab,DC=example,DC=com.”

Group search base

Enter the **dsquery group** command with a known group name to determine the base distinguished name. For example, the following command uses the group name Employees to return the distinguished name:

```
C:\>dsquery group -name "Employees"
"CN=Employees,CN=Users,DC=csc-lab,DC=example,DC=com"
```

The group base DN would be “DC=csc-lab,DC=example,DC=com.”

You can also use the ADSI Edit program to browse the Active Directory structure (**Start > Run > adsiedit.msc**). In ADSI Edit, right click any object, such as an organizational unit (OU), group, or user, and choose **Properties** to view the distinguished name. You can then copy the string of DC values as the base.

To verify that you have the correct base:

1. Click the Test Connection button in the directory properties to verify connectivity. Resolve any problems, and save the directory properties.
2. Commit changes to the device.
3. Create an access rule, select the **Users** tab, and try to add known user and group names from the directory. You should see auto-complete suggestions as you type for matching users and groups in the realm that

contains the directory. If these suggestions appear in a drop-down list, then the system was able to query the directory successfully. If you see no suggestions, and you are certain the string you typed should appear in a user or group name, you need to correct the corresponding search base.

Configuring AD Identity Realms

An identity realm is a directory server plus other attributes required to provide authentication services. The directory server contains information about the users and user groups who are allowed access to your network.

For Active Directory, a realm is equivalent to an Active Directory domain. Create separate realms for each AD domain you need to support.

Realms are used in the following policies:

- **Identity**—The realm provides user identity and group membership information, which you can then use in access control rules. The system downloads updated information about all users and groups every day in the last hour of the day (UTC). The directory server must be reachable from the management interface.
- **Remote access VPN**—The realm provides authentication services, which determine whether a connection is allowed. The directory server must be reachable from the RA VPN outside interface.
- **Access Control, SSL Decryption**—You can select the realm in the user criteria for the rule to apply the rule to all users within the realm.

Work with your directory administrator to get the values required to configure the directory server properties.



Note If the directory server is not on an attached network or available through the default route, create a static route for the server. Select **Device > Routing > View Configuration** to create static routes. Alternatively, select the appropriate interface when defining the server.

The following procedure explains how you can create and edit objects directly through the Objects page. You can also create identity realm objects while editing a realm property by clicking the **Create New Identity Realm** link shown in the object list.

Before you begin

Ensure that time settings are consistent among the directory servers, the threat defense device, and clients. A time shift among these devices can prevent successful user authentication. "Consistent" means that you can use different time zones, but the time should be the same relative to those zones; for example, 10 AM PST = 1 PM EST.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an AD realm, click + > **AD**.
- To edit a realm, click the edit icon (🔍) for the realm.

To delete an unreferenced object, click the trash can icon (🗑️) for the object.

Step 3 Configure the basic realm properties.

- **Name**—A name for the directory realm.
- **Type**—The type of directory server. Active Directory is the only supported type, and you cannot change this field.
- **Directory Username, Directory Password**—The distinguished username and password for a user with appropriate rights to the user information you want to retrieve. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

Note The system generates ldap-login-dn and ldap-login-password from this information. For example, Administrator@example.com is translated as cn=admin, cn=users, dc=example, dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name “users” folder.

- **Base DN**—The directory tree for searching or querying user and group information, that is, the common parent for users and groups. For example, cn=users,dc=example,dc=com. For information on finding the base DN, see [Determining the Directory Base DN, on page 4](#).
- **AD Primary Domain**—The fully qualified Active Directory domain name that the device should join. For example, example.com.

Step 4 Configure the directory server properties.

