Settings

System-level settings available to you vary depending on your role. For example, only users with **Site Admin** and **Customer Support user** role, can view the **Users** option.

Threat Intelligence

To manage threat intelligence, in the left navigation pane, click **Manage** > **Service Settings** > **Threat Intelligence**.

The **Threat Intelligence** feature provides the most up-to-date datasets for the Secure Workload pipeline that identifies and quarantines threats by inspecting the datacenter workloads against externally known malware command and control addresses, security flaws in processes and geographical location.

The Threat Intelligence dashboard displays the updated status of threat intelligence datasets. These datasets are updated automatically.

Warning The Threat Intelligence feature requires a connection to Cisco Secure Workload servers to automatically update. Your enterprise outbound HTTP request may require:

- Allow the following domain from the enterprise firewall outbound rules: `uas.tetrationcloud.com`

- Configure your Outbound HTTP Connection.

In environments without an outbound connection, upload the datasets directly. See the **Manual Uploads** section.

Table 1: Datasets

Dataset	Description
NVD CVEs	Security related software flaws, CVSS base score, vulnerable product configuration, and weakness categorization
MaxMind Geo	Identification of the location and other characteristics of source IPs
NIST RDS	NIST Reference Data Set of digital signatures of known, traceable software applications
Team Cymru	Insight on 3,000+ botnet command and control IPs
Hash Verdict	Verdict of Secure Workload on process hashes (only available with the Automatic Updates section).

Note In case the MaxMind Geo dataset is manually uploaded in an earlier release, you must reupload the corresponding RPM to view the location and related information on the Flow Visibility page.

Automatic Updates

The threat dataset updates are triggered from the appliance to synchronize with the global dataset that is hosted on the Internet at uas.tetrationcloud.com, everyday between 3-4 a.m. UTC. The global dataset is refreshed weekly on Fridays or Mondays. The Threat Intelligence dashboard lists the datasets and the date on which the dataset is last updated.

Figure 1: Dashboard

Automatic Updates

Status

Tetration Cloud Connection

Automatic updates are not active. An Outbound HTTP Proxy may need to be configured.

Threat Datasets   Auto Refresh

Name ↑	Version ↑↓	File Name ↑↓	Status ↑↓	Start Date ↑↓	Install Date ↑↓	Source ↑↓	History
CVE Data	201807161119	tetration_os_supplemental_data_pack_cve_k9-201807161119-1.noarch.rpm	Installed	Aug 10 4:00:12pm	Jul 3 12:45:00pm		
MaxMind Geo	201804070620	tetration_os_supplemental_data_pack_geo_k9-201807161119-1.noarch.rpm	Installed	Aug 10 4:00:12pm	Jul 3 12:45:00pm		
NIST RDS	201809200819	tetration_os_supplemental_data_pack_rds_k9-201807161119-1.noarch.rpm	Installed	Aug 10 4:00:12pm	Jul 3 12:45:00pm		

Upload Threat Dataset

Select Supplemental RPM

Threat Datasets Supplemental RPMs can be downloaded from Cisco Tetration Update Portal.
Learn More

Manual Uploads

Attention	**Scheduling Manual Uploads**—Dataset RPM files are published to Secure Workload Update Portal weekly. It is recommended to install the latest releases periodically by configuring a schedule for an administrator.

Downloading Updated Datasets

The datasets can be downloaded from Secure Workload Update Portal.

Uploading Datasets Manually

To upload dataset RPM files:

Before you begin

Log in as a **Site Administrator** or **Customer Support**.

Step 1 In the left navigation pane, click **Manage** > **Service Settings** > **Threat Intelligence**.

Step 2 Under the **Upload Threat Dataset** section, click **Select Supplemental RPM**.

Step 3 Upload the RPM file downloaded from Secure Workload Update Portal.

Step 4 Click **Upload**.

The RPM upload process is initiated and the status is displayed on a progress bar. After the upload, the RPM file is processed and installed in the background. The table is updated after the installation is complete.

Figure 2: Threat Datasets

Name ↑	Version ↑↓	File Name ↑↓	Status ↑↓	Start Date ↑↓	Install Date ↑↓	Source ↑↓	History
MaxMind Geo	202108060000	tetration_os_supplemental_data_pack_geo_k9-202108060000-1.noarch.rpm	Failed	Aug 10 5:22:47pm		⬆	☰
Team Cymru	202108060000	tetration_os_supplemental_data_pack_zeus_k9-202108060000-1.noarch.rpm	Failed	Aug 10 5:23:12pm		⬆	☰

Threat Datasets — Auto Refresh 🔘

Change Log

Site Admins can access the **Change Log** page under the **Manage** menu in the navigation bar at the left side of the window. This page displays the most recent changes that are made within Cisco Secure Workload.

Figure 3: Change Log Page

The details of each change log entry can be viewed by clicking on the link in the **Change At** column. This page includes a **Before** and **After** snapshot of the fields changed. The fields may include technical names that require some interpretation to understand how they are surfaced elsewhere throughout Secure Workload.

Figure 4: Change Log Details Page

The complete list of changes for an entity can be viewed by clicking the button in the upper-right corner, titled **Full log for this <entity type>**. This page displays the details of each change. It also includes the **Current State** of the entity, when available.

