

Identity Policies

You can use identity policies to collect user identity information from connections. You can then view usage based on user identity in the dashboards, and configure access control based on user or user group.

- Identity Policy Overview, on page 1
- How to Implement the Identity Policy, on page 3
- Configuring Identity Policies, on page 4
- Enabling Transparent User Authentication, on page 9
- Monitoring Identity Policies, on page 12
- Examples for Identity Policies, on page 13

Identity Policy Overview

You can use identity policies to detect the user who is associated with a connection. By identifying the user, you can correlate threat, endpoint, and network intelligence with user identity information. By linking network behavior, traffic, and events directly to individual users, the system can help you identify the source of policy breaches, attacks, or network vulnerabilities.

For example, you can identify who owns the host targeted by an intrusion event, and who initiated an internal attack or port scan. You can also identify high bandwidth users and users who are accessing undesirable web sites or applications.

User detection goes beyond collecting data for analysis. You can also write access rules based on user name or user group name, selectively allowing or blocking access to resources based on user identity.

You can obtain user identity using the following methods:

- Passive authentication—For all types of connections, obtain user identity from other authentication services without prompting for username and password.
- Active authentication—For HTTP connections only, prompt for username and password and authenticate against the specified identity source to obtain the user identity for the source IP address.

The following topics provide more information on user identity.

Establishing User Identity Through Passive Authentication

Passive authentication gathers user identity without prompting the user for username and password. The system obtains the mappings from the identity sources you specify.

You can passively obtain user-to-IP address mappings from the following sources:

- Remote access VPN logins. The following user types are supported for passive identity:
  - User accounts defined in an external authentication server.
  - Local user accounts that are defined in the Firewall Device Manager.
- Cisco Identity Services Engine (ISE); Cisco Identity Services Engine Passive Identity Connector (ISE PIC).

If a given user is identified through more than one source, the RA VPN identity takes precedence.

Establishing User Identity through Active Authentication

Authentication is the act of confirming the identity of a user.

With active authentication, when an HTTP traffic flow comes from an IP address for which the system has no user-identity mapping, you can decide whether to authenticate the user who initiated the traffic flow against the directory configured for the system. If the user successfully authenticates, the IP address is considered to have the identity of the authenticated user.

Failure to authenticate does not prevent network access for the user. Your access rules ultimately decide what access to provide these users.

Dealing with Unknown Users

When you configure the directory server for the identity policy, the system downloads user and group membership information from the directory server. This information is refreshed every 24 hours at midnight, or whenever you edit and save the directory configuration (even if you do not make any changes).

If a user succeeds in authenticating when prompted by an active authentication identity rule, but the user's name is not in the downloaded user identity information, the user is marked as Unknown. You will not see the user's ID in identity-related dashboards, nor will the user match group rules.

However, any access control rules for the Unknown user will apply. For example, if you block connections for Unknown users, these users are blocked even though they succeeded in authenticating (meaning that the directory server recognizes the user and the password is valid).

Thus, when you make changes to the directory server, such as adding or deleting users, or changing group membership, these changes are not reflected in policy enforcement until the system downloads the updates from the directory.

If you do not want to wait until the daily midnight update, you can force an update by editing the directory realm information (from **Objects** > **Identity Sources**, then edit the realm). Click **Save**, then deploy changes. The system will immediately download the updates.



Note You can check whether new or deleted user information is on the system by going to **Policies** > **Access Control**, clicking the **Add Rule (+)** button, and looking at the list of users on the **Users** tab. If you cannot find a new user, or you can find a deleted user, then the system has old information.

How to Implement the Identity Policy

To enable user identity acquisition, so that the user associated with an IP address is known, you need to configure several items. When configured correctly, you will be able to see usernames in the monitoring dashboards and events. You will also be able to use user identity in access control and SSL decryption rules as a traffic match criteria.

The following procedure provides an overview of what you must configure to get identity policies to work.

Procedure

Step 1 Configure the AD identity realm.

Whether you collect user identity actively (by prompting for user authentication) or passively, you need to configure the Active Directory (AD) server that has the user identity information. See Configuring AD Identity Realms.