- **Hostname/IP Address**—The hostname or IP address of the directory server. If you use an encrypted connection to the server, you must enter the fully-qualified domain name, not the IP address.
- **Interface**—The interface through which the AD server should be reached. If you do not select an interface, the data routing table is used to find the appropriate interface. If you want to use a management-only interface, you must choose it specifically; you cannot use a route lookup from the management-only routing table.
- **Port**—The port number used for communications with the server. The default is 389. Use port 636 if you select LDAPS as the encryption method.
- **Encryption**—To use an encrypted connection for downloading user and group information, select the desired method, **STARTTLS** or **LDAPS**. The default is **None**, which means that user and group information is downloaded in clear text.
 - **STARTTLS** negotiates the encryption method, and uses the strongest method supported by the directory server. Use port 389. This option is not supported if you use the realm for remote access VPN.
 - **LDAPS** requires LDAP over SSL. Use port 636.
- **Trusted CA Certificate**—If you select an encryption method, upload a Certificate Authority (CA) certificate to enable a trusted connection between the system and the directory server. If you are using a certificate to authenticate, the name of the server in the certificate must match the server Hostname / IP Address. For example, if you use 10.10.10.250 as the IP address but ad.example.com in the certificate, the connection fails.



- Step 5** If there are multiple servers for the realm, click **Add Another Configuration** and enter the properties for each additional server.
- You can add up to 10 AD servers to the realm. These servers need to be duplicates of each other and support the same AD domain.
- You can collapse and expand each server entry for your convenience. The sections are labeled with the hostname/IP address and port.
- Step 6** Click the **Test** button to verify the system can contact the server.
- The system uses separate processes and interfaces to access the server, so you might get errors indicating that the connection works for one type of use but not another, for example, available for Identity policies but not for remote access VPN. If the server cannot be reached, verify that you have the right IP address and host name, that the DNS server has an entry for the hostname, and so forth. You might need to configure a static route for the server. For more information, see [Troubleshooting Directory Server Connections, on page 8](#).
- Step 7** Click **OK**.
-

Configuring an AD Realm Sequence

You can use an AD realm sequence in a passive identity rule so that the system can try to match users from more than one AD server. In a realm sequence, you configure an ordered list of AD realms where each AD server manages a different realm or domain, for example, `engineering.example.com` and `marketing.example.com`.

Realm sequences are useful only if you support more than one AD domain, and users from different domains might send traffic through the threat defense device. The realms are used in order to find the identity for a passively-authenticated user session. The order of the realms is used to resolve identity conflicts, in the rare cases where a conflict might arise.

Procedure

- Step 1** Select **Objects**, then select **Identity Sources** from the table of contents.
- Step 2** Do one of the following:
- To create an AD realm sequence, click + > **AD Realm Sequence**.
 - To edit an AD realm sequence, click the edit icon () for the object.
- To delete an unreferenced object, click the trash can icon () for the object.
- Step 3** Configure the realm sequence properties:
- **Name**—The name of the object.
 - **Description**—An optional description of the object.
 - **AD Realms**—Click + to add AD realm objects to the sequence. After you add the realms, click and drag and drop the realms into the desired ordered sequence.
- Step 4** Click **OK**.

You can now select the AD realm sequence in a passive identity rule.

Troubleshooting Directory Server Connections

The system uses different processes to communicate with your directory server depending on the feature. Thus, a connection for identity policies might work, whereas one for remote access VPN fails.

These processes use different interfaces to communicate with the directory server. You must ensure connectivity from these interfaces.

- Management interface, for: identity policies.
- Data interface, for: remote access VPN (outside interface).

When you configure the identity realm, use the **Test** button to verify that the connection can work. Failure messages should indicate the feature that is having connection problems. The following are the general issues you might encounter, based on authentication attributes and routing/interface configuration.

Directory user authentication issues.

If the problem is that the system could not log into the directory server because of the username or password, ensure that the name and password are correct and valid on the directory server. For Active Directory, the user does not need elevated privileges. You can specify any user in the domain. The username must be fully qualified; for example, Administrator@example.com (not simply Administrator).

Also, the system generates ldap-login-dn and ldap-login-password from the username and password information. For example, Administrator@example.com is translated as cn=administrator,cn=users,dc=example,dc=com. Note that cn=users is always part of this translation, so you must configure the user you specify here under the common name “users” folder.

The directory server is accessible through a data interface.

If the directory server is on a network that is either directly connected to a data interface (such as a GigabitEthernet interface), or routeable from a directly-connected network, you must ensure that there is a route between the virtual management interface and the directory server.