Figure 5: Full Change Log for Entity

Collection Rules

Site Admins and **Customer Support users** can access the **Collection Rules** page under the **Manage > Service Settings** menu in the navigation bar at the left side of the window. This page displays the hardware collection rules by VRF that is used by switches running the Cisco Secure Workload agent. There is a row in the table for each VRF.

Rules

Click the **Edit** button on a VRF to modify its collection rules. By default, every VRF is configured with two default catch-all rules, one for IPv4 (`0.0.0.0/0 INCLUDE`) and one for IPv6 (`::/0 INCLUDE`). *These default rules can be removed, but do so with caution.*

Extra include and exclude rules can be added. Enter a valid subnet, select include or exclude, and click **Add Rule**. The priority of these rules can be adjusted via drag-and-drop. Click-and-hold on a rule in the list and drag it to adjust the order.

Changes may take several minutes to propagate to your switches. Click the **Back** button in the upper-right corner to return to the VRF list.

Priority

Collection Rules are ordered in decreasing order priority. No longest prefix match is done to determine the priority. The rule appearing first has higher priority over all the subsequent rules. Example:

1. 1.1.0.0/16 INCLUDE

2. 1.0.0.0/8 EXCLUDE

3. 0.0.0.0/0 INCLUDE

In the earlier example, all addresses belonging to 1.0.0.0/8 subnet are excluded except subnet 1.1.0.0/16 which is included.

Another Example with changed order:

1. 1.0.0.0/8 EXCLUDE

2. 1.1.0.0/16 INCLUDE

3. 0.0.0.0/0 INCLUDE

In the above example, all addresses belonging to 1.0.0.0/8 subnet are excluded. Rule number-2 does not get exercised here because of a higher-order rule already defined for its subnet.

Collectors

Site Admins and **Customer Support users** can access the **Collectors** page under the **Platform** menu in the navigation bar at the left side of the window. This page displays the currently configured collectors. The Secure Workload agents send flow data to the commissioned collectors, so it is important for all of the commissioned collectors to be available. By default, all collectors are periodically checked for their health and they are either commissioned or decommissioned based on their health. You can opt out of this automated process using the toggle **Auto Commission Opt Out**. With this toggle on, the **Play** and **Stop** icons under the far right column can be used to commission and decommission respectively.

Figure 6: Collectors Page

Name ↑↓	IP ↑↓	TCP Port ↑↓	UDP Port ↑↓	Health ↑↓	Health Details ↑↓	Status ↑↓	Auto Commission Opt Out	Manual Action
collectorDatamover-1	172.21.156.182	5640	5640	Healthy		Commissioned		
collectorDatamover-2	172.21.156.183	5640	5640	Healthy		Commissioned		

Session Configuration

UI User Authentication idle session timeout can be configured here. This config applies to all the users of the appliance. The default idle session duration is 1 hour. The idle session duration can be set within the range of 5 minutes to 24 hours. The session timeout takes effect on a user's authenticated session when this value is saved.

Site Admins and **Customer Support users** can access this setting. In the left navigation pane, click **Manage** > **Service Settings** > **Session Configuration**.

Company

You can set the following company-wide (per Secure Workload cluster) configurations.

Outbound HTTP Connection

To ensure the latest Threat Intelligence Datasets are retrieved from Cisco Cloud, we highly recommend that you set up an outbound HTTP connection.

⚠️

Warning Your enterprise outbound HTTP request may require allowing traffic to **periscope.tetrationcloud.com** and **uas.tetrationcloud.com** from enterprise firewall outbound rules in addition to setting up the HTTP Proxy as shown below.

The TLS connection to **periscope.tetrationcloud.com** is used to transport Threat Intelligence Data for identifying known vulnerabilities. Therefore, it is essential for Cisco Secure Workload to verify the authenticity of the domain name by verifying the domain's X.509 certificate's signing CA cert against reputable root CA certificates included with Secure Workload. Tampering with the X.509 trust chain prevents the feature from working correctly.

Figure 7: Outbound HTTP Connection

Site Admins and **Customer Support users** can access Outbound HTTP settings. In the navigation bar on the left, click **Platform** > **Outbound HTTP**.

Field	Description
Status	Indicates whether Secure Workload appliance can reach to Secure Workload Cloud to retrieve Threat Intelligence Dataset updates. The status check can be retriggered by clicking on the refresh button. The following HTTP proxy settings can be used to configure HTTP Proxy settings based on your Secure Workload deployment.
Enable HTTP Proxy	All external HTTP connections use HTTP proxy if this option is enabled
Host	HTTP proxy host address
Port	HTTP proxy port number
Username	Required only if your HTTP proxy server uses basic authentication

Field	Description
password	Required only if your HTTP proxy server uses basic authentication

Login Page Message

Site Admins and **Customer Support users** can enter a message of up to 1600 characters that users see on the sign-in page.

To create or change the login page message:

1. In the left navigation page, click **Platform** > **Login Page Message**.

2. Enter or edit the message. The character limit is less than or equal to 1600 characters.

3. Click **Save**.

Configure External Authentication

If this option is enabled, authentication can be handed off to an external system. The current options for authentication are Lightweight Directory Access Protocol (LDAP) and Single Sign-On (SSO). This means that once this is enabled all users signing in will use the chosen mechanism to authenticate. It is important to establish that the LDAP connection is configured correctly, especially if no users are on the 'Use Local Authentication' Option. The recommended approach is to have at least one locally authenticated user with **Site Admin** credentials by turning on the 'Use Local Authentication' Option. This user can make sure that the LDAP configuration is set up correctly. Once the connection is successfully set up, this user can also be transitioned to external authentication by unchecking the 'Use Local Authentication' option in the user edit flow.

