

User Identity Sources

The following topics describe Firepower System user *identity sources*, which are sources for *user awareness*. These users can be controlled with identity and access control policies:

- About User Identity Sources, on page 1
- The User Agent Identity Source, on page 2
- The ISE Identity Source, on page 5
- The Captive Portal Identity Source, on page 10
- The Traffic-Based Detection Identity Source, on page 23

About User Identity Sources

The following table provides a brief overview of the user identity sources supported by the Firepower System. Each identity source provides a store of users for user awareness. These users can then be controlled with identity and access control policies.

User Identity Source	Policy	Server Requirements	Type	Authentication Type	User Awareness?	User Control?	For more information, see...
User Agent	Identity	Microsoft Active Directory	Authoritative logins	Passive	Yes	Yes	The User Agent Identity Source, on page 2
ISE	Identity	Microsoft Active Directory	Authoritative logins	Passive	Yes	Yes	The ISE Identity Source, on page 5
Captive portal	Identity	Microsoft Active Directory	Authoritative logins	Active	Yes	Yes	The Captive Portal Identity Source, on page 10
Identity	RADIUS	Authoritative logins	Active	Yes	No		

User Identity Source	Policy	Server Requirements	Type	Authentication Type	User Awareness?	User Control?	For more information, see...
Traffic-based detection	Network discovery	n/a	Non-authoritative logins	n/a	Yes	No	The Traffic-Based Detection Identity Source, on page 23

Consider the following when selecting identity sources to deploy:

- You must use traffic-based detection for non-LDAP user logins. For example, if you are using only user agents to detect user activity, restricting non-LDAP logins has no effect.
- You must use traffic-based detection or captive portal to record failed login or authentication activity. A failed login or authentication attempt does not add a new user to the list of users in the database.
- The captive portal identity source requires a managed device with a routed interface. You *cannot* use an inline (also referred to as tap mode) interface with captive portal.

Data from those identity sources is stored in the Firepower Management Center's users database and the user activity database. You can configure Firepower Management Center-server user downloads to automatically and regularly download new user data to your databases.

After you configure identity rules using the desired identity source, you must associate each rule with an access control policy and deploy the policy to managed devices for the policy to have any effect. For more information about access control policies and deployment, see User, Realm, and ISE Attribute Conditions (User Control).

For general information about user identity in the Firepower System, see About User Identity.

Video icon YouTube videos for configuring identity sources.

The User Agent Identity Source

The Cisco Firepower User Agent is a passive authentication method; it is an authoritative identity source, meaning user information is supplied by a trusted Active Directory server. When integrated with the Firepower System, the user agent monitors users when they log in and out of hosts with Active Directory credentials. The data gained from the User Agent can be used for user awareness and user control.

The user agent associates each user with an IP address, which allows access control rules with user conditions to trigger. You can use one user agent to monitor user activity on up to five Active Directory servers and send encrypted data to up to five Firepower Management Centers.

The User Agent does not report failed login attempts.

Video User agent setup video on YouTube.

User Agent Guidelines

The User Agent requires a multi-step configuration that includes the following:

- At least one computer with the user agent installed.
- Connections between a Firepower Management Center and the computers or Active Directory servers with the user agent installed.
- An identity realm configured in each Firepower Management Center that receives user data from a user agent.

For detailed information about the multi-step User Agent configuration and a complete discussion of the server requirements, see the *Cisco Firepower User Agent Configuration Guide*.

Note Make sure the time on your computer or Active Directory server is synchronized with the time on the Firepower Management Center. If the appliances are not synchronized, the system might perform user timeouts at unexpected intervals.

The Firepower Management Center connection not only allows you to retrieve metadata for the users whose logins and logoffs were detected by User Agents, but also is used to specify the users and groups you want to use in access control rules. If the user agent is configured to exclude specific user names, login data for those user names are not reported to the Firepower Management Center. User agent data is stored in the user database and user activity database on the Firepower Management Center.

Note User Agents cannot transmit Active Directory user names ending with the `$` character to the Firepower Management Center. You must remove the final `$` character if you want to monitor these users.

If multiple users are logged into a host using remote sessions, the agent might not detect logins from that host properly. For information about how to prevent this, see the *Cisco Firepower User Agent Configuration Guide*.

Configure the User Agent for User Control

For more information about the User Agent, see The User Agent Identity Source, on page 2.

Before you begin

- Configure and enable an Active Directory realm for your User Agent connection as described in Create a Realm.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **System** > **Integration**.

Step 3 Click **Identity Sources**.

Step 4 Click **User Agent** for the **Service Type** to enable the User Agent connection.

Note To disable the connection, click **None**.

Step 5 Click **New Agent** to add a new agent.

Step 6 Enter the **Hostname** or **Address** of the computer where you plan to install the agent. You must use an IPv4 address; you cannot configure the Firepower Management Center to connect to a User Agent using an IPv6 address.

Step 7 Click **Add**.

Step 8 To delete a connection, click **Delete** ( ) and confirm that you want to delete it.

What to do next

- Continue User Agent setup as described in the *Cisco Firepower User Agent Configuration Guide*.
- Configure an identity rule as described in Create an Identity Rule.
- Associate the identity policy with an access control policy as discussed in Associating Other Policies with Access Control.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes.

Related Topics

Troubleshoot the User Agent Identity Source, on page 4
Access Control Policies

Troubleshoot the User Agent Identity Source

If you experience issues with the User Agent connection, see the *Cisco Firepower User Agent Configuration Guide*.

For related troubleshooting information in this guide, see Troubleshoot Realms and User Downloads and Troubleshoot User Control.