- Using **data-interfaces** as the management gateway should make routing successful.
- If you have an explicit gateway on the management interface, that gateway router needs to have a route to the directory server.
- If there is a router between the directly-connected network and the network that hosts the directory server, configure a static route for the directory server (**Device > Routing**).
- Verify that the data interface has the correct IP address and subnet mask.

The directory server is on an external network.

If the directory server is on a network on the other side of the outside (uplink) interface, you might need to configure a site-to-site VPN connection. For the detailed procedure, see [How to Use a Directory Server on an Outside Network with Remote Access VPN](#).

RADIUS Servers and Groups

You can use RADIUS servers to authenticate and authorize remote access VPN connections, and the device manager and the threat defense CLI administration users. For example, if you also use Cisco Identity Services Engine (ISE) and its RADIUS server, you can use that server with the device manager.

When you configure a feature to use RADIUS servers, you select a RADIUS group instead of individual servers. A RADIUS group is a collection of RADIUS servers that are copies of each other. If a group has more than one server, they form a chain of backup servers to provide redundancy in case one server becomes unavailable. But even if you have only one server, you must create a one-member group to configure RADIUS support for a feature.

The following topics explain how to configure RADIUS servers and groups, so that they are available for use in the supported features.

Configure RADIUS Servers

RADIUS servers provide AAA (authentication, authorization, and accounting) services. If you use RADIUS servers to authenticate and authorize users, you can use those servers with the device manager.

After creating objects for each of your RADIUS servers, create RADIUS server groups to contain each group of duplicate servers.


Before you begin


If you want to configure a redirect ACL for RA VPN, you must use Smart CLI to create the extended ACL before creating or editing the server object. You cannot create the ACL while editing the object.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an object, click + > **RADIUS Server**.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Configure the following properties:

- **Name**—The name of the object. This does not have to match anything configured on the server.
- **Server Name or IP Address**—The fully-qualified host name (FQDN) or IP address of the server. For example, radius.example.com or 10.100.10.10.
- **Authentication Port**—The port on which RADIUS authentication and authorization are performed. The default is 1812.
- **Timeout**—The length of time, 1-300 seconds, that the system waits for a response from the server before sending the request to the next server. The default is 10 seconds. If you are using this server as a secondary

authentication source for remote access VPN, for example, to prompt for an authentication token, increase this timeout to 60 seconds at least. This provides time for the user to obtain and enter the token.

- **Server Secret Key**—(Optional.) The shared secret that is used to encrypt data between the threat defense device and the RADIUS server. The key is a case-sensitive, alphanumeric string of up to 64 characters, with no spaces. The key must start with an alphanumeric character or an underscore, and it can contain the special characters: \$ & - _ . + @. The string must match the one configured on the RADIUS server. If you do not configure a secret key, the connection is not encrypted.

Step 4 (Optional.) If you are using the server for remote access VPN Change of Authorization configuration, you can click the **RA VPN Only** link and configure the following options.

- **Redirect ACL**—Select the extended ACL to use for the RA VPN redirect ACL. Create extended ACLs using the Smart CLI **Extended Access List** object on the **Device > Advanced Configuration > Smart CLI > Objects** page.

The purpose of the redirect ACL is to send initial traffic to Cisco Identity Services Engine (ISE) so that ISE can assess the client posture. The ACL should send HTTPS traffic to ISE, but not traffic that is already destined for ISE, or traffic that is directed to a DNS server for name resolution. For an example, see [Configure Change of Authorization on the Threat Defense Device](#).

- **Interface Used to Connect to RADIUS Server**—Which interface to use when communicating with the server. If you select **Resolve via Route Lookup**, the system always uses the data routing table to determine the interface to use. If you select **Manually Choose Interface**, the system will always use the interface you select. If you want to use a management-only interface, you must choose it specifically; you cannot use a route lookup for the management-only routing table.

If you are configuring Change of Authorization, you must select a specific interface so that the system can correctly enable the CoA listener on the interface.

If you also use this server for the device manager administrative access, this interface is ignored. Administrative access attempts are always authenticated through the management IP address.

Step 5 (Optional, when editing the object only.) Click **Test** to check whether the system can connect to the server.

