



User Identity Policies

The following topics discuss how to create and manage identity rules and identity policies:

- [About Identity Policies, on page 1](#)
- [License Requirements for Identity Policies, on page 2](#)
- [Requirements and Prerequisites for Identity Policies, on page 2](#)
- [Create an Identity Policy, on page 3](#)
- [Identity Rule Conditions, on page 5](#)
- [Create an Identity Rule, on page 11](#)
- [Manage an Identity Policy, on page 13](#)
- [Manage an Identity Rule, on page 13](#)
- [Troubleshoot User Control, on page 14](#)

About Identity Policies

Identity policies contain identity rules. Identity rules associate sets of traffic with a realm and an authentication method: passive authentication, active authentication, or no authentication.

With the exception noted in the following paragraphs, you must configure realms and authentication methods you plan to use before you can invoke them in your identity rules:

- You configure realms outside of your identity policy, at **System > Integration > Realms**. For more information, see [Create an LDAP Realm or an Active Directory Realm and Realm Directory](#).
- You configure ISE/ISE-PIC, a passive authentication identity source, at **System > Integration > Identity Sources**.
- You configure the TS Agent, a passive authentication identity source, outside the system. For more information, see the *Cisco Terminal Services (TS) Agent Guide*.
- You configure captive portal, an active authentication identity source, within the identity policy. For more information, see [How to Configure the Captive Portal for User Control](#).
- You configure Remote Access VPN, an active authentication identity source, in Remote Access VPN policies. For more information, see [Remote Access VPN Authentication](#).

After you add multiple identity rules to a single identity policy, order the rules. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles the traffic.

You can also filter traffic by network object, which limits the network each device monitors in the event your devices are at or near their memory limits.

After you configure one or more identity policies, you must associate one identity policy with your access control policy. When traffic on your network matches the conditions in your identity rule, the system associates the traffic with the specified realm and authenticates the users in the traffic using the specified identity source.

If you do not configure an identity policy, the system does not perform user authentication.

Exception to creating an identity policy

An identity policy is not required if all of the following are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies. For more information, see [ISE SGT vs Custom SGT Rule Conditions](#).

Related Topics

[How to Set Up an Identity Policy](#)

License Requirements for Identity Policies

Threat Defense License

Any

Classic License

Control

Requirements and Prerequisites for Identity Policies

Model Support

Any.

Supported Domains

Any

User Roles

- Admin
- Access Admin
- Network Admin

Create an Identity Policy

This task discusses how to create an identity policy.

Before you begin

An identity policy is required to use users and groups in a realm in access control policies. Create and enable one or more realms as described in [Create an LDAP Realm or an Active Directory Realm and Realm Directory](#).

(Optional.) If a particular managed device monitors a large number of user groups, the system might drop user mappings based on groups due to managed device memory limitations. As a result, rules with realm or user conditions might not perform as expected. Provided the devices run version 6.7 or later, you can configure the identity rule to monitor traffic by one network or network group object only. To create a network object, see [Creating Network Objects](#).

An identity policy is not required if all of the following are true:

- You use the ISE/ISE-PIC identity source.
- You do not use users or groups in access control policies.
- You use Security Group Tags (SGT) in access control policies. For more information, see [ISE SGT vs Custom SGT Rule Conditions](#).

Procedure

-
- | | |
|---------------|--|
| Step 1 | Log in to the management center. |
| Step 2 | Click Policies > Access Control heading > Identity and click New Policy . |
| Step 3 | Enter a Name and, optionally, a Description . |
| Step 4 | Click Save . |
| Step 5 | To add a rule to the policy, click Add Rule as described in Create an Identity Rule, on page 11 . |
| Step 6 | To create a rule category, click Add Category . |
| Step 7 | To configure captive portal active authentication, click Active Authentication and see Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule . |
| Step 8 | (Optional.) To filter traffic by network object, click the Identity Source tab. From the list, click the network object to use to filter traffic for this identity policy. Click Add (+) to create a new network object. |
| Step 9 | Click Save to save the identity policy. |
-

What to do next

- Add rules to your identity policy that specify which users to match and other options; see [Create an Identity Rule, on page 11](#).
- Associate the identity policy with an access control policy to allow or block selected users from accessing specified resources; see [Associating Other Policies with Access Control](#).
- Deploy configuration changes to managed devices; see [Deploy Configuration Changes](#).