If you configure passive identity, you can create AD realm sequences that allow the system to pull from identities in more than one AD realm. This is useful if you have multiple AD domains in your network.

Step 2 If you want to use passive authentication identity rules, configure the passive identity sources.

You can configure any of the following, based on the services you are implementing in the device and the services available to you in your network.

- Remote access VPN—If you intend to support remote access VPN connections to the device, user logins can provide the identity based on the AD server or on local users (those defined within the Firewall Device Manager). For information on configuring RA VPN, see Configuring Remote Access VPN.
- Cisco Identity Services Engine (ISE) or Cisco Identity Services Engine Passive Identity Connector (ISE PIC)—If you use these products, you can configure the device as a pxGrid subscriber, and obtain user identity from ISE. See Configure Identity Services Engine.

Step 3 Choose **Policies** > **Identity**, and enable the identity policy. See Configuring Identity Policies, on page 4.

Step 4 Configure Identity Policy Settings, on page 4.

The passive identity sources are automatically selected based on the sources you configured in the system. If you want to configure active authentication, you must configure the certificates for captive portal and SSL re-sign decryption (if you have not already enabled the SSL Decryption policy).

Step 5 Configure the Identity Policy Default Action, on page 6.

If your intention is to use passive authentication only, you can set the default action to passive authentication and there is no need to create specific rules.

Step 6 Configure Identity Rules, on page 6.

Create rules that will collect passive or active user identities from the relevant networks.

Configuring Identity Policies

You can use identity policies to collect user identity information from connections. You can then view usage based on user identity in the dashboards, and configure access control based on user or user group.

The following is an overview of how to configure the elements required to obtain user identity through identity policies.

Procedure

Step 1 Select **Policies** > **Identity**.

If you have not yet defined an identity policy, click **Enable Identity Policy** and configure settings as described in Configure Identity Policy Settings, on page 4.

Step 2 Manage the identity policy.

After you configure identity settings, this page lists all rules in order. Rules are matched against traffic from top to bottom with the first match determining the action to apply. You can do the following from this page:

- To enable or disable the identity policy, click the **Identity Policy** toggle.
- To change the identity policy settings, click the **Identity Policy Configuration** button ( ).
- To change the **Default Action**, click the action and select the desired action. See Configure the Identity Policy Default Action, on page 6.
- To move a rule, edit it and select the new location from the **Order** drop-down list.
- To configure rules:
  - To create a new rule, click the + button.
  - To edit an existing rule, click the edit icon ( ) for the rule (in the Actions column). You can also selectively edit a rule property by clicking on the property in the table.
  - To delete a rule you no longer need, click the delete icon ( ) for the rule (in the Actions column).

For more information on creating and editing identity rules, see Configure Identity Rules, on page 6.

Configure Identity Policy Settings

For identity policies to work, you must configure the sources that provide user identity information. The settings you must configure differ based on the type of rules you will configure: passive, active, or both.

The settings dialog box shows these settings in separate sections. Depending on how you access the dialog box, you will see both sections, or just one section. The dialog box appears automatically if you try to create a rule for an authentication type without having already configured the required settings.

The following procedure covers the full dialog box.

Before you begin

Ensure that time settings are consistent among the directory servers, Firewall 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 **Policies** > **Identity**.

Step 2 Click the **Identity Policy Configuration** button ( ).

Step 3 Configure the **Passive Authentication** options.

The dialog box shows you the passive authentication sources that you have already configured.

If necessary, you can configure ISE through this dialog box. If you have not configured an ISE object yet, you can click the **Integrate ISE** link and create it now. If the object exists, it is listed along with its state: Enabled or Disabled.

You must have configured at least one enabled passive identity source to create passive authentication rules.

Step 4 Configure the **Active Authentication** options.

When an identity rule requires active authentication for a user, the user is redirected to the captive portal port on the interface through which they are connected and then they are prompted to authenticate.

- **Server Certificate**—Select the internal certificate to present to users during active authentication. If you have not already created the required certificate, click **Create New Internal Certificate** from the bottom of the drop-down list .