You are prompted for a username and password. The test confirms whether the server can be contacted, and if it can, that the username can be authenticated.

Step 6 Click **OK**.

Configure RADIUS Server Groups


A RADIUS server group contains one or more RADIUS server objects. The servers within a group must be copies of each other. These servers form a chain of backup servers, so that if the first server is unavailable, the system can try the next server in the list.


When you configure RADIUS support in a feature, you must select a server group. Thus, even if you have just one RADIUS server, you must create a server group to contain it.

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an object, click + > **RADIUS Server Group**.
- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3 Configure the following properties:

- **Name**—The name of the object. This does not have to match anything configured on the servers.
- **Dead Time**—Failed servers are reactivated only after all servers have failed. The dead time is how long to wait, from 0 - 1440 minutes, after the last server fails before reactivating all servers. Deadtime applies only if you configure fallback to the local database; authentication is attempted locally until the deadtime elapses. The default is 10 minutes.
- **Maximum Failed Attempts**—The number of failed AAA transactions (that is, requests that do not get a response) sent to a RADIUS server in the group before trying the next server. You can specify 1-5, and the default is 3. When the maximum number of failed attempts is exceeded, the system marks the server as Failed.

For a given feature, if you configured a fallback method using the local database, and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. The server group remains marked as unresponsive for the duration of the dead time, so that additional AAA requests within that period do not attempt to contact the server group, and the fallback method is used immediately.

- **Dynamic Authorization (for RA VPN only), Port**—If you enable RADIUS dynamic authorization or change of authorization (CoA) services for this RADIUS server group, the group will be registered for CoA notification and listen on the specified port for CoA policy updates from Cisco Identity Services Engine (ISE). The default listening port is 1700, or you can specify a different port in the range 1024 to 65535. Enable dynamic authorization only if you are using this server group in a remote access VPN in conjunction with ISE
- **Realm that Supports the RADIUS Server**—If the RADIUS server is configured to use an AD server for authenticating users, select the AD realm that specifies the AD server used in conjunction with this RADIUS server. If the realm does not already exist, click **Create New Identity Realm** at the bottom of the list and configure it now.
- **RADIUS Server list**—Select up to 16 RADIUS server objects that define the servers for the group. Add these objects in priority order. The first server in the list is used until it becomes unresponsive. After adding the objects, you can drag and drop to rearrange them. If the object you need does not yet exist, click **Create New RADIUS Server** and add it now.

You can also click the **Test** link to verify the system can connect to the server. You are prompted for a username and password. The test confirms whether the server can be contacted, and if it can, that the username can be authenticated.

Step 4 (Optional.) Click the **Test All Servers** button to check connectivity to each server in the group.

You are prompted for a username and password. The system checks whether each server can be contacted, and whether the username can be authenticated on each server.

Step 5 Click **OK**.

Troubleshoot RADIUS Servers and Groups

Following are some things you can check if external authorization does not work.

- Use the **Test** buttons in the RADIUS server and server group objects to verify that the servers can be contacted from the device. Ensure that you save the objects before testing. If the test fails:
 - Please understand that the test ignores the interface configured for the server, and always uses the management interface. The test is expected to fail if the RADIUS authentication proxy is not configured to respond to requests from the management IP address.
 - Verify that you are entering a correct username/password combination during the test. You should get a Bad Credentials message if they are incorrect.
 - Verify the secret key, port, and IP address for the server. If you are using a hostname, verify that DNS is configured for the management interface. Consider the possibility that the secret key was changed on the RADIUS server but not in the device configuration.
 - If the test continues to fail, you might need to configure a static route to the RADIUS servers. Try pinging the server from the CLI Console or an SSH session to see if it can be reached.
- If external authentication has been working, but has stopped working, consider the possibility that all servers are in the dead time. If you configure fallback to local authentication, when all the RADIUS servers within a group have failed, the dead time is the number of minutes the system waits before trying the first server again. During the dead time, local authentication is used, so a given user's username and password would be the local username/password. The default is 10 minutes, but you can configure as long as 1440 minutes.
- If HTTPS external authentication works for some users but not others, evaluate the cisco-av-pair attribute defined in the RADIUS server for each user account. This attribute might be configured incorrectly. A missing or incorrect attribute will block all HTTPS access for that user account.
- If SSH external authentication works for some users but not for others, evaluate the Service-Type attribute defined in the RADIUS server for each user account. This attribute might be configured incorrectly. A missing or incorrect attribute will block all SSH access for that user account.