If you encounter issues, see [Troubleshoot User Control](#), on page 14.

Related Topics

[Configure the Captive Portal Part 2: Create an Identity Policy and Active Authentication Rule Captive Portal Fields](#)

[Troubleshoot User Control](#), on page 14

[Create an Identity Mapping Filter](#), on page 4

Create an Identity Mapping Filter

An identity mapping filter can be used to limit the networks to which an identity rule applies. For example, if your management center manages FTDs that have a limited amount of memory, you can limit the networks they monitor.

You must create separate identity mapping filters for IPv4 and IPv6 addresses.

You can also optionally exclude subnets from the following:

- Receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE.

You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.

Before you begin

Perform the following tasks:

1. Create a realm, which is required for an identity policy. See [Create an LDAP Realm or an Active Directory Realm and Realm Directory](#).
2. Create an identity policy. See [Create an Identity Policy](#), on page 3.
3. Create a network object or network group object as discussed in [Creating Network Objects](#). The network object or group you create should define the network you want managed devices to monitor in identity policies.

Procedure

-
- | | |
|---------------|--|
| Step 1 | Log in to the management center. |
| Step 2 | Click Policies > Identity . |
| Step 3 | Click Edit (✎). |
| Step 4 | Click the Identity Source tab. |
| Step 5 | From the Identity Mapping Filter list, choose the name of a network object to use as a filter . |

To create a new network object, see [Creating Network Objects](#).

Note

To restrict traffic to IPv6 addresses, you must add at least one address, network, or group to the filter.

- | | |
|---------------|---------------------|
| Step 6 | Click Save . |
|---------------|---------------------|

Step 7 Deploy configuration changes to managed devices; see [Deploy Configuration Changes](#).

What to do next

Associate the identity policy with an access control policy as discussed in [Associating Other Policies with Access Control](#).

To check or change ISE identity mapping filters (also referred to as *subnet filters*), use the following commands:

```
show identity-subnet-filter
configure identity-subnet-filter { add | remove } subnet
```

Identity Rule Conditions

Rule conditions enable you to fine-tune your identity policy to target the users and networks you want to control. See one of the following sections for more information.

Related Topics

[Security Zone Rule Conditions](#)

[Network Rule Conditions](#)

[VLAN Tags Rule Conditions](#)

[Port Rule Conditions](#)

[Realm & Settings Rule Conditions](#), on page 9

Security Zone Rule Conditions

Security zones segment your network to help you manage, classify, and decrypt traffic flow by grouping interfaces across multiple devices.

Security zones control or decrypt traffic by its source and destination security zones. If you add both source and destination zones to a zone condition, matching traffic must originate from an interface in one of the source zones and leave through an interface in one of the destination zones.

Just as all interfaces in a zone must be of the same type (all inline, passive, switched, or routed), all zones used in a zone condition must be of the same type. Because devices deployed passively do not transmit traffic, you cannot use a zone with passive interfaces as a destination zone.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.



Tip

Constraining rules by zone is one of the best ways to improve system performance. If a rule does not apply to traffic through any of device's interfaces, that rule does not affect that device's performance.

Security Zone Conditions and Multitenancy

In a multidomain deployment, a zone created in an ancestor domain can contain interfaces that reside on devices in different domains. When you configure a zone condition in an descendant domain, your configurations apply to only the interfaces you can see.

Network Rule Conditions

Networks control or decrypt traffic by its source and destination IP address, using inner headers. Tunnel rules, which use outer headers, have tunnel endpoint conditions instead of network conditions.

You can use predefined objects to build network conditions, or manually specify individual IP addresses or address blocks.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.