If you experience issues with user data reported by the User Agent, note:

- After the system detects activity from a User Agent user whose data is not yet in the database, the system retrieves information about them from the server. That user's activity is not handled by rules, and is not displayed in the web interface until the system successfully retrieves information about them in a user download.
- If you have Firepower Management Center high availability configured and the primary fails, all logins reported by a User Agent cannot be identified during failover downtime, even if the users were previously seen and downloaded to the Firepower Management Center. The unidentified users are logged as Unknown users on the Firepower Management Center. After the downtime, the Unknown users are re identified and processed according to the rules in your identity policy.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

The ISE Identity Source

You can integrate your Cisco Identity Services Engine (ISE) deployment with the Firepower System to use ISE for passive authentication.

ISE is an authoritative identity source, and provides user awareness data for users who authenticate using Active Directory (AD), LDAP, RADIUS, or RSA. Additionally, you can perform user control on Active Directory users. ISE does not report failed login attempts or the activity of ISE Guest Services users.

Note The Firepower System does not parse IEEE 802.1x machine authentication but it *does* parse 802.1x user authentication. If you are using 802.1x with ISE, you must include user authentication. 802.1x machine authentication will not provide a user identity to the FMC that can be used in policy.

For more information on Cisco ISE, see the *Cisco Identity Services Engine Administrator Guide*.

Note We strongly recommend you use the latest version of ISE to get the latest feature set and the most number of issue fixes.

How to Configure ISE for User Control

You can use ISE in any of the following configurations:

- With a realm, identity policy, and associated access control policy.

  Use a realm to control *user* access to network resources in policy. You can still use ISE Security Group Tags (SGT) metadata in your policies.

- With an access control policy only. No realm or identity policy are necessary.

  Use this method to control network access using SGT metadata alone.

ISE Guidelines and Limitations

Use the guidelines discussed in this section when configuring ISE with the Firepower System.

ISE Version and Configuration Compatibility

Your ISE version and configuration affects its integration and interaction with Firepower, as follows:

- We strongly recommend you use the latest version of ISE to get the latest feature set.
- Synchronize the time on the ISE server and the Firepower Management Center. Otherwise, the system might perform user timeouts at unexpected intervals.
- To implement user control using ISE data, configure and enable a realm for the ISE server assuming the pxGrid persona as described in Create a Realm.
- Each Firepower Management Center host name that connects to an ISE server must be unique; otherwise, the connection to one of the Firepower Management Centers will be dropped.

- Version 1.3 of ISE does not include support for IPv6-enabled endpoints. With this version of ISE, you cannot gather user identity data or perform remediations on IPv6-enabled endpoints.
- If ISE Endpoint Protection Service (EPS) is enabled and configured in your ISE deployment, you can use your ISE connection to run ISE EPS remediations on the source or destination host involved in a correlation policy violation.
- If you configured your ISE deployment to update a user's SGT after the user's EPSStatus changes, your ISE EPS remediations also update the SGT on the Firepower Management Center.

For the specific versions of ISE that are compatible with this version of the system, see the *Cisco Firepower Compatibility Guide*.

IPv6 support

Version 2.0 (patch 4) and later of ISE includes support for IPv6-enabled endpoints.

Approve clients in ISE

Before a connection between the ISE server and the Firepower Management Center succeeds, you must manually approve the clients in ISE. (Typically, there are two clients: one for the connection test and another for ISE agent.)

You can also 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*.

Security Group Tags (SGT)

A Security Group Tag (SGT) specifies the privileges of a traffic source within a trusted network. Cisco ISE and Cisco TrustSec use a feature called Security Group Access (SGA) to apply SGT attributes to packets as they enter the network. These SGTs correspond to a user's assigned security group within ISE or TrustSec. If you configure ISE as an identity source, the Firepower System can use these SGTs to filter traffic.



Note In some rules, custom SGT conditions can match traffic tagged with SGT attributes that were *not* assigned by ISE. This is not considered user control, and works only if you are not using ISE as an identity source; see Custom SGT Conditions.

ISE and High Availability

When the primary Firepower Management Center fails, the following occur:

- Until the standby is promoted to primary, the user database on the secondary Firepower Management Center is read-only.

  Users added to the repository (for example, Active Directory) are not downloaded to the Firepower Management Center and those users are identified as Unknown.

  New SGTs are not used.

- After the standby is promoted to primary, all operations return to normal; that is, users are downloaded, new SGTs are used, and users are identified if possible.

When the ISE primary server fails, you must manually promote the secondary to primary; there is no automatic failover.

Endpoint Location (or Location IP)

An Endpoint Location attribute is the IP address of the network device that used ISE to authenticate the user, as identified by ISE.

ISE Attributes

Configuring an ISE connection populates the Firepower Management Center database with ISE attribute data. You can use the following ISE attributes for user awareness and user control.

Endpoint Profile (or Device Type)

An Endpoint Profile attribute is the user's endpoint device type, as identified by ISE.

Configure ISE for User Control

The following procedure discusses how to configure the ISE identity source. You must be in the global domain to perform this task.

Before you begin

- To get user sessions from a Microsoft Active Directory Server or supported LDAP server, configure and enable a realm for the ISE server, assuming the pxGrid persona, as discussed in Create a Realm.

Procedure

Step 1 Log in to the Firepower Management Center.

Step 2 Click **System** > **Integration**.

Step 3 Click **Identity Sources**.

Step 4 Click **Identity Services Engine** for the **Service Type** to enable the ISE connection.

Note To disable the connection, click **None**.

Step 5 Enter a **Primary Host Name/IP Address** and, optionally, a **Secondary Host Name/IP Address**.

Step 6 Click the appropriate certificate authorities from the **pxGrid Server CA** and **MNT Server CA** lists, and the appropriate certificate from the **FMC Server Certificate** list. You can also click **Add** ( ) to add a certificate.

Note The **FMC Server Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.

Step 7 (Optional.) Enter an **ISE Network Filter** using CIDR block notation.

Step 8 To test the connection, click **Test**.

If the test fails, click **Additional Logs** for more information about the connection failure.