Site Admin can enable more debug messages which are useful to debug external connection issues, user sign-in failures and so on. This can be enabled by checking the 'External Auth Debug' option. Once this is turned on, more descriptive log messages are written into a separate log file titled 'external_auth_debug.log'. The recommendation is to turn off 'External Auth Debug' once debugging is done to prevent extra logs being written into the log file.

Note Users can bypass external authentication once it is enabled on a per user basis as indicated in 'Use Local Authentication' Option. This option can also be enabled by going to the user edit flow from link through the warning message when external auth is enabled as well.

External Authentication with SSO is the recommended authentication approach if Federation is enabled.

Note Starting from 3.7.1.5 release and later, external authentication session for eviction time is increased from six hours to nine hours. This setting is applicable for external authentication or on-premises only.

Site Admins and **Customer Support users** can configure external authentication. In the navigation bar on the left, click **Platform** > **External Authentication**.
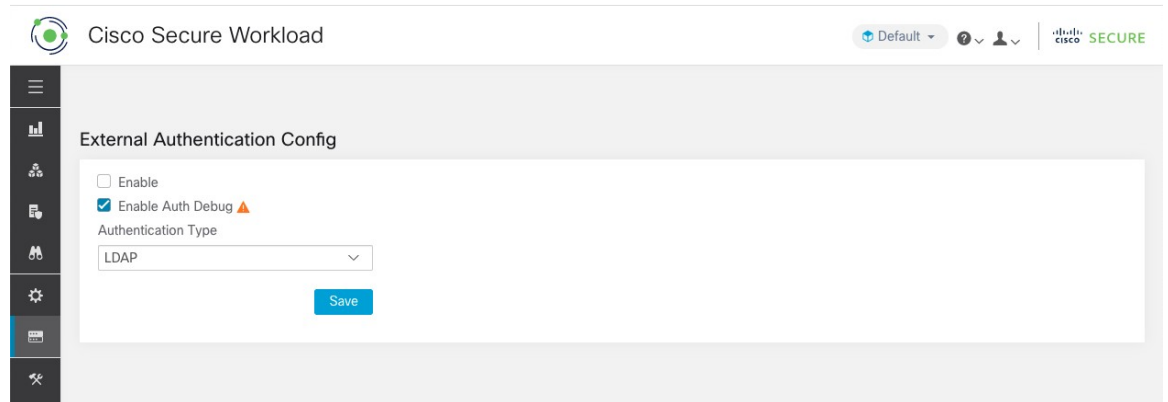
Figure 8: Configuring External Authentication
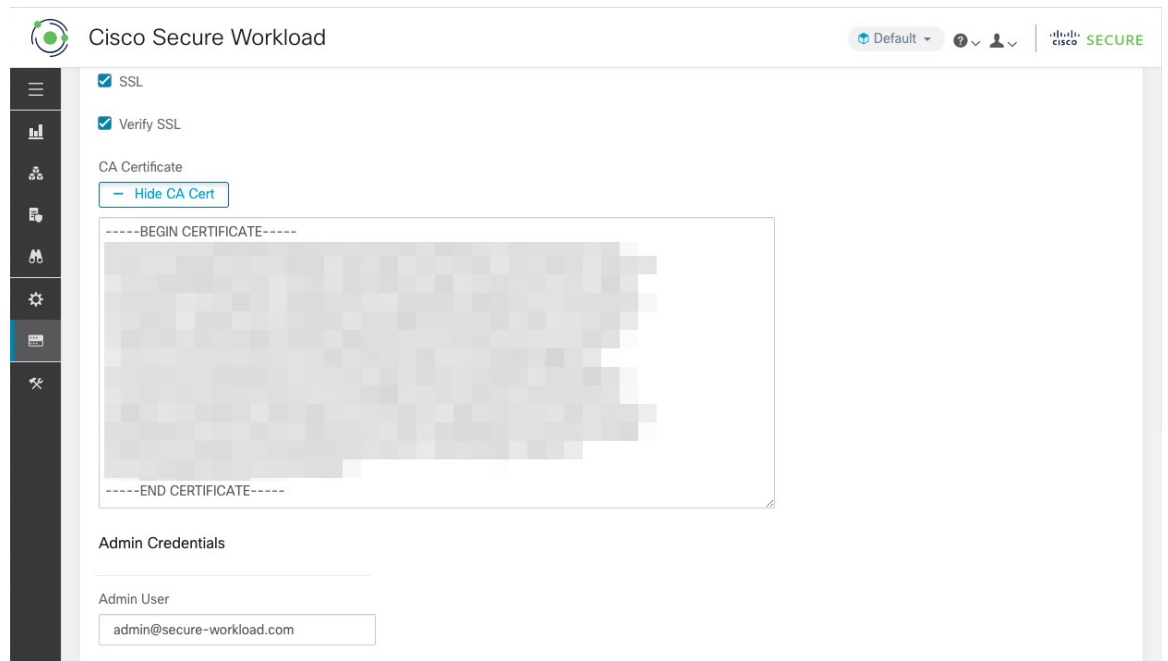

Figure 9: Configuring External Authentication Continued

Figure 10: Configuring External Authentication Continued

Figure 11: External Authentication Warning

Configure Lightweight Directory Access Protocol

Choose the Lightweight Directory Access Protocol (LDAP) option to authenticate users. This means that once this is enabled all users will be logged out and subsequent signing in will use their LDAP email and password to authenticate.