Identity Services Engine (ISE)

You can integrate your Cisco Identity Services Engine (ISE) or ISE Passive Identity Connector (ISE-PIC) deployment with the threat defense device to use ISE/ISE-PIC for passive authentication.

ISE/ISE-PIC is an authoritative identity source, and provides user awareness data for users who authenticate using Active Directory (AD), LDAP, RADIUS, or RSA. However, for threat defense, you can use ISE for user identity awareness in conjunction with AD only. You can use the user identity in access control and SSL decryption policies as matching criteria, in addition to seeing user information in the various monitoring dashboards and events.

For more information on Cisco ISE/ISE-PIC, see the *Cisco Identity Services Engine Administrator Guide* (<https://www.cisco.com/c/en/us/support/security/identity-services-engine/tsd-products-support-series-home.html>) and the *Identity Services Engine Passive Identity Connector (ISE-PIC) Installation and Administrator Guide* (<https://www.cisco.com/c/en/us/support/security/ise-passive-identity-connector/tsd-products-support-series-home.html>).

Guidelines and Limitations for ISE

- The firewall system does not support 802.1x device authentication alongside Active Directory authentication because the system does not associate device authentication with users. If you use 802.1x active logins, configure ISE to report only 802.1x active logins (both device and user). That way, a device login is reported only once to the system.
- ISE/ISE-PIC does not report the activity of ISE Guest Services users.
- Synchronize the time on the ISE/ISE-PIC server and the device. Otherwise, the system might perform user timeouts at unexpected intervals.
- If you configure ISE/ISE-PIC to monitor a large number of user groups, the system might drop user mappings based on groups due to memory limitations. As a result, rules with realm or user conditions might not perform as expected.
- For the specific versions of ISE/ISE-PIC that are compatible with this version of the system, see the *Cisco Secure Firewall Compatibility Guide*, <https://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-device-support-tables-list.html>.
- Use the IPv4 address of the ISE server, unless you confirm that your version of ISE supports IPv6.

Configure Identity Services Engine

To use the Cisco Identity Services Engine (ISE) or Cisco Identity Services Engine Passive Identity Connector (ISE PIC) as a passive identity source, you must configure the connection to the ISE Platform Exchange Grid (pxGrid) server.

Before you begin

- Export the pxGrid and MNT server certificates from ISE. For example, in ISE PIC 2.2, you find these on the **Certificates > Certificate Management > System Certificates** page. The MNT (Monitoring and Troubleshooting node) is shown as Admin in the Used By column in the certificates list. You can either upload them as trusted CA certificates on the **Objects > Certificates** page, or upload them during the following procedure. These nodes might be using the same certificate.
- You must also configure an AD identity realm. The system obtains the list of users from AD, and from ISE it gets information on the user-to-IP address mappings.
- If you will use security group tags (SGT) for access control, with or without static security group tag mappings, and listen to the SXP topic, you also need to configure SXP and these mappings in ISE. See [Configure Security Groups and SXP Publishing in ISE](#).

Procedure

Step 1 Select **Objects**, then select **Identity Sources** from the table of contents.

Step 2 Do one of the following:

- To create an object, click + > **Identity Services Engine**. You can create at most one ISE object.
- To edit an object, click the edit icon (🔗) for the object.

To delete an unreferenced object, click the trash can icon (🗑️) for the object.