Note When you run two ISE pxGrid 1.0 nodes, it is normal for one host to show Success and one to show Failure. Because pxGrid 1.0 only runs actively on one ISE node at a time, the likelihood of success depends on which node in ISE is the active pxGrid node.

What to do next

- Specify users to control and other options using an identity policy as described in Create an Identity Policy.
- Associate the identity rule with an access control policy, which filters and optionally inspects traffic, as discussed in Associating Other Policies with Access Control.
- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes.
- Monitor user activity as discussed in Using Workflows.

Related Topics

Troubleshoot the Captive Portal Identity Source, on page 22
Trusted Certificate Authority Objects
Internal Certificate Objects

ISE Configuration Fields

The following fields are used to configure a connection to ISE.

Primary and Secondary Host Name/IP Address

The hostname or IP address for the primary and, optionally, the secondary pxGrid ISE servers.

The ports used by the host names you specify must be reachable by both ISE and the Firepower Management Center.

pxGrid Server CA

The certificate authority 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

The certificate authority for the ISE certificate when performing bulk downloads. If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.

FMC Server Certificate

The certificate and key that the Firepower Management Center must provide to ISE to connect to ISE or to perform bulk downloads.



Note The **FMC Server Certificate** must include the clientAuth extended key usage value, or it must not include any extended key usage values.

ISE Network Filter

An optional filter you can set to restrict the data that ISE reports to the Firepower Management Center. If you provide a network filter, ISE reports data from the networks within that filter. You can specify a filter in the following ways:

- Leave the field blank to specify **`any`**.
- Enter a single IPv4 address block using CIDR notation.

- Enter a list of IPv4 address blocks using CIDR notation, separated by commas.

Note This version of the Firepower System does not support filtering using IPv6 addresses, regardless of your ISE version.

Related Topics

Trusted Certificate Authority Objects
Internal Certificate Objects

Troubleshoot ISE or Cisco TrustSec Issues

Troubleshoot Cisco TrustSec issues

A device interface can be configured to propagate Security Group Tags (SGTs) either from ISE or from a Cisco device on the network (referred to as Cisco TrustSec.) On the device management page (**Devices** > **Device Management**), the **Propagate Security Group Tag** check box for an interface is checked after a device reboot. If you do not want the interface to propagate TrustSec data, uncheck the box.

FMC health monitor issue

The ISE Connection Status Monitor (health monitor) displays `check connectivity error` if ISE uses pxgrid v1 even though there is nothing wrong with the connection.

Troubleshoot ISE issues

For other related troubleshooting information, see Troubleshoot Realms and User Downloads and Troubleshoot User Control.

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

- The pxGrid Identity Mapping feature in ISE must be enabled before you can successfully integrate ISE with the Firepower System.
- When the primary server fails, you must manually promote the secondary to primary; there is no automatic failover.
- Before a connection between the ISE server and the Firepower Management Center succeeds, you must manually approve the clients in ISE. (Typically, there are two clients: one for the connection test and another for ISE agent.)

  You can also 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 **FMC Server Certificate** must include the **clientAuth** extended key usage value, or it must not include any extended key usage values.
- The time on your ISE server must be synchronized with the time on the Firepower Management Center. If the appliances are not synchronized, the system may perform user timeouts at unexpected intervals.
- If your deployment includes a primary and a secondary pxGrid node,
  - The certificates for both nodes must be signed by the same certificate authority.

  - The ports used by the host name must be reachable by both the ISE server and by the Firepower Management Center.

- If your deployment includes a primary and a secondary MNT node, the certificates for both nodes must be signed by the same certificate authority.
- If you updated to Version 6.1.x from Version 6.0.x and you are experiencing issues with your ISE connection, check your pxGrid server certificate. Version 6.1 is compliant with RFC6125-6.4.4, which states that certificate CNs should be ignored if there are SAN values specified. If the pxGrid server certificate in your ISE deployment is configured with a CN value and one or more SAN values, remove the CN value and add it as an additional SAN.

To exclude subnets from receiving user-to-IP and Security Group Tag (SGT)-to-IP mappings from ISE, use the **configure identity-subnet-filter** {**add** | **remove**} command. You should typically do this for lower-memory managed devices to prevent Snort identity health monitor memory errors.

If you experience issues with user data reported by ISE, 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 web interface 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 Firepower Management Center does not receive user data for ISE Guest Services users.
- Your ISE version and configuration impact how you can use ISE in the Firepower System. For more information, see The ISE Identity Source, on page 5.
- If you have Firepower Management Center high availability configured and the primary fails, see the section on ISE and High Availability in ISE Guidelines and Limitations, on page 5.
- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.

If you experience issues with supported functionality, see The ISE Identity Source, on page 5 for more information about version compatibility.

Troubleshoot ISE user timeout

If you're setting up ISE/ISE-PIC without a realm, be aware there is a user session timeout that affects how users are seen by the Firepower Management Center. For more information, see Realm Fields.

The Captive Portal Identity Source

Captive portal is one of the authoritative identity sources supported by the Firepower System. It is the only active authentication method supported by the Firepower System, where users can authenticate onto the network using a managed device.

You typically use captive portal to require authentication to access the internet or to access restricted internal resources; you can optionally configure guest access to resources. After the system authenticates captive portal

users, it handles their user traffic according to access control rules. Captive portal performs authentication on HTTP and HTTPS traffic only.

Note HTTPS traffic must be decrypted before captive portal can perform authentication.

Captive portal also records failed authentication attempts. A failed attempt does not add a new user to the list of users in the database. The user activity type for failed authentication activity reported by captive portal is **Failed Auth User**.

The authentication data gained from captive portal can be used for user awareness and user control.