LDAP is currently not recommended as the authentication mechanism if 'Federation' is enabled.

If LDAP is enabled the recommended workflow for new user creation is as follows.

Site Admins are encouraged to first create new users with their emails and assign the appropriate roles by Configure LDAP Authorization (AD Authorization) before new users logs in via LDAP for the first time. If a new user logs in via LDAP without the appropriate role, no default role is assigned to the user.

Figure 12: Configuring Lightweight Directory Access Protocol

| Field	Description
Auto Create Users	Turning on 'Auto Create Users' will create users if they don't exist at first login. This saves the site admins from having to preprovision users before allowing users to log in. This option should be turned off if Secure Workload access should be limited to users manually created on the Users page.
Host	LDAP Host which will be used for authentication.
Port	LDAP Port which will be used for authentication.
Email Attribute	LDAP attribute name which represents email for the organization.
Base	LDAP base dn from where users will be searched.
SSL	Enable encryption and use 'ldaps://'.
SSL Verify	Verify server's SSL attributes such as Fully Qualified Domain Name (FQDN) based on server's certificate.
SSL Certificate Authority Cert	Signing cert for LDAP server's SSL Cert. Required if server cert chain cannot be publicly verified.
Admin User	LDAP Admin user (not Secure Workload user) name used to bind against the LDAP server. For example: [User]@[Domain] or [Domain]\\[User]

Field	Description
Admin Password	LDAP Admin password that is used to bind against the LDAP server.
Ldap Authorization	LDAP Authorization can be enabled and configured as explained in Configure LDAP Authorization (AD Authorization).

Once the LDAP config is enabled all users except users with 'Use Local Authentication' Option enabled will be logged out of their sessions.

The LDAP config can be saved once the **'Save'** button is clicked. We recommend that you wait for a minute after the LDAP config is saved successfully before attempting to test the LDAP connection.

The LDAP connection can be tested out after the LDAP config has been saved using the **'Test Connection'** button. This tries a bind against the LDAP server with the admin credentials entered.

Figure 13: Authentication Workflow
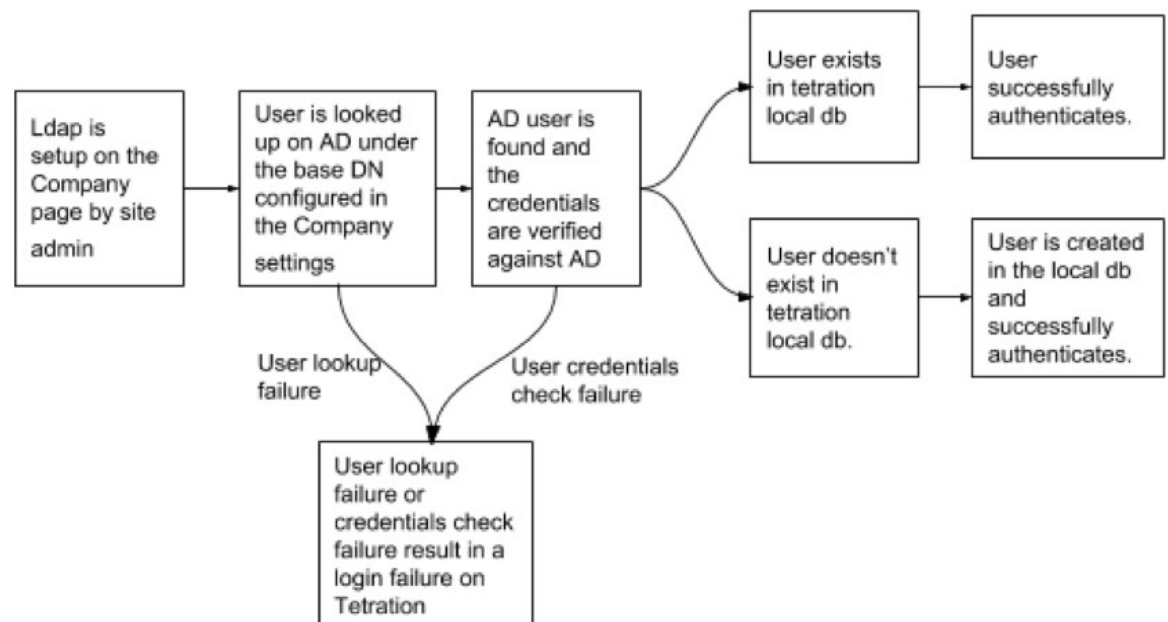
Troubleshoot LDAP Issues

If an error is raised when you test the ldap connection, check the following:

- Check whether the LDAP admin credentials are correct.

- Check the connection params such as host, port, ssl and so on.

- Check whether the LDAP server can be reached from Secure Workload UI VIPs.

- Check whether the AD server is up.

- Use command-line tools such as **'ldapsearch'** with the connection details to see whether a bind can be made.

If an error is raised during login for a user, check the following:

- Check whether the user can log in with their LDAP credentials to other company websites which use LDAP authentication.