Step 3 Configure the following properties:

- **Name**—The name of the object.
- **Status**—Click the toggle to enable or disable the object. When disabled, you cannot use ISE as an identity source in your identity rules.
- **Description**—An optional description of the object.
- **Primary Node Hostname/IP Address**—The hostname or IP address for the primary pxGrid ISE server. Do not specify an IPv6 address unless you verify that your version of ISE supports IPv6.
- **Secondary Node Hostname/IP Address**—If you set up a secondary ISE server for high availability, click **Add Secondary Node Hostname/IP Address** and enter the hostname or IP address of the secondary pxGrid ISE server.
- **pxGrid Server CA Certificate**—The trusted Certificate Authority certificate for the pxGrid framework. If your deployment includes a primary and a secondary pxGrid node, the certificates for both nodes must be signed by the same certificate authority.
- **MNT Server CA Certificate**—The trusted Certificate Authority certificate for the ISE certificate when performing bulk downloads. This can be the same as the pxGrid server certificate if your MNT (Monitoring and Troubleshooting) server is not separate. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.
- **Server Certificate**—The internal identity certificate that the threat defense device must provide to ISE when connecting to ISE or performing bulk downloads.
- **Subscribe To**—Select which ISE pxGrid topics should be subscribed to. Subscribing to a topic means you will download data related to the topic.
 - **Session Directory Topic**—Whether to obtain information about user sessions, including SGT mappings for user sessions. This option is enabled by default. You should select this option if you want to obtain passive user identity for use in security policies and for visibility in the monitoring dashboards.
 - **SXP Topic**—Whether to obtain static SGT-to-IP address mappings. Select this topic if you want to write access control rules based on security group tags (SGT).
- **ISE Network Filters**—An optional filter you can set to restrict the data that ISE reports to the system. If you provide a network filter, ISE reports data from the networks within the filter only. Click +, select the network objects that identify the networks, and click **OK**. Click **Create New Network** if you need to create the objects. Configure IPv4 network objects only.

Step 4 Click the **Test** button to verify that the system can connect to your ISE server.

If the test fails, click the **See Logs** link to read the detailed error messages. For example, the following message indicates that the system could not connect to the server at the required port. The problem might be no route to the host, that the ISE server is not using the expected port, or that you have access control rules that prevent the connection.

```
Captured Jabberwerx log:2018-05-11T16:10:30 [ ERROR]: connection timed out while
trying to test connection to host=10.88.127.142:ip=10.88.127.142:port=5222
```

Step 5 Click **OK** to save the object.

What to do next

After you configure ISE, enable the identity policy, configure passive authentication rules, and deploy the configuration. Then, you must go into ISE/ISE PIC and accept the device as a subscriber. If you configure ISE/ISE PIC to automatically accept subscribers, you do not need to manually accept the subscription.

Troubleshoot the ISE/ISE-PIC Identity Source

ISE/ISE-PIC Connections

If you experience issues with the ISE or ISE-PIC connection, check the following:

- The pxGrid Identity Mapping feature in ISE must be enabled before you can successfully integrate ISE with the threat defense device.
- Before a connection between the ISE server and the threat defense device succeeds, you must manually approve the clients in ISE.

Alternatively, you can enable **Automatically approve new accounts** in ISE as discussed in the chapter on Managing users and external identity sources in the *Cisco Identity Services Engine Administrator Guide*.

- The threat defense device (server) certificate must include the **clientAuth** extended key usage value, or it must not include any extended key usage values. If the clientAuth extended key usage is set, then there must also either be no key usage set, or the Digital Signature key usage value must be set. The self-signed identity certificates you can create using the device manager meet these requirements.
- The time on your ISE server must be synchronized with the time on the threat defense device. If the appliances are not synchronized, the system might perform user timeouts at unexpected intervals.

ISE/ISE-PIC User Data

If you experience issues with user data reported by ISE or ISE-PIC, note the following:

- After the system detects activity from an ISE user whose data is not yet in the database, the system retrieves information about them from the server. Activity seen by the ISE user is not handled by access control rules, and is not displayed in the dashboards until the system successfully retrieves information about them in a user download.

- You cannot perform user control on ISE users who were authenticated by an LDAP, RADIUS, or RSA domain controller.
- The system does not receive user data for ISE Guest Services users.

SAML Servers

You can configure Security Assertion Markup Language 2.0 (SAML 2.0) servers to use as single sign-on (SSO) authentication sources for remote access VPN connections and device manager users. SAML is an open standard for exchanging authentication and authorization data between parties, specifically an Identity Provider (IdP) and Service Provider (SP).