Note You *cannot* use FQDN network objects in identity rules.

Redirect to Host Name Network Rule Conditions

(Snort 3.0 only.)—You can use a network object that contains the fully-qualified host name (FQDN) of the interface that captive portal can use for active authentication requests.

The FQDN must resolve to the IP address of one of the interfaces on a managed device. By using an FQDN, you can assign a certificate for active authentication that the client will recognize, thus avoiding the untrusted certificate warning users get when being redirected to a managed device's IP address.

The certificate can specify one FQDN, a wildcard FQDN, or multiple FQDNs in the Subject Alternate Names (SAN) in the certificate.

If an identity rule requires active authentication for a user, but you do not specify a redirect FQDN, the user is redirected to the captive portal port on the managed device interface to which they are connected.

If you do not supply a Redirect to Host Name FQDN, the HTTP Basic, HTTP Response Page, and NTLM authentication methods redirect the user 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.directory-server-domain-name`. To use HTTP Negotiate without a Redirect to Host Name FQDN, 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.

We recommend that you always provide a Redirect to Host Name FQDN to ensure consistent behavior regardless of authentication method.

VLAN Tags Rule Conditions



Note VLAN tags in access rules only apply to inline sets. Access rules with VLAN tags do not match traffic on firewall interfaces.

VLAN rule conditions control VLAN-tagged traffic, including Q-in-Q (stacked VLAN) traffic. The system uses the innermost VLAN tag to filter VLAN traffic, with the exception of the prefilter policy, which uses the outermost VLAN tag in its rules.

Note the following Q-in-Q support:

- Firewall Threat Defense on Firepower 4100/9300—Does not support Q-in-Q (supports only one VLAN tag).
- Firewall Threat Defense on all other models:
 - Inline sets and passive interfaces—Supports Q-in-Q, up to 2 VLAN tags.
 - Firewall interfaces—Does not support Q-in-Q (supports only one VLAN tag).

You can use predefined objects to build VLAN conditions, or manually enter any VLAN tag from 1 to 4094. Use a hyphen to specify a range of VLAN tags.

In a cluster, if you encounter problems with VLAN matching, edit the access control policy advanced options, Transport/Network Preprocessor Settings, and select the **Ignore the VLAN header when tracking connections** option.

Port Rule Conditions

Port conditions allow you to control traffic by its source and destination ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic.

Application filtering is also recommended for applications, like Firewall Threat Defense, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as source port conditions in a single access control rule.

Port, Protocol, and ICMP Code Rule Conditions

Port conditions match traffic based on the source and destination ports. Depending on the rule type, “port” can represent any of the following:

- **TCP and UDP**—You can control TCP and UDP traffic based on the port. The system represents this configuration using the protocol number in parentheses, plus an optional associated port or port range. For example: TCP(6)/22.
- **ICMP**—You can control ICMP and ICMPv6 (IPv6-ICMP) traffic based on its internet layer protocol plus an optional type and code. For example: ICMP(1):3:3.
- **Protocol**—You can control traffic using other protocols that do not use ports.

Minimize the number of matching criteria whenever possible, especially those for security zones, network objects, and port objects. When you specify multiple criteria, the system must match against *every* combination of the contents of the criteria you specify.

Best Practices for Port-Based Rules

Specifying ports is the traditional way to target applications. However, applications can be configured to use unique ports to bypass access control blocks. Thus, whenever possible, use application filtering criteria rather than port criteria to target traffic. Note that application filtering is not available in prefilter rules.

Application filtering is also recommended for applications, like FTP, that open separate channels dynamically for control vs. data flow. Using port-based access control rules can prevent these kinds of applications from performing correctly, and could result in blocking desirable connections.

Using Source and Destination Port Constraints

If you add both source and destination port constraints, you can only add ports that share a single transport protocol (TCP or UDP). For example, if you add DNS over TCP as a source port, you can add Yahoo Messenger Voice Chat (TCP) as a destination port but not Yahoo Messenger Voice Chat (UDP).