- Check whether the 'base' dn that is specified in the Company LDAP settings is correct. This can be done by using command-line tools such as **'ldapsearch'** to look up the user against the base dn.

Example **'ldapsearch'** query to search a user by email:

```
ldapsearch -H "ldap://<host>:<port>" -b "<base-dn>" -D "<ldap-admin-user>" -w
<ldap->admin-password> "(mail=<users-email-address>)"
```
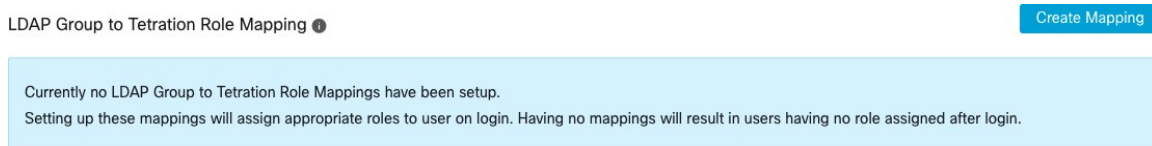
Configure LDAP Authorization (AD Authorization)

Active Directory Authorization can be configured by enabling the 'LDAP Authorization' checkbox in the 'Admin Credentials' section of the External Authentication LDAP configuration. Once this setting is enabled, Site Admin must set up mappings of LDAP 'MemberOf' groups to Secure Workload Roles in the section below. By default, without this configuration, Active Directory users must be preconfigured with one or more Secure Workload roles prior to a login attempt.

LDAP MemberOf Group to Secure Workload Role Mapping must be set up if LDAP external authentication is enabled. 'Create Mapping' allows setting up an LDAP MemberOf group value to be mapped to a Secure Workload Role. The roles in the role dropdown are prepopulated based on the scope that is selected in the scope selector. Once these mappings are saved, all users get authorized based on these values on their subsequent login.

These mappings can be reordered, edited, or deleted. Any modifications to the mappings will be reflected on the roles assigned to users on their subsequent login. A maximum of 50 LDAP MemberOf Group to Secure Workload Role Mappings can be created.

Duplicate LDAP MemberOf group names are not allowed. However multiple LDAP MemberOf groups can map to the same role. If more than one group maps to the same role, the last mapping will be stored in the user as the matched LDAP MemberOf to Secure Workload role.

Figure 14: LDAP Group to Secure Workload Role SetUp

LDAP Group to Tetration Role Mapping ⓘ   Create Mapping

Currently no LDAP Group to Tetration Role Mappings have been setup.
Setting up these mappings will assign appropriate roles to user on login. Having no mappings will result in users having no role assigned after login.

Figure 15: LDAP Group to Secure Workload Role Mapping

LDAP Group to Tetration Role Mapping ⓘ   Create Mapping

☰   Apply member group ░░░░░░░░░░░░░░░░░░░░   to Tetration role **Site Admin**   Edit   Delete

☰   Apply member group ░░░░░░░░░░░░░░░░░░░░   to Tetration role **Global Application Enforcement**   Edit   Delete

A site admin user can reconcile the assignment of roles based on the above role mapping with the help of external user's information that is obtained from the user's last successful login.

Note   Users can bypass external authentication once it is enabled on a per user basis as indicated in 'Use Local Authentication' Option. These users will also bypass the authorization process set up for AD authorization.

Figure 16: External User Information
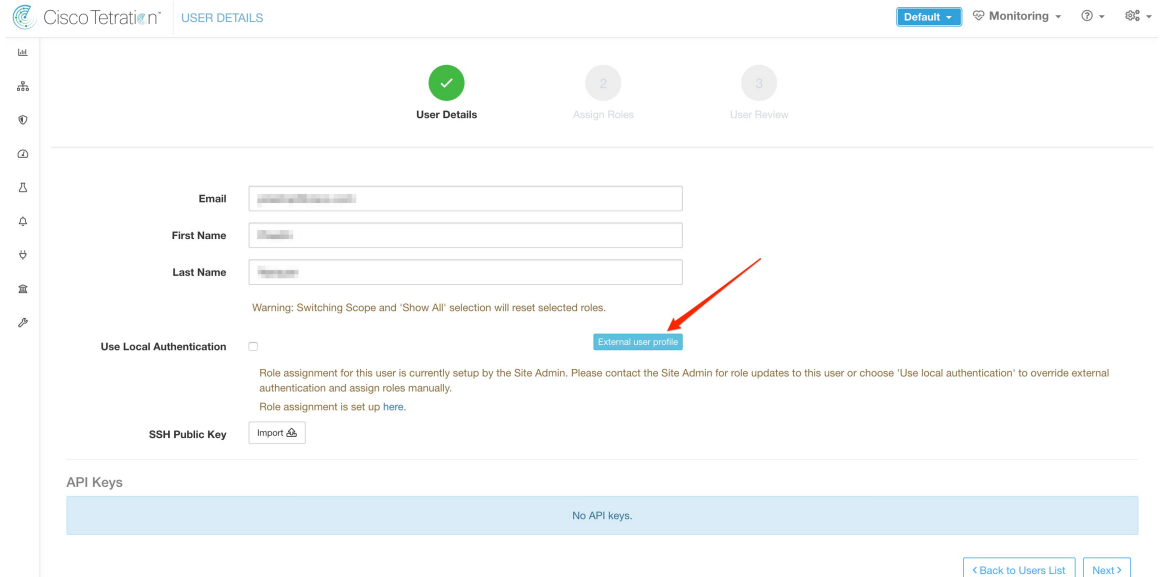
Once authorization is enabled, manual Secure Workload Role selection in the user creation (Adding a New User Account, on page 54) and user edit flows (Editing a User Account) is **disallowed**.