Configure SAML Servers

You can configure Security Assertion Markup Language 2.0 (SAML 2.0) servers to use as single sign-on (SSO) authentication sources for remote access VPN connections and device manager users. For example, the Duo Access Gateway (DAG) is a SAML server.

When you use a SAML server as the authentication method, the SAML server acts as the Identity Provider (IdP), whereas the threat defense device acts as the Service Provider (SP).

For RA VPN, you can use a SAML server as the primary authentication source, but you cannot configure a secondary authentication source, nor can you configure a fallback source.

For device manager login, you can use Common Access Card (CAC) for login when using a SAML server if you configure the SAML server to support it.


Before you begin


Obtain the following information from the SAML server identity provider. If possible, download the information from the user in an XML file for easy upload.

- Entity ID URL, which provides the SAML server metadata.
- Sign-in URL.
- Sign-out URL.
- Identity provider certificate.

Procedure

- Step 1** Do one of the following to get to the SAML Servers page:
- Select **Objects**, then select **Identity Sources** from the table of contents.
 - Select **Device > Remote Access VPN > SAML Servers**.
- Step 2** Do one of the following:
- To create an object, click + > **SAML Server**.

- To edit an object, click the edit icon () for the object.

To delete an unreferenced object, click the trash can icon () for the object.

Step 3

Configure the following properties:

- **Name**—The name of the object.
- **Description**—An optional description of the object.
- **Identity Provider (IDP) Entity ID URL**—The URL for a page that serves a metadata XML that describes how the SAML Issuer is going to respond to requests. Some SAML server products call this the Entity ID, others call it the Metadata URL. The URL must be from 4-256 characters, including the protocol, https://. For example, https://191.168.2.21/dag/saml2/idp/metadata.php.

Note If you downloaded the information from the SAML server in an XML file, click **Populate from XML file** and select the file. This field plus **Sign-In URL** and **Identity Provider Certificate** can be populated from the XML file.

- **Sign-In URL**—The URL for signing into the identity provider SAML server. The URL must be between 4-500 characters, including protocol. Both http:// and https:// are allowed. For example, https://191.168.2.21/dag/saml2/idp/SSOService.php.
- **Sign-Out URL**—The URL for signing out of the identity provider SAML server. The URL must be between 4-500 characters, including protocol. Both http:// and https:// are allowed. For example, https://191.168.2.21/dag/saml2/idp/SingleLogoutService.php.
- **Service Provider Certificate**—The internal certificate to use for the threat defense device. Ideally, you have already uploaded a certificate signed by a recognized third party, and you can select it now. You can also use the built-in DefaultInternalCertificate, or click **Create New Internal Certificate** and upload a signed certificate now. The SAML server identity provider will have to trust this certificate, so you might need to upload it to the SAML server. Consult the SAML server documentation for information on how to upload certificates or otherwise enable a trust relationship with a service provider.
- **Identity Provider Certificate**—The trusted CA certificate for the SAML server identity provider. Download this certificate from the SAML server. If you have not already uploaded it, click **Create New Trusted CA Certificate** and upload it now.
- **Request Signature**—The encryption algorithm to use when signing login request. Select None to disable encryption. Otherwise, choose one of the following, which are ordered from weakest to strongest: SHA1, SHA256, SHA384, SHA512.
- **Request Timeout**—SAML assertions have a valid period: the user must complete the single sign-on request within the valid period. You can set a timeout, in seconds, to change this period. If you set a timeout that is longer than the assertion's NotOnOrAfter condition, your timeout is ignored and NotOnOrAfter is honored. The range is 1-7200 seconds. The default is 300 seconds.
- **This SAML identity provider (IDP) is on an internal network**—Whether the SAML server is operating on an internal network, rather than external to the protected networks.
- **Request IDP re-authentication at login**—Select this option for the user to re-authenticate at each login, rather than have the SAML server re-use a previous authentication session. This option is enabled by default.

Step 4

Click **User Roles** and configure the RBAC authorization roles for external users.