If you add only source ports or only destination ports, you can add ports that use different transport protocols. For example, you can add both DNS over TCP and DNS over UDP as destination port conditions in a single access control rule.

Matching Non-TCP Traffic with Port Conditions

You can match non-port-based protocols. By default, if you do not specify a port condition, you are matching IP traffic. Although you can configure port conditions to match non-TCP traffic, there are some restrictions:

- **Access control rules**—For Classic devices, you can match GRE-encapsulated traffic with an access control rule by using the GRE (47) protocol as a destination port condition. To a GRE-constrained rule, you can add only network-based conditions: zone, IP address, port, and VLAN tag. Also, the system uses outer headers to match **all** traffic in access control policies with GRE-constrained rules. For Firewall Threat Defense devices, use tunnel rules in the prefilter policy to control GRE-encapsulated traffic.
- **Decryption rules**—These rules support TCP port conditions only.

- **ICMP echo**—A destination ICMP port with the type set to 0 or a destination ICMPv6 port with the type set to 129 only matches unsolicited echo replies. ICMP echo replies sent in response to ICMP echo requests are ignored. For a rule to match on any ICMP echo, use ICMP type 8 or ICMPv6 type 128.

Realm & Settings Rule Conditions

The **Realm & Settings** tab page enables you to choose a realm or realm sequence to which to apply the identity rule. If you are using captive portal, you have additional options.

Authentication Realm

From the **Realm** list, click a realm or realm sequence.

The realm or realm sequence containing the users you want to perform the specified **Action** on. You must fully configure a realm or realm sequence before selecting it as the realm in an identity rule.



Note If remote access VPN is enabled and your deployment is using a RADIUS server group for VPN authentication, make sure you specify the realm associated with this RADIUS server group.

Active authentication only: other options

If you either choose **Active Authentication** as the authentication type or if you check the box, **Use active authentication if passive or VPN identity cannot be established**, you have the following options.

Use active authentication if passive or VPN identity cannot be established

(Passive authentication rule only.) Selecting this option authenticates users using captive portal active authentication if a passive or a VPN authentication fails to identify them. You must configure an Active Authentication rule in your identity policy in order to select this option. (That is, users must authenticate using the captive portal.)

If you disable this option, users that do not have a VPN identity or that passive authentication cannot identify are identified as Unknown.

Also see the discussion of the **Authentication Realm** list later in this topic,

Identify as Special Identities/Guest if authentication cannot identify user

Selecting this option allows users who fail captive portal active authentication the specified number of times to access your network as a guest. These users appear in the Firewall Management Center identified by their username (if their username exists on the AD or LDAP server) or by **Guest** (if their user name is unknown). Their realm is the realm specified in the identity rule. (By default, the number of failed logins is 3.)

This field is displayed only if you configure **Active Authentication** (that is, captive portal authentication) as the rule **Action**.

Authentication Protocol

The method to use to perform captive portal active authentication. .

The selections vary depending on the type of realm, LDAP or AD:

- Choose **HTTP Basic** if you want to authenticate users using an unencrypted HTTP Basic Authentication (BA) connection. Users log in to the network using their browser's default authentication pop-up window.

Most web browsers cache the credentials from **HTTP Basic** logins and use the credentials to seamlessly begin a new session after an old session times out.

- Choose **NTLM** to authenticate users using a NT LAN Manager (NTLM) connection. This selection is available only when you select an AD realm. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication pop-up window.
- Choose **Kerberos** to authenticate users using a Kerberos connection. This selection is available only when you select an AD realm for a server with secure LDAP (LDAPS) enabled. If transparent authentication is configured in a user's browser, the user is automatically logged in. If transparent authentication is not configured, users log in to the network using their browser's default authentication pop-up window.



Note The **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.



Note If you are creating an identity rule to perform Kerberos captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the host name you provided when configuring DNS.

For Firewall Threat Defense devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.

- Choose **HTTP Negotiate** to allow the captive portal server to choose between HTTP Basic, Kerberos, or NTLM for the authentication connection. This type is available only when you select an AD realm.