  Users will have to accept the certificate if you do not upload a certificate that their browsers already trust.

- **Port**—The captive portal port. The default is 885 (TCP). If you configure a different port, it must be in the range 1025-65535.

  Note
  For the HTTP Basic, HTTP Response Page, and NTLM authentication methods, the user is redirected to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate.

Step 5 (Active authentication only.) In **Decrypt Re-Sign Certificate**, select the internal CA certificate to use for rules that implement decryption with re-signed certificates.

You can use the pre-defined NGFW-Default-InternalCA certificate, or one that you created or uploaded. If the certificate does not yet exist, click **Create Internal CA** to create it.

If you have not already installed the certificate in client browsers, click the download button ( ) to obtain a copy. See the documentation for each browser for information on how to install the certificate. Also see Downloading the CA Certificate for Decrypt Re-Sign Rules.

Note
You are prompted for SSL Decryption settings only if you have not already configured the SSL decryption policy. To change these settings after enabling the identity policy, edit the SSL decryption policy settings.

Step 6 Click **Save**.

Configure the Identity Policy Default Action

The identity policy has a default action, which is implemented for any connections that match no individual identity rules.

In fact, having no rules is a valid configuration for your policy. If you intend to use passive authentication on all traffic sources, then simply configure Passive Authentication as your default action.

Procedure

Step 1 Select **Policies** > **Identity**.

Step 2 Click in the **Default Action** and choose one of the following:

- **Passive Auth (Any Identity Source)**—User identity will be determined using all configured passive identity sources for connections that do not match any identity rules. If you do not configure any passive identity sources, using Passive Auth as the default is the same as using No Auth.
- **No Auth (No Authentication Required)**—User identity will not be determined for connections that do not match any identity rules.

Configure Identity Rules

Identity rules determine whether user identity information should be collected for matching traffic. You can configure No Authentication if you do not want to get user identity information for matching traffic.

Keep in mind that regardless of your rule configuration, active authentication is performed on HTTP traffic only. Thus, you do not need to create rules to exclude non-HTTP traffic from active authentication. You can simply apply an active authentication rule to all sources and destinations if you want to get user identity information for all HTTP traffic.

Note Also keep in mind that a failure to authentication has no impact on network access. Identity policies collect user identity information only. You must use access rules if you want to prevent users who failed to authenticate from accessing the network.

Before you begin

Rules are evaluated top-down. For a connection that matches the specified network criteria of a given rule, the user is evaluated against the identity realm specified in the rule. If the user is not part of that realm, they will be marked as unknown, and no further rules in the identity policy will be evaluated. Thus, if you have more than one realm that needs to be evaluated, be sure to use realm sequences instead of a single realm.

Procedure

Step 1 Select **Policies** > **Identity**.

Step 2 Do any of the following:

- To create a new rule, click the + button.
- To edit an existing rule, click the edit icon ( ) for the rule.

To delete a rule you no longer need, click the delete icon ( ) for the rule.

Step 3 In **Order**, select where you want to insert the rule in the ordered list of rules.

Rules are applied on a first-match basis, so you must ensure that rules with highly specific traffic matching criteria appear above policies that have more general criteria that would otherwise apply to the matching traffic.

The default is to add the rule to the end of the list. If you want to change a rule's location later, edit this option.

Step 4 In **Title**, enter a name for the rule.

Step 5 Select the **Action** and if necessary, **AD Identity Source**.

You must select the AD identity realm that includes the user accounts for passive and active authentication rules. If the realm you need does not yet exist, click **Create New Identity Realm** and create it now. For passive authentication, you can select an AD realm sequence rather than a single AD realm object.

- **Passive Auth**—Use passive authentication to determine user identity. All configured identity sources are shown. The rule automatically uses all configured sources.
- **Active Auth**—Use active authentication to determine user identity. Active authentication is applied to HTTP or decrypted HTTPS traffic only. If any other type of traffic matches an identity policy that requires or allows active authentication, then active authentication will not be attempted.
- **No Auth**—Do not obtain user identity. Identity-based access rules will not be applied to this traffic. These users are marked as **No Authentication Required**.