Related Topics

How to Configure the Captive Portal for User Control, on page 13

Captive Portal Guidelines and Limitations

When you configure and deploy captive portal in an identity policy, users from specified realms authenticate through the following device to access your network:

- Virtual routers on 7000 and 8000 Series devices
- ASA FirePOWER devices in routed mode running Version 9.5(2) or later
- Firepower Threat Defense devices in routed mode

Routed Interface Required

Captive portal active authentication can be performed only by a device with a routed interface configured. If you are configuring the rule for captive portal and your captive portal device contains inline and routed interfaces, you must configure an interface condition to target only the routed interfaces on the device.

If the identity policy referenced by your access control policy contains one or more captive portal identity rules and you deploy the policy on a Firepower Management Center that manages:

- One or more devices with routed interfaces configured, the policy deployment succeeds and the routed interfaces perform active authentication.

  The system does not validate the type of interface in ASA with FirePOWER devices. If you apply a captive portal policy to an inline (tap mode) interface on an ASA with FirePOWER device, the policy deployment succeeds but users in traffic matching those rules are identified as Unknown.

- One or more NGIPSv devices, the policy deployment fails.

Captive Portal and Policies

You configure captive portal in your identity policy and invoke active authentication in your identity rules. Identity policies are associated with access control policies.

You configure some captive portal identity policy settings on the access control policy's **Active Authentication** tab page and configure the rest in an identity rule associated with the access control policy.

An active authentication rule has either an **Active Authentication** rule action, or a **Passive Authentication** rule action with **Use active authentication if passive authentication cannot identify user** selected. In each case the system transparently enables or disables SSL decryption, which restarts the Snort process.

Captive portal authenticates any user in the associated realm, even if that user does not belong to a downloaded group. The system identifies users in non-downloaded groups as Unknown; Unknown users match no identity rules. To avoid that, configure the realm to download users in all groups you expect to authenticate with captive portal.

To make sure the system downloads all users in a realm, make sure the groups are in the Available Groups list in the realm's configuration.

For more information about users and groups, see Download Users and Groups.

Caution Adding the first or removing the last active authentication rule when 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 target device handles traffic. See Snort® Restart Traffic Behavior for more information.

Captive Portal Requirements and Limitations

Note the following requirements and limitations:

- The system supports up to 20 captive portal logins per second.
- There is a maximum five minute limit between failed login attempts for a failed login attempt to be counted toward the count of maximum login attempts. The five minute limit is not configurable.

  (Maximum login attempts are displayed in connection events: **Analysis** > **Connections** > **Events**.)

  If more than five minutes elapse between failed logins, the user will continue to be redirected to captive portal for authentication, will not be designated a failed login user or a guest user, and will not be reported to the Firepower Management Center.
- The only way to be sure a user logs out is to close and reopen the browser. Unless that happens, in some cases, the user can log out of captive portal and be able to access the network without authenticating again using the same browser.
- If a realm is created for a parent domain and the managed device detects a login to a child of that parent domain, the user's subsequent logout is not detected by the managed device.
- To use an ASA FirePOWER device (in routed mode and running ASA version 9.5(2) or later) for captive portal, use the **captive-portal** ASA CLI command to enable captive portal for active authentication and define the port as described in the *ASA Firewall Configuration Guide* (Version 9.5(2) or later): https://www.cisco.com/c/en/us/support/security/adaptive-security-appliance-asa-software/products-installation-and-configuration-guides-list.html.
- You must allow traffic destined for the IP address and port of the device you plan to use for captive portal.
- To perform captive portal active authentication on HTTPS traffic, you must use an SSL policy to decrypt the traffic from the users you want to authenticate. You cannot decrypt the traffic in the connection between a captive portal user's web browser and the captive portal daemon on the managed device; this connection is used to authenticate the captive portal user.

- To limit the amount of non-HTTP or HTTPS traffic that is allowed through the managed device, you should enter typical HTTP and HTTPS ports in the identity policy's **Ports** tab page.

  The managed device changes a previously unseen user from **Pending** to **Unknown** when it determines that the incoming request does not use the HTTP or HTTPS protocol. As soon as the managed device changes a user from **Pending** to another state, access control, Quality of Service, and SSL policies can be applied to that traffic. If your other policies don't permit non-HTTP or HTTPS traffic, configuring ports on the captive portal identity policy can prevent undesired traffic from being allowed through the managed device.

How to Configure the Captive Portal for User Control

High-level overview of how to control user activity with captive portal:

Before you begin

To use the captive portal for active authentication, you must set up an AD or LDAP realm, access control policy, an identity policy, an SSL policy, and associate the identity and SSL policies with the access control policy. Finally, you must deploy the policies to managed devices. This topic provides a high-level summary of those tasks.

An example of the entire procedure begins in Configure the Captive Portal Part 1: Create an Identity Policy, on page 14.

Perform the following tasks first:

- Confirm that your Firepower Management Center manages one or more devices with a routed interface configured.

  In particular, if your Firepower Management Center manages ASA with FirePOWER devices, see Captive Portal Guidelines and Limitations, on page 11.

- To use encrypted authentication with the captive portal, either create a PKI object or have your certificate data and key available on the machine from which you're accessing the Firepower Management Center. To create a PKI object, see PKI Objects.

Procedure

Step 1 Create and enable a realm as discussed in the following topics:

- Configure a Realm Directory
- Download Users and Groups

Captive portal authenticates any user in the associated realm, even if that user does not belong to a downloaded group. The system identifies users in non-downloaded groups as Unknown; Unknown users match no identity rules. To avoid that, configure the realm to download users in all groups you expect to authenticate with captive portal.

To make sure the system downloads all users in a realm, make sure the groups are in the Available Groups list in the realm's configuration.

For more information about users and groups, see Download Users and Groups.

Step 2 Create an active authentication identity policy for captive portal.

The identity policy enables selected users in your realm access resources after authenticating with the captive portal.

For more information, see Configure the Captive Portal Part 1: Create an Identity Policy, on page 14.

Step 3 Configure an access control rule for the captive portal that allows traffic on the captive portal port (by default, TCP 885).