Note The **Realm** you choose must be configured with an **AD Join Username** and **AD Join Password** for **HTTP Negotiate** to choose Kerberos captive portal active authentication.



Note If you are creating an identity rule to perform **HTTP Negotiate** captive portal and you have DNS resolution configured, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN of the device you are using for captive portal must match the hostname you provided when configuring DNS.

- Choose **HTTP Response Page** to enable users to choose a realm to log in to.

You can optionally customize the response page; for example, to conform to company style standards.

Create an Identity Rule

For details about configuration options for identity rules, see [Identity Rule Fields, on page 12](#).

Before you begin

You must create and enable a realm or realm sequence.

- Create a Microsoft Active Directory realm and realm directory as discussed in [Create an LDAP Realm or an Active Directory Realm and Realm Directory](#).
- Download users and groups and enable the realm as discussed in [Synchronize Users and Groups](#).
- (Optional.) Create a realm sequence as discussed in [Create a Realm Sequence](#).
- 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.



Caution

Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include a decryption policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See [Snort Restart Traffic Behavior](#) for more information.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

Procedure

- Step 1** Log in to the management center.
- Step 2** Click **Policies > Access Control heading > Identity**.
- Step 3** Click **Edit** (✎) next to the identity policy to which to add the identity rule.

If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Step 4** Click **Add Rule**.
- Step 5** Enter a **Name**.
- Step 6** If the Specified rule is applicable, check the check box of **Enabled**.

- Step 7** To add the rule to an existing category, indicate where you want to **Insert** the rule. To add a new category, click **Add Category**.
- Step 8** Choose a rule **Action** from the list.
- Step 9** If you're configuring captive portal, see [How to Configure the Captive Portal for User Control](#).
- Step 10** (Optional) To add conditions to the identity rule, see [Identity Rule Conditions, on page 5](#).
- Step 11** Click **Add**.
- Step 12** In the policy editor, set the rule position. Click and drag or use the right-click menu to cut and paste. Rules are numbered starting at 1. The system matches traffic to rules in top-down order by ascending rule number. The first rule that traffic matches is the rule that handles that traffic. Proper rule order reduces the resources required to process network traffic and prevents rule preemption.
- Step 13** Click **Save**.
-

What to do next

- Deploy configuration changes; see [Deploy Configuration Changes](#).

Identity Rule Fields

Use the following fields to configure identity rules.

Enabled

Enabling this option enables the identity rule in the identity policy. Unselecting this option disables the identity rule.

Action

Specify the type of authentication you want to perform on the users in the specified realm: **Passive Authentication** (default), **Active Authentication**, or **No Authentication**. You must fully configure the authentication method, or *identity source*, before selecting it as the action in an identity rule.

Additionally, if VPN is enabled (configured on at least one managed device), remote access VPN sessions are actively authenticated by VPN. Other sessions use the rule action. This means that, if VPN is enabled, VPN identity determination is performed first for all sessions regardless of the selected action. If a VPN identity is found on the specified realm, this is the identity source used. No additional captive portal active authentication is done, even if selected.

If the VPN identity source is not found, the process continues according to the specified action. You cannot restrict the identity policy to VPN authentication only because if the VPN identity is not found, the rule is applied according to the selected action.

**Caution**






Adding the first or removing the last active authentication rule when TLS/SSL decryption is disabled (that is, when the access control policy does not include an SSL policy) restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the assigned device handles traffic. See [Snort Restart Traffic Behavior](#) for more information.

Note that an active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive or VPN identity cannot be established** selected.

For information about which passive and active authentication methods are supported in your version of the system, see [About User Identity Sources](#).