Step 6 (Active Authentication only.) Select the authentication method (**Type**) supported by your directory server.

- **HTTP Basic**—Authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication popup window. This is the default.
- **NTLM**—Authenticate users using an NT LAN Manager (NTLM) connection. This selection is only available when you select an AD realm. Users log in to the network using their browser's default authentication popup window, although you can configure IE and Firefox browsers to transparently authenticate using their Windows domain login (see Enabling Transparent User Authentication, on page 9).
- **HTTP Negotiate**—Allow the device to negotiate the method between the user agent (the application the user is using to initiate the traffic flow) and the Active Directory server. Negotiation results in the

strongest commonly supported method being used, in order, NTLM, then basic. Users log in to the network using their browser's default authentication popup window.

- **HTTP Response Page**—Prompt users to authenticate using a system-provided web page. This is a form of HTTP Basic authentication.

Note
For the HTTP Basic, HTTP Response Page, and NTLM authentication methods, the user is redirected to the captive portal using the IP address of the interface. However, for HTTP Negotiate, the user is redirected using the fully-qualified DNS name *firewall-hostname.AD-domain-name*. If you want to use HTTP Negotiate, you must also update your DNS server to map this name to the IP addresses of all inside interfaces where you are requiring active authentication. Otherwise, the redirection cannot complete, and users cannot authenticate.

Step 7 (Active authentication only.) Select **Fall Back as Guest** > **On/Off** to determine whether users who fail active authentication are labeled as Guest users.

Users get 3 chances to successfully authenticate. If they fail, your selection for this option determines how the user is marked. You can write access rules based on these values.

- **Fall Back as Guest** > **On**—Users are marked as **Guest**.
- **Fall Back as Guest** > **Off**—Users are marked as **Failed Authentication**.

Step 8 Define the traffic matching criteria on the **Source/Destination** tab.

Keep in mind that active authentication will be attempted with HTTP traffic only. Therefore, there is no need to configure No Auth rules for non-HTTP traffic, and there is no point in creating Active Authentication rules for any non-HTTP traffic. However, passive authentication is valid for any type of traffic.

The Source/Destination criteria of an identity rule define the security zones (interfaces) through which the traffic passes, the IP addresses or the country or continent (geographical location) for the IP address, or the protocols and ports used in the traffic. The default is any zone, address, geographical location, protocol, and port.

To modify a condition, you click the **+** button within that condition, select the desired object or element, and click **OK** in the popup dialog box. If the criterion requires an object, you can click **Create New *Object*** if the object you require does not exist. Click the **x** for an object or element to remove it from the policy.

You can configure the following traffic matching criteria.

Source Zones, Destination Zones

The security zone objects that define the interfaces through which the traffic passes. You can define one, both, or neither criteria: any criteria not specified applies to traffic on any interface.

- To match traffic leaving the device from an interface in the zone, add that zone to the **Destination Zones**.
- To match traffic entering the device from an interface in the zone, add that zone to the **Source Zones**.
- If you add both source and destination zone conditions to a rule, matching traffic must originate from one of the specified source zones and egress through one of the destination zones.

Use this criteria when the rule should apply based on where the traffic enters or exits the device. For example, if you want to ensure that user identity is collected from all traffic originating from inside networks, select an inside zone as the **Source Zones** while leaving the destination zone empty.

Note

You cannot mix passive and routed security zones in a single rule. In addition, you can specify passive security zones as source zones only, you cannot specify them as destination zones.

Source Networks, Destination Networks

The network objects or geographical locations that define the network addresses or locations of the traffic.

- To match traffic from an IP address or geographical location, configure the **Source Networks**.
- To match traffic to an IP address or geographical location, configure the **Destination Networks**.
- If you add both source and destination network conditions to a rule, matching traffic must originate from one of the specified IP addresses and be destined for one of the destination IP addresses.