You can choose any available TCP port for the captive portal to use. Whatever your choice, you must create a rule that allows traffic on that port.

For more information, see Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule, on page 16.

Step 4 Add another access control rule to allow users in the selected realms to access resources using the captive portal.

This enables users to authenticate with captive portal.

For more information, see Configure the Captive Portal Part 3: Create a User Access Control Rule, on page 17.

Step 5 Configure an SSL decrypt - resign policy for the **Unknown** user so captive portal users can access web pages using the HTTPS protocol.

The captive portal can authenticate users only if the HTTPS traffic is decrypted before the traffic is sent to the captive portal. Captive portal is seen by the system as the **Unknown** user.

For more information, see Configure Captive Portal Part 4: Create an SSL Decrypt-Resign Policy, on page 18.

Step 6 Asscociate the identity and SSL policies with the access control policy from step 2.

This final step enables the system to authenticate users with the captive portal.

For more information, see Configure Captive Portal Part 5: Associate Identity and SSL Policies with the Access Control Policy, on page 19.

What to do next

See Configure the Captive Portal Part 1: Create an Identity Policy, on page 14.

Related Topics

Exclude Applications from Captive Portal, on page 21
Internal Certificate Objects
Troubleshoot the Captive Portal Identity Source, on page 22
Snort® Restart Scenarios

Configure the Captive Portal Part 1: Create an Identity Policy

Before you begin

This five-part procedure shows how to set up the captive portal using the default TCP port 885 and using a Firepower Management Center server certificate for both the captive portal and for SSL decryption. Each part of this example explains one task required to enable the captive portal to perform active authentication.

If you follow all the steps in this procedure, you can configure captive portal to work for users in your domains. You can optionally perform additional tasks, which are discussed in each part of the procedure.

For an overview of the entire procedure, see How to Configure the Captive Portal for User Control, on page 13.

Procedure

Step 1 Log in to the Firepower Management Center if you have not already done so.

Step 2 Click **Policies** > **Access Control** > **Identity** and create or edit an identity policy.

Step 3 (Optional.) Click **Add Category** to add a category for the captive portal identity rules and enter a **Name** for the category.

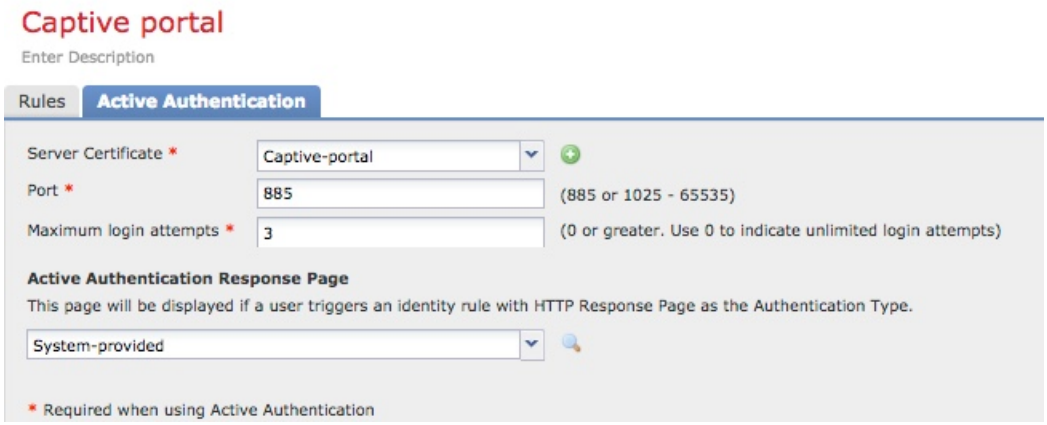
Step 4 Click **Active Authentication**.

Step 5 Choose the appropriate **Server Certificate** from the list or click **Add** ( ) to add a certificate.

Note Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

Step 6 Enter **885** in the **Port** field and specify the **Maximum login attempts**.

Step 7 (Optional.) Choose an **Active Authentication Response Page** as described in Captive Portal Fields, on page 20.
The following figure shows an example.



Step 8 Click **Save**.

Step 9 Click **Rules**.

Step 10 Click **Add Rule** to add a new captive portal identity policy rule, or click **Edit** ( ) to edit an existing rule.

Step 11 Enter a **Name** for the rule.

Step 12 From the **Action** list, choose **Active Authentication**.

The system can enforce captive portal active authentication on HTTP and HTTPS traffic only. If an identity rule **Action** is **Active Authentication** (you are using captive portal) or if you are using passive authentication and you check the option on **Realms & Settings** page to **Use active authentication if passive authentication cannot identify user**, use TCP ports constraints only.

Step 13 Click **Realm & Settings**.

Step 14 From the **Realms** list, choose a realm to use for user authentication.

Step 15 (Optional.) Check **Identify as Guest if authentication cannot identify user**. For more information, see Captive Portal Fields, on page 20.

Step 16 Choose an **Authentication Type** from the list.

Step 17 (Optional.) To exempt specific application traffic from captive portal, see Exclude Applications from Captive Portal, on page 21.

Step 18 Add conditions to the rule (port, network, and so on) as discussed in Rule Condition Types.

Step 19 Click **Add**.

Step 20 At the top of the page, click **Save**.

What to do next

Continue with Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule, on page 16.

Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule

This part of the procedure shows how to create an access control rule that allows the captive portal to communicate with clients using TCP port 885, which is the captive portal's default port. You can choose another port if you wish, but the port must match the one you chose in Configure the Captive Portal Part 1: Create an Identity Policy, on page 14.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 13.

Procedure

Step 1 Log in to the Firepower Management Center if you have not already done so.

Step 2 If you haven't done so already, create a certificate for the captive portal as discussed in PKI Objects.

Step 3 Click **Policies** > **Access Control** > **Access Control** and create or edit an access control policy.