Manage an Identity Policy

Procedure

-
- Step 1** Log in to the management center.
 - Step 2** Click **Policies > Access Control heading > Identity**.
 - Step 3** To delete a policy, click **Delete** (). If the controls are dimmed, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
 - Step 4** To edit a policy, click **Edit** () next to the policy and make changes as described in [Create an Identity Policy, on page 3](#). If **View** () appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
 - Step 5** To copy a policy, click **Copy** ().
 - Step 6** To generate a report for the policy, click **Report** () as described in [Generate Current Policy Reports](#).
 - Step 7** To compare policies, see [Compare Policies](#).
 - Step 8** To create a folder in which to organize policies, click **Add Category**.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#).

Manage an Identity Rule

Procedure

-
- Step 1** Log in to the management center.

- Step 2** Click **Policies > Access Control heading > Identity** .
- Step 3** Click **Edit** (✎) next to the policy you want to edit. If **View** (👁) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.
- Step 4** To edit an identity rule, click **Edit** (✎) and make changes as described in [Create an Identity Policy, on page 3](#).
- Step 5** To delete an identity rule, click **Delete** (🗑).
- Step 6** To create a rule category, click **Add Category** and choose the position and the rule.
- Step 7** Click **Save**.
-

What to do next

- Deploy configuration changes; see [Deploy Configuration Changes](#).

Troubleshoot User Control

If you notice unexpected user rule behavior, consider tuning your rule, identity source, or realm configurations. For other related troubleshooting information, see:

- [Troubleshoot the ISE/ISE-PIC or Cisco TrustSec Issues](#)
- [Troubleshoot the TS Agent Identity Source](#)
- [Troubleshoot the Captive Portal Identity Source](#)
- [Troubleshoot Realms and User Downloads](#)

Rules targeting realms, users, or user groups are not matching traffic

If you configure a TS Agent or ISE/ISE-PIC device to monitor a large number of user groups, or if you have a very large number of users mapped to hosts on your network, the system may drop user records due to your Firewall Management Center user limit. As a result, rules with user conditions may not match traffic as expected.

Rules targeting user groups or users within user groups are not matching traffic as expected

If you configure a rule with a user group condition, your LDAP or Active Directory server must have user groups configured. The system cannot perform user group control if the server organizes the users in basic object hierarchy.

Rules targeting users in secondary groups are not matching traffic as expected

If you configure a rule with a user group condition that includes or excludes users who are members of a secondary group on your Active Directory server, your server may be limiting the number of users it reports.

By default, Active Directory servers limit the number of users they report from secondary groups. You must customize this limit so that all of the users in your secondary groups are reported to the Firewall Management Center and eligible for use in rules with user conditions.

Rules are not matching users when seen for the first time

After the system detects activity from a previously-unseen user, the system retrieves information about them from the server. Until the system successfully retrieves this information, activity seen by this user is *not* handled by matching rules. Instead, the user session is handled by the next rule it matches (or the policy's default action, if applicable).

For example, this might explain when:

- Users who are members of user groups are not matching rules with user group conditions.
- Users who were reported by a TS Agent or ISE device are not matching rules, when the server used for user data retrieval is an Active Directory server.

Note that this might also cause the system to delay the display of user data in event views and analysis tools.

Rules are not matching all ISE/ISE-PIC users

This is expected behavior. You can perform user control on ISE/ISE-PIC users who were authenticated by an Active Directory domain controller. You cannot perform user control on ISE/ISE-PIC users who were authenticated by an LDAP, RADIUS, or RSA domain controller.

Users and groups using too much memory

If processing users and groups is using too much memory, health alerts are displayed. Remember that all user sessions are propagated to all devices managed by the Firewall Management Center. If your Firewall Management Center manages devices with different amounts of memory, the device with the least amount of memory determines the number of user sessions the system can handle without errors.

It's not possible to tune memory allocated to identity processes; even if a device has available memory, it can report out-of-memory issues. If issues persist, you have the following options:

- Segregate lower capacity managed devices on subnets and configure ISE/ISE-PIC to not report passive authentication data to those subnets.

See the chapter on managing network devices in the *Cisco Identity Services Engine Administrator Guide*.

- Unsubscribe from Security Group Tags (SGTs).
- Upgrade your managed device to a model with more memory.