When you add this criteria, you select from the following tabs:

- **Network**—Select the network objects or groups that define the source or destination IP addresses for the traffic you want to control.
- **Geolocation**—Select the geographical location to control traffic based on its source or destination country or continent. Selecting a continent selects all countries within the continent. Besides selecting geographical location directly in the rule, you can also select a geolocation object that you created to define the location. Using geographical location, you could easily restrict access to a particular country without needing to know all of the potential IP addresses used there.

  Note
  To ensure you are using up-to-date geographical location data to filter your traffic, Cisco strongly recommends that you regularly update the geolocation database (GeoDB).

Source Ports, Destination Ports/Protocols

The port objects that define the protocols used in the traffic. For TCP/UDP, this can include ports.

- To match traffic from a protocol or port, configure the **Source Ports**. Source ports can be TCP/UDP only.
- To match traffic to a protocol or port, configure the **Destination Ports/Protocols**.
- To match traffic both originating from specific TCP/UDP ports and destined for specific TCP/UDP ports, configure both. If you add both source and destination ports to a condition, you can only add ports that share a single transport protocol, TCP or UDP. For example, you could target traffic from port TCP/80 to port TCP/8080.

Step 9 Click **OK**.

Enabling Transparent User Authentication

If you configure the identity policy to allow for active authentication, you can use the following authentication methods to acquire user identity:

HTTP Basic

With HTTP basic authentication, users are always prompted to authenticate with their directory username and password. The password is transmitted in clear text. For that reason, basic authentication is not considered a secure form of authentication.

Basic is the default authentication mechanism.

HTTP Response Page

This is a type of HTTP basic authentication, where the user is presented with a login browser page.

NTLM, HTTP Negotiate (Integrated Windows Authentication for Active Directory)

With integrated Windows authentication, you take advantage of the fact that users log into a domain to use their workstation. The browser tries to use this domain login when accessing a server, including the Firewall Threat Defense captive portal during active authentication. The password is not transmitted. If authentication is successful, the user is transparently authenticated; the user is unaware that any authentication challenge was made or satisfied.

If the browser cannot satisfy an authentication request using the domain login credentials, the user is prompted for username and password, which is the same user experience as basic authentication. Thus, if you configure integrated Windows authentication, it can reduce the need for users to supply credentials when accessing the network or servers in the same domain.

Note that HTTP Negotiate picks the strongest method supported by both the Active directory server and the user agent. If negotiation selects HTTP Basic as the authentication method, you will not get transparent authentication. The order of strength is NTLM, then basic. Negotiation must select NTLM for transparent authentication to be possible.

You must configure client browsers to support integrated Windows authentication to enable transparent authentication. The following sections explain the general requirements and basic configuration of integrated Windows authentication for some commonly used browsers that support it. Users should consult the help for their browser (or other user agent) for more detailed information, because the techniques can change between software releases.



Tip Not all browsers support integrated Windows authentication, such as Chrome and Safari (based on the versions available when this was written). Users will be prompted for username and password. Consult the browser's documentation to determine if support is available in the version you use.

Requirements for Transparent Authentication

Users must configure their browser or user agent to implement transparent authentication. They can do this individually, or you can configure it for them and push the configuration to client workstations using your software distribution tools. If you decide to have users do it themselves, ensure that you provide the specific configuration parameters that work for your network.

Regardless of browser or user agent, you must implement the following general configuration:

- Add the Firewall Threat Defense interface through which users connect to the network to the Trusted Sites list. You can use the IP address or if available, the fully-qualified domain name (for example, inside.example.com). You can also use wildcards or partial addresses to create a generalized trusted site. For example, you can typically cover all internal sites using *.example.com or simply example.com, trusting all servers in your network (use your own domain name). If you add the specific address of the

interface, you might need to add several addresses to the trusted sites to account for all user access points to the network.

- Integrated Windows authentication does not work through a proxy server. Therefore, you must either not use a proxy, or you must add the Firewall Threat Defense interface to the addresses excluded from going through the proxy. If you decide that you must use a proxy, users will be prompted for authentication even if you use NTLM.

Tip Configuring transparent authentication is not a requirement, but a convenience to end users. If you do not configure transparent authentication, users are presented with a login challenge for all authentication methods.