Step 4 Click **Add Rule**.

Step 5 Enter a **Name** for the rule.

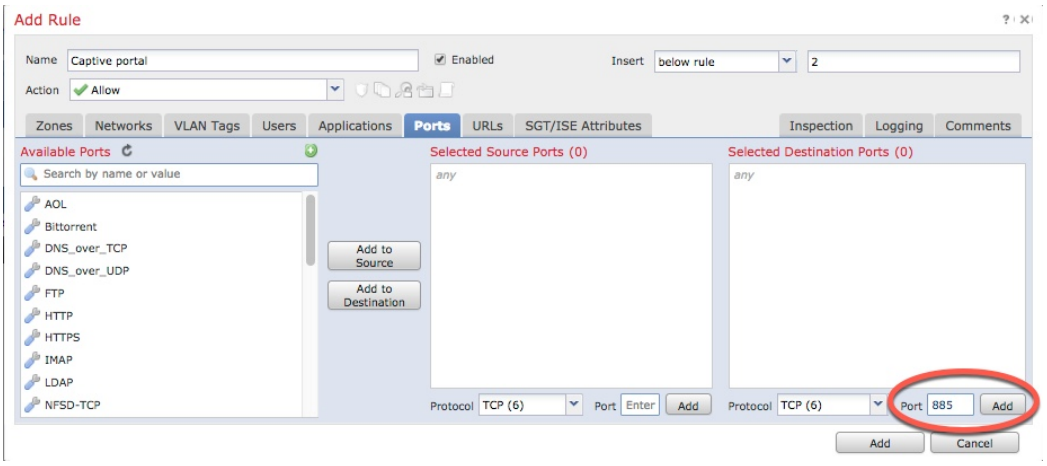
Step 6 Choose **Allow** from the **Action** list.

Step 7 Click **Ports**.

Step 8 From the **Protocol** list under the **Selected Destination Ports** field, choose **TCP**.

Step 9 In the **Port** field, enter **885**.

Step 10 Click **Add** next to the **Port** field.
The following figure shows an example.



Step 11 Click **Add** at the bottom of the page.

What to do next

Continue with Configure the Captive Portal Part 3: Create a User Access Control Rule, on page 17.

Configure the Captive Portal Part 3: Create a User Access Control Rule

This part of the procedure discusses how to add an access control rule that enables users in a realm to authenticate using captive portal.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 13.

Procedure

Step 1 In the rule editor, click **Add Rule**.

Step 2 Enter a **Name** for the rule.

Step 3 Choose **Allow** from the **Action** list.

Step 4 Click **Users**.

Step 5 In the **Available Realms** list, click the realms to allow.

Step 6 If no realms display, click **Refresh** ( ).

Step 7 In the **Available Users** list, choose the users to add to the rule and click **Add to Rule**.

Step 8 (Optional.) Add conditions to the access control policy as discussed in Rule Condition Types.

Step 9 Click **Add**.

Step 10 On the access control rule page, click **Save**.

Step 11 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.

What to do next

Continue with Configure Captive Portal Part 4: Create an SSL Decrypt-Resign Policy, on page 18.

Configure Captive Portal Part 4: Create an SSL Decrypt-Resign Policy

This part of the procedure discusses how to create an SSL access policy to decrypt and resign traffic before the traffic reaches the captive portal. The captive portal can authenticate traffic only after it has been decrypted.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 13.

Procedure

Step 1 If you haven't done so already, create a certificate object to decrypt SSL traffic as discussed in PKI Objects.

Step 2 Click **Policies** > **Access Control** > **SSL**.

Step 3 Click **New Policy**.

Step 4 Enter a **Name** and choose a **Default Action** for the policy. Default actions are discussed in SSL Policy Default Actions.

Step 5 Click **Save**.

Step 6 Click **Add Rule**.

Step 7 Enter a **Name** for the rule.

Step 8 From the **Action** list, choose **Decrypt - Resign**.

Step 9 From the **with** list, choose your PKI object.

Step 10 Click **Users**.

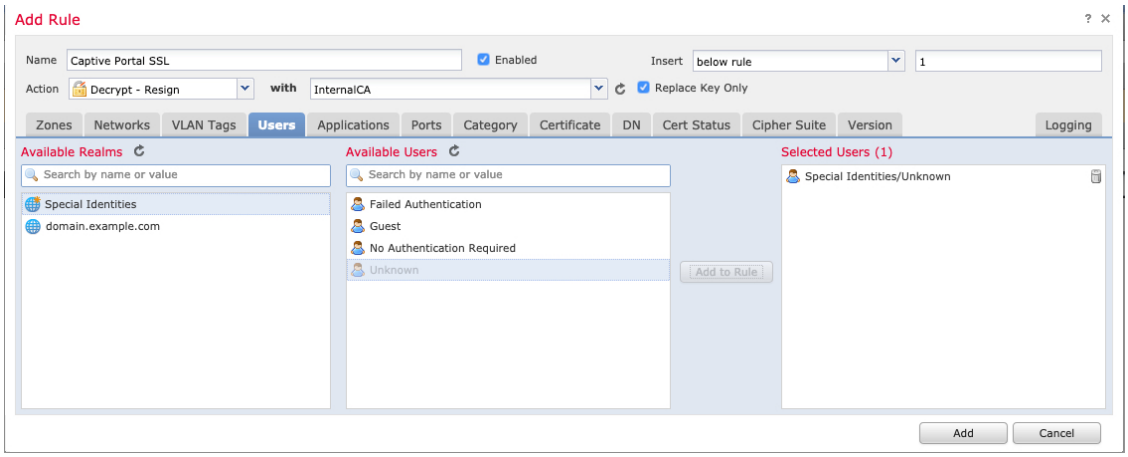
Step 11 Above the **Available Realms** list, click **Refresh** ( ).