Figure 17: Users Page

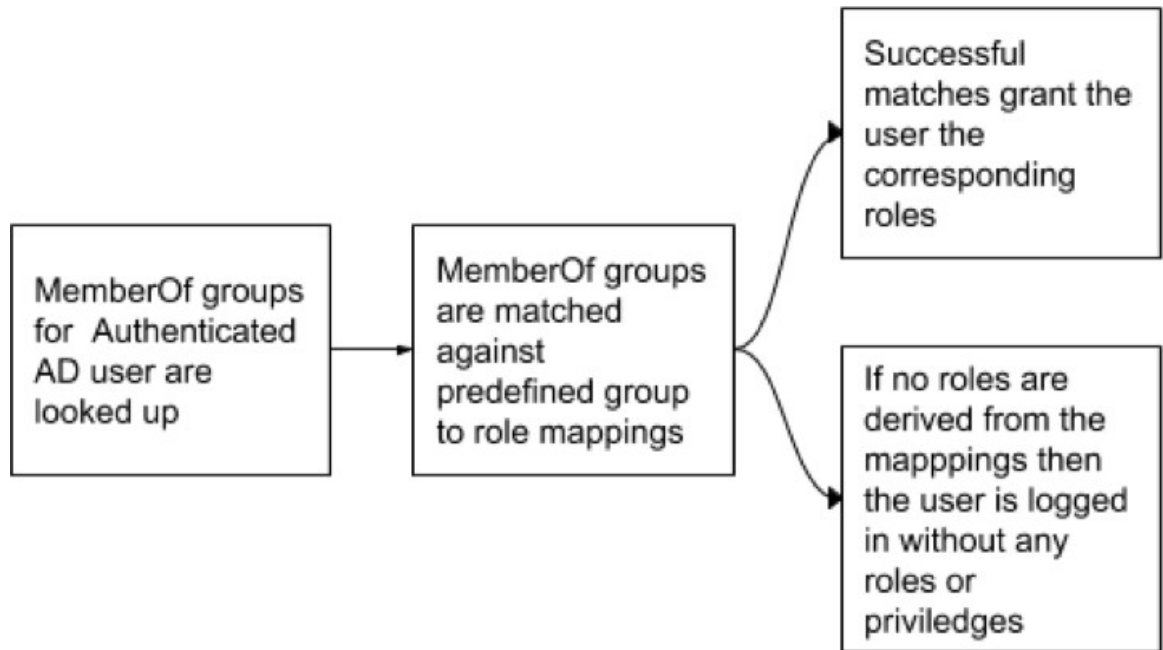The mapped LDAP MemberOf groups to Secure Workload Roles are visible on the user profile page.

Figure 18: User Profile Page

Figure 19: Authorization Workflow

If LDAP Authorization is enabled, access to OpenAPI via API Keys cease to work seamlessly because Secure Workload roles that are derived from LDAP MemberOf groups are reassessed once the user session terminates. Hence to ensure uninterrupted OpenAPI access, we recommend that users with API Keys have 'Use Local Authentication' Option enabled.

Figure 20: LDAP Authorization API Key Warning

API Keys

Ensure that you have 'Use Local Authentication' enabled for the user to allow seamless API access using API keys when LDAP authorization is enabled.

API Key	Capabilities	Description ↑↓	Created At ↑	Last Used ↑↓	
8aac707bc10743d0995b725ceb37ce4e	• sensor_management • software_download • flow_inventory_query		Aug 11 02:38:07 pm (EEST)		🗑

Figure 21: LDAP Authorization API Key Warning on Users Page

Troubleshoot LDAP Authorization Issues

If the roles are not getting assigned to users based on the mappings defined in the 'External Authentication', 'LDAP Group to Role Mappings' section, check the role mappings setup and format once more.

- Group string must be of the string format. For example: CN=group.jacpang,OU=Organizational,OU=Cisco Groups,DC=stage,DC=cisco,DC=com

- Group names must be exact from what is present in AD with no spaces or extra characters.

- Role mapping for the group must be selected from the role selector.

User Role Mapping Debug Steps

- You must have two users, one that is Site Admin, the email of this user should not be the same as the AD user.

- This user is called 'SA User' for the steps below.

- SA user has previously set up the role-mapping configs on the Company page External Auth Config as described earlier. Let's assume 'SA User' will be logging in with [site-admin]@[Domain].

  - We assume that 'AD User' is [ad-user]@[Domain]. We assume that the LDAP setup is done and the AD user is able to log in but not getting his role that is assigned.

- As AD User, log in using incognito browser session. This splits the browser state from SA User session.

- As SA User, login and go to Users page.

- Click on the Edit Icon for the AD User that must have Role Mapping configured.

- Click the 'External User Profile' button on the User Profile page.

- You will see an External Auth Profile Table that includes a 'memberof' section.

- This is one of the 'memberof' values that you can use for role mapping on Company page, External Auth Config, LDAP Group to Role-Mapping section.

- You must provide the whole 'memberof' per-line string to match. Once you create this role mapping, anyone who has the same attribute 'memberof' will be assigned the mapped role.