Configuring Internet Explorer for Transparent Authentication

To configure Internet Explorer for NTLM transparent authentication:

Procedure

Step 1 Select **Tools** > **Internet Options**.

Step 2 Select the **Security** tab, select the **Local Intranet** zone, then do the following:

a) Click the **Sites** button to open the list of trusted sites.

b) Ensure that at least one of the following options is selected:

- **Automatically detect intranet network**. If you select this option, all other options are disabled.
- **Include all sites that bypass the proxy**.

c) Click **Advanced** to open the Local Intranet Sites dialog box, then paste the URL you want to trust into the **Add Site** box and click **Add**.

Repeat the process if you have more than one URL. Use wildcards to specify a partial URL, such as `http://*.example.com` or simply `*.example.com`.

Close the dialog boxes to return to the Internet Options dialog box.

d) With **Local Intranet** still selected, click **Custom Level** to open the Security Settings dialog box. Find the **User Authentication** > **Logon** setting and select **Automatic logon only in Intranet zone**. Click **OK**.

Step 3 In the Internet Options dialog box, click the **Connections** tab, then click **LAN Settings**.

If **Use a proxy server for your LAN** is selected, you need to ensure that the Firewall Threat Defense interface bypasses the proxy. Do any of the following as appropriate:

- Select **Bypass proxy server for local addresses**.
- Click **Advanced** and enter the address into the **Do not use proxy server for addresses beginning with** box. You can use wildcards, for example, `*.example.com`.

Configuring Firefox for Transparent Authentication

To configure Firefox for NTLM transparent authentication:

Procedure

Step 1 Open **about:config**. Use the filter bar to help you locate the preferences that you need to modify.

Step 2 To support NTLM, modify the following preferences (filter on network.automatic):

- **network.automatic-ntlm-auth.trusted-uris**—Double-click the preference, enter the URL, and click **OK**. You can enter multiple URLs by separating them with commas; including the protocol is optional. For example:

  ```
  http://host.example.com, http://hostname, myhost.example.com
  ```

  You can also use partial URLs. Firefox matches the end of the string, not a random substring. Thus, you could include your entire internal network by specifying just your domain name. For example:

  ```
  example.com
  ```

- **network.automatic-ntlm-auth.allow-proxies**—Ensure that the value is **true**, which is the default. Double-click to change the value if it is currently false.

Step 3 Check the HTTP proxy settings. You can find these by selecting **Tools** > **Options**, then click the **Network** tab in the Options dialog box. Click the **Settings** button in the Connection group.

- If **No Proxy** is selected, there is nothing to configure.
- If **Use System Proxy Settings** is selected, you need to modify the **network.proxy.no_proxies_on** property in about:config to add the trusted URIs you included in **network.automatic-ntlm-auth.trusted-uris**.
- If **Manual Proxy Configuration** is selected, update the **No Proxy For** list to include these trusted URIs.
- If one of the other options is selected, ensure that the properties used for those configurations exclude the same trusted URIs.

Monitoring Identity Policies

If identity policies that require authentication are working correctly, you should see user information on the **Monitoring** > **Users** dashboard and other dashboards that include user information.

In addition, events shown in **Monitoring** > **Events** should include user information.

If you do not see any user information, verify that the directory server is functioning correctly. Use the **Test** button in the directory server configuration dialog box to verify connectivity.

If the directory server is functioning and usable, verify that the traffic matching criteria on the identity rules that require active authentication are written in a way that will match your users. For example, ensure that the source zone contains the interfaces through which your user traffic will enter the device. The active

authentication identity rules match HTTP traffic only, so users must be sending that type of traffic through the device.

For passive authentication, use the **Test** button in the ISE object if you are using that source. If you are using remote access VPN, verify that the service is functioning correctly and that users can make VPN connections. See the troubleshooting topics for these features for more detailed information on identifying and resolving issues.

Examples for Identity Policies

The use case chapter includes an example of implementing identity policies. Please see How to Gain Insight Into Your Network Traffic.