Step 12 In the **Available Realms** list, click **Special Identities**.

Step 13 In the **Available Users** list, click **Unknown**.

Step 14 Click **Add to Rule**.

The following figure shows an example.



Step 15 (Optional.) Set other options as discussed in TLS/SSL Rule Conditions.

Step 16 Click **Add**.

Step 17 At the top of the page, click **Save**.

What to do next

Continue with Configure Captive Portal Part 5: Associate Identity and SSL Policies with the Access Control Policy, on page 19.

Configure Captive Portal Part 5: Associate Identity and SSL Policies with the Access Control Policy

This part of the procedure discusses how to associate the identity policy and SSL **Decrypt - Resign** rule with the access control policy you created earlier. After this, users can authenticate using the captive portal.

Before you begin

For an overview of the entire captive portal configuration, see How to Configure the Captive Portal for User Control, on page 13.

Procedure

Step 1 Click **Policies** > **Access Control** > **Access Control** and edit the access control policy you created as discussed in Configure the Captive Portal Part 2: Create a TCP Port Access Control Rule, on page 16. If **View** ( ) appears instead, the configuration belongs to an ancestor domain, or you do not have permission to modify the configuration.

Step 2 Either create a new access control policy or edit an existing policy.

Step 3 At the top of the page, click the link next to **Identity Policy**.

Step 4 From the list, choose the name of your identity policy and, at the top of the page, click **Save**.

Step 5 Repeat the preceding steps to associate your captive portal SSL policy with the access control policy.

Step 6 If you haven't done so already, target the policy at managed devices as discussed in Setting Target Devices for an Access Control Policy.

What to do next

- Deploy your identity and access control policies to managed devices as discussed in Deploy Configuration Changes.
- Monitor user activity as discussed in Using Workflows.

Captive Portal Fields

Use the following fields to configure captive portal on the **Active Authentication** tab page of your identity policy. See also Identity Rule Fields and Exclude Applications from Captive Portal, on page 21.

Server Certificate

The server certificate presented by the captive portal daemon.

Note Captive portal does *not* support the use of Digital Signature Algorithm (DSA) or Elliptic Curve Digital Signature Algorithm (ECDSA) certificates.

Port

The port number to use for the captive portal connection. If you plan to use an ASA FirePOWER device for captive portal, the port number in this field must match the port number you configured on the ASA FirePOWER device using the **captive-portal** CLI command.

Maximum login attempts

The maximum allowed number of failed login attempts before the system denies a user's login request.

Active Authentication Response Page

The system-provided or custom HTTP response page you want to display to captive portal users. After you select an **Active Authentication Response Page** in your identity policy active authentication settings, you also must configure one or more identity rules with **`HTTP Response Page`** as the **Authentication Type**.

The system-provided HTTP response page includes **Username** and **Password** fields, as well as a **Login as guest** button to allow users to access the network as guests. To display a single login method, configure a custom HTTP response page.

Choose the following options:

- To use a generic response, click **System-provided**. You can click **View** ( ) to view the HTML code for this page.
- To create a custom response, click **Custom**. A window with system-provided code is displayed that you can replace or modify. When you are done, save your changes. You can edit a custom page by clicking **Edit** ( ).

Related Topics

Internal Certificate Objects

Exclude Applications from Captive Portal

You can select applications (identified by their HTTP `User-Agent` strings) and exempt them from captive portal active authentication. This allows traffic from the selected applications to pass through the identity policy without authenticating.

Note Only applications with the **`User-Agent Exclusion` Tag** are displayed in this list.

Procedure

Step 1 If you haven't done so already, log in to the Firepower Management Center.

Step 2 Click **Policies** > **Access Control** > **Identity** .

Step 3 Edit the identity policy that contains the captive portal rule.

Step 4 On **Realm & Settings** tab page, use the filters in the **Application Filters** list to narrow the applications you want to add to the filter.

- Click the arrow next to each filter type to expand and collapse the list.
- Right-click a filter type and click **Check All** or **Uncheck All**. Note that the list indicates how many filters you have selected of each type.
- To narrow the filters that are displayed, type a search string in the **Search by name** field; this is especially useful for categories and tags. To clear the search, click **Clear** ( ).
- To refresh the filters list and clear any selected filters, click **Reload** ( ).
- To clear all filters and search fields, click **Clear All Filters**.

Note The list displays 100 applications at a time.

Step 5 Choose the applications that you want to add to the filter from the **Available Applications** list:

- To narrow the individual applications that appear, enter a search string in the **Search by name** field. To clear the search, click **Clear** ( ).
- Use paging at the bottom of the list to browse the list of individual available applications.
- To refresh the applications list and clear any selected applications, click **Reload** ( ).

Step 6 Add the selected applications to exclude from external authentication. You can click and drag, or you can click **Add to Rule**. The result is the combination of the application filters you selected.

What to do next

- Continue configuring the identity rule as described in Create an Identity Rule.

Troubleshoot the Captive Portal Identity Source

For other related troubleshooting information, see Troubleshoot Realms and User Downloads and Troubleshoot User Control.

If you experience issues with captive portal, check the following:

- The time on your captive portal server must be synchronized with the time on the Firepower Management Center.
- If you have DNS resolution configured and you create an identity rule to perform **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) captive portal, you must configure your DNS server to resolve the fully qualified domain name (FQDN) of the captive portal device. The FQDN must match the hostname you provided when configuring DNS.