- **Default User Role**—The authorization role to assign a user if one cannot be otherwise determined by the settings on this page.
- **Group Member Attribute**—The user attribute in the SAML server that defines the RBAC authorization role for a user.
- **Role Mapping**—For each role, type in the string that will appear in the group member attribute in the SAML user records that should correspond to the role.
 - **Administrator**—Users who have full read-write access to all aspects of the application.
 - **Cryptographic Admin**—Users who can configure encryption-related features such as certificates, decryption policies, and secret keys. Read-only access to other features.
 - **Audit Admin**—Users who can view user login history and the audit log and perform auditing-related actions. Read-only access to configuration features.
 - **Read-Write**—Users who can do everything a read-only user can do, but also edit and deploy the configuration. The only restrictions are for system-critical actions, which include installing upgrades, creating and restoring backups, viewing the audit log, and ending the sessions of other device manager users.
 - **Read-Only**—Users who can view dashboards and the configuration, but cannot make any changes. If you try to make a change, the error message explains that this is due to lack of permission.

Step 5 Click **OK**.

What to do next

If you enabled **Request Signature** to encrypt communications, you need to upload the device manager information to the SAML server. From the identity sources list, click the **Download** (📄) button for the server, and save the XML file. Then, log into the SAML server and upload the information. See your SAML provider documentation for detailed information.

If you use the server for device manager login, and it does not work, verify the SAML server configuration.

- Log In into the SAML IdP and verify that the device manager SAML Response consumer is configured correctly. The value should be: `https://<FDM_URL>/api/fdm/latest/fdm/token`
- If signing is enabled in the SAML server object, make sure that the device manager public certificate is uploaded into the SAML application, and that encryption is enabled. Uploading the device manager XML file should add the certificate to the SAML server. You can also retrieve the device manager certificate through the FDM API: `https://<FDM_URL>/saml/metadata`

Local Users

The local user database (LocalIdentitySource) includes users that you have defined in the device manager.

You can use locally-defined users for the following purposes:

- Remote Access VPN, as a primary or fallback identity source.

- Management access, as a primary or secondary source for the device manager users.

The **admin** user is a system-defined local user. However, the admin user cannot log into a remote access VPN. You cannot create additional local administrative users.

If you define external authentication for management access, external users who log into the device appear on the local users list.

- Identity policy, indirectly, as a passive identity source to collect user identity from remote access VPN logins.

The following topic explains how to configure local users.

Configure Local Users

You can create user accounts directly on the device for use with remote access VPN. You can use the local user accounts instead of, or in addition to, an external authentication source.

If you use the local user database as a fallback authentication method for remote access VPN, ensure that you configure the same usernames/passwords in the local database as the names in the external database. Otherwise, the fallback mechanism will be ineffective.

The users defined here cannot log into the device CLI.


Procedure


Step 1 Select **Objects > Users**.

The list shows the usernames and service types, which can be:

- MGMT—For administrative users who can log into the device manager. The admin user is always defined, and you cannot delete it. You also cannot configure additional MGMT users. However, if you define external authentication for management access, external users who log into the device appear on the local users list as MGMT users.
- RA VPN—For users who can log into a remote access VPN configured on the device. You must also select the local database for the primary or secondary (fallback) source.

Step 2 Do one of the following:

- To add a user, click +.
- To edit a user, click the edit icon () for the user.

If you no longer need a particular user account, click the delete icon () for the user.

Step 3 Configure the user properties:

The name and password can contain any printable ASCII alphanumeric or special character except spaces and question marks. Printable characters are ASCII codes 33-126.

- **Name**—The username for logging into the remote access VPN. The name can be 4-64 characters, and it cannot contain spaces. For example, johndoe.

- **Password, Confirm Password**—Enter the password for the account. The password must be 8-16 characters long. It cannot contain consecutive letters that are the same. It must also contain at least one of each of the following: number, upper and lower case characters, and a special character.

Note Users cannot change their passwords. Notify them of their passwords, and when they need to change them, you must edit the user account. Also, do not update the password for external MGMT users: the passwords are controlled by the external AAA server.

Step 4 Click **OK**.