- For the AD User to be granted the newly mapped role, the user needs to log out then log back in to allow re-evaluation of this mapping profile.

- Once a user logs in and has roles that are assigned successfully as a result of group role mappings, the matching rules are visible on the 'Preferences' page for that user.

Configure Single Sign-On (SSO)

If this option is selected, SSO can be used to authenticate users. This means that once this is enabled all users will be redirected to the identity provider sign-in page to authenticate. Users with 'Use Local Authentication' Option enabled can use the email and password sign-in form in the sign-in page to authenticate.

It is important to establish that the SSO configuration is set up correctly, especially if no users are on the 'Use Local Authentication' Option. The recommended approach is to have at least one locally authenticated user with **Site Admin** credentials by turning on the 'Use Local Authentication' Option. This user can make sure that the SSO configuration is set up correctly. Once the connection is successfully set up, this user can also be transitioned to external authentication by unchecking the 'Use Local Authentication' option in the user edit flow.

If SSO is enabled the recommended workflow for new user creation is as follows.

Site Admins and **Scope Owners** are encouraged to first create new users with their emails and assign the appropriate roles and scopes before the new user logs in via SSO for the first time. If a new user logs in via SSO without the appropriate role, no default role is assigned to the user.

The following table describes the fields that must be set up to configure SSO on Secure Workload. Secure Workload is the Service Provider (SP) in this case.

Figure 22: Configuring Single Sign-On

Field	Description
SSO Target Url	SSO IdP target url to which users will be redirected to for sign in.
SSO Is-suer	SSO Entity Id of your SP, a URL that uniquely identifies your SP. This is generally the metadata for the SP. In this case it is: `https://<tetration-cluster-fqdn>/h4_users/saml/metadata`
SSO Certifi-cate	SSO certificate that is provided by the Identity Provider (IdP).
SSO AuthN Context	Choice for SSO AuthN Context which is specified in the SAML Request. The default option is 'Password Protected Transport'. The other choices are 'Integrated Windows Authentication' and 'X.509 Certificate' for Windows and PIV-based authentication.

After the SSO configuration is enabled, all users, except the users who have enabled the Use Local Authentication option, are logged out of their sessions.

The SSO configuration is saved after the **Save** button is clicked.

Figure 23: Authentication Workflow
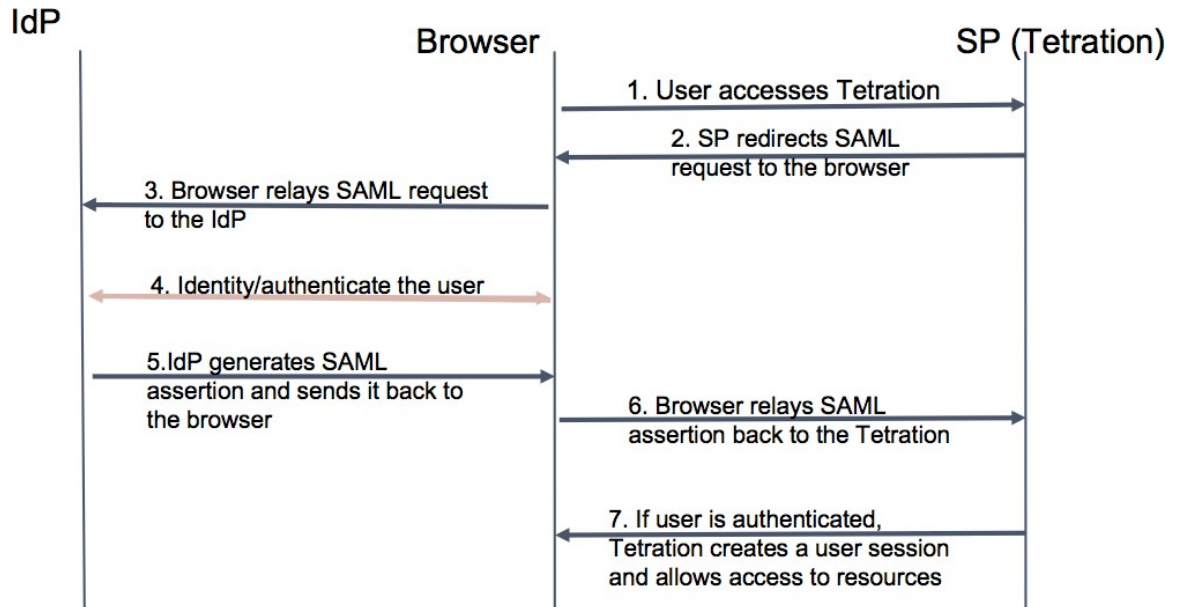
Information Shared to Identity Provider (IdP)

The IdP requires some information from Secure Workload (SP) to set up SSO for authentication. The following table describes the fields that must be set up.