  For ASA with FirePOWER Services and Firepower Threat Defense devices, the FQDN must resolve to the IP address of the routed interface used for captive portal.
- If you're using Kerberos authentication, the managed device's host name must be less than 15 characters (it's a NetBIOS limitation set by Windows); otherwise, captive portal authentication fails. You set the managed device host name when you set up the device. For more information, see an article like this one on the Microsoft documentation site: Naming conventions in Active Directory for computers, domains, sites, and OUs.
- DNS must return a response of 512 bytes or less to the hostname; otherwise, testing the connection the AD connection fails. This limit applies in both directions and is discussed in RFC 6891 section-6.2.5.
- If you select **`Kerberos`** (or **`HTTP Negotiate`**, if you want Kerberos as an option) as the **Authentication Type** in an identity rule, the **Realm** you select must be configured with an **AD Join Username** and **AD Join Password** to perform Kerberos captive portal active authentication.
- If you select **`HTTP Basic`** as the **Authentication Type** in an identity rule, users on your network might not notice their sessions time out. 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.
- If the connection between your Firepower Management Center and a managed device fails, no captive portal logins reported by the device can be identified during the downtime, unless the users were previously seen and downloaded to the Firepower Management Center. The unidentified users are logged as Unknown users on the Firepower Management Center. After the downtime, the Unknown users are reidentified and processed according to the rules in your identity policy.
- If the device you want to use for captive portal contains both inline and routed interfaces, you must configure a zone condition in your captive portal identity rules to target only the routed interfaces on the captive portal device.
- The system does not validate the type of interface in ASA with FirePOWER devices. If you apply a captive portal policy to an inline (tap mode) interface on an ASA with FirePOWER device, the policy deployment succeeds but users in traffic matching those rules are identified as Unknown.
- The host name of the managed device must be less than 15 characters for Kerberos authentication to succeed.
- The only way to be sure a user logs out is to close and reopen the browser. Unless that happens, in some cases, the user can log out of captive portal and be able to access the network without authenticating again using the same browser.

- Active FTP sessions are displayed as the **Unknown** user in events. This is normal because, in active FTP, the server (not the client) initiates the connection and the FTP server should not have an associated user name. For more information about active FTP, see RFC 959.
- Captive portal authenticates any user in the associated realm, even if that user does not belong to a downloaded group. The system identifies users in non-downloaded groups as Unknown; Unknown users match no identity rules. To avoid that, configure the realm to download users in all groups you expect to authenticate with captive portal.

  To make sure the system downloads all users in a realm, make sure the groups are in the Available Groups list in the realm's configuration.

  For more information about users and groups, see Download Users and Groups.

The Traffic-Based Detection Identity Source

Traffic-based detection is the only non-authoritative identity source supported by the Firepower System. When configured, managed devices detect LDAP, AIM, POP3, IMAP, Oracle, SIP (VoIP), FTP, HTTP, MDNS, and SMTP logins on the networks you specify. The data gained from traffic-based detection can be used only for user awareness. Unlike authoritative identity sources, you configure traffic-based detection in your network discovery policy as described in Configuring Traffic-Based User Detection.

Note the following limitations:

- Traffic-based detection interprets only Kerberos logins for LDAP connections as LDAP authentications. Managed devices cannot detect encrypted LDAP authentications using protocols such as SSL or TLS.
- Traffic-based detection detects AIM logins using the OSCAR protocol only. They cannot detect AIM logins using TOC2.
- Traffic-based detection cannot restrict SMTP logging. This is because users are not added to the database based on SMTP logins; although the system detects SMTP logins, the logins are not recorded unless there is already a user with a matching email address in the database.

Traffic-based detection also records failed login attempts. A failed login attempt does not add a new user to the list of users in the database. The user activity type for detected failed login activity detected by traffic-based detection is **Failed User Login**.

Note The system cannot distinguish between failed and successful HTTP logins. To see HTTP user information, you must enable **Capture Failed Login Attempts** in the traffic-based detection configuration.

Caution Enabling or disabling non-authoritative, traffic-based user detection over the HTTP, FTP, or MDNS protocols, using the network discovery 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 target device handles traffic. See Snort® Restart Traffic Behavior for more information.

Traffic-Based Detection Data

When a device detects a login using traffic-based detection, it sends the following information to the Firepower Management Center to be logged as user activity:

- the user name identified in the login
- the time of the login
- the IP address involved in the login, which can be the IP address of the user's host (for LDAP, POP3, IMAP, and AIM logins), the server (for HTTP, MDNS, FTP, SMTP and Oracle logins), or the session originator (for SIP logins)
- the user's email address (for POP3, IMAP, and SMTP logins)
- the name of the device that detected the login

If the user was previously detected, the Firepower Management Center updates that user's login history. Note that the Firepower Management Center can use the email addresses in POP3 and IMAP logins to correlate with LDAP users. This means that, for example, if the Firepower Management Center detects a new IMAP login, and the email address in the IMAP login matches that for an existing LDAP user, the IMAP login does not create a new user; rather, it updates the LDAP user's history.

If the user was previously undetected, the Firepower Management Center adds the user to the users database. Unique AIM, SIP, and Oracle logins always create new user records, because there is no data in those login events that the Firepower Management Center can correlate with other login types.

The Firepower Management Center does **not** log user activity or user identities in the following cases:

- if you configured the network discovery policy to ignore that login type
- if a managed device detects an SMTP login, but the users database does not contain a previously detected LDAP, POP3, or IMAP user with a matching email address

The user data is added to the users table.

Traffic-Based Detection Strategies

You can restrict the protocols where user activity is discovered to reduce the total number of detected users so you can focus on users likely to provide the most complete user information. Limiting protocol detection helps minimize user name clutter and preserve storage space on your Firepower Management Center.

Consider the following when selecting traffic-based detection protocols:

- Obtaining user names through protocols such as AIM, POP3, and IMAP may introduce user names not relevant to your organization due to network access from contractors, visitors, and other guests.
- AIM, Oracle, and SIP logins may create extraneous user records. This occurs because these login types are not associated with any of the user metadata that the system obtains from an LDAP server, nor are they associated with any of the information contained in the other types of login that your managed devices detect. Therefore, the Firepower Management Center cannot correlate these users with other types of users.

Related Topics

Configuring Traffic-Based User Detection