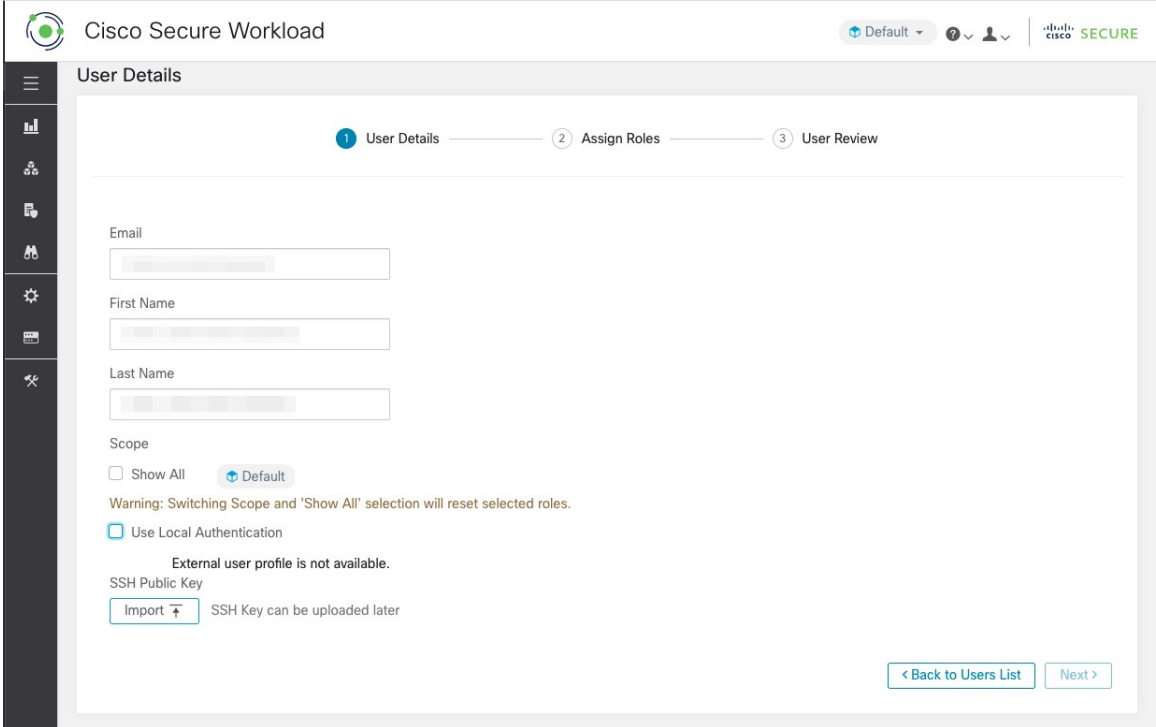| Field	Description
SSO Url	The authentication endpoint (url) which will consume the SAML assertion (response from the IdP). In our case, it will be: `https://<tetration-cluster-fqdn>/h4_users/saml/auth`
Entity Id	This is the metadata for the SP. In this case it is: `https://<tetration-cluster-fqdn>/ h4_users/saml/metadata`
Name ID format	NameId is email i.e `'urn:oasis:names:tc:SAML:1. 1:nameid-format:emailAddress'`
Attributes	User attributes are fetched from the IDP. We fetch these attributes as part of authentication: • email • firstName • lastName Ensure that the attribute names are as specified previously.

Troubleshoot SSO Issues

- Set up some downtime for this SSO config setup since the only way to verify authentication works (from the Service Provider) it is after setting it up.

- Check and validate the IdP metadata generated.

- Check all configuration parameters that are exchanged between IdP and SP.

  - Config at the IdP - SSO url, Audience, Name ID, attributes and so on

  - Config on Secure Workload Company page - SSO Target url, SSO issuer, and SSO certificate.

- Get a sample SAML assertion returned from the IdP from the server app logs. Validate it against a SAML validator to make sure it is a valid SAML response.

- Errors in the SP SSO setup may result in an error that is generated from the IdP. Using the browser inspect element, you can see the network requests being made.

- If a user has issues logging in, have the IdP admin check whether the user has access to the Secure Workload app.

'Use Local Authentication' Option

Once the config is set up, it is possible for site admins to allow users not to use external authentication. This can be done on a per user basis by enabling a flag 'Use Local Authentication' in the user edit section. Selecting this field for the user will log that user out of all sessions.

Figure 24: Use Local Authentication

⚠️

Warning   **Ensure that at least one user has local authentication access!**

If the 'Use Local Authentication' option is removed (i.e unchecked) for a user and this user happens to be the last user with the option, then no user has local authentication access to sign in to Secure Workload. This means that no user can sign in if there is any disruption with the external authentication system, such as config issues, connectivity issues, and so on. You see a warning if you try to delete the last locally authenticated user.

Users logging via external authentication has shorter sessions and will be prompted to log in when the session expires. Users logging via external authentication cannot reset their password on the site (they have to do it on their company website). However if the 'Use Local Authentication' flag is set for the user, password reset is possible.

SSL Certificate and Key

To enable fully verifiable HTTPS access to the Secure Workload UI, an SSL certificate specific to the UI's domain name and the RSA private key that matches the SSL certificate's public key can be uploaded into the cluster.

An SSL Certificate can be obtained in two ways depending on the format of the Fully Qualified Domain Name (FQDN) used to refer to the Secure Workload UI Virtual IP (VIP) address. If the Secure Workload FQDN is based on an enterprise domain name such as tetration.cisco.com, your enterprise Certificate Authority (CA) who owns the base domain issues you an SSL Certificate. Otherwise, you may use a reputable SSL Certificate vendor to issue you an SSL Certificate for your FQDN.

✎

Note   It is important to note that although the Secure Workload UI supports Server Name Indication (SNI), subject alternative names (SANs) specified in the certificate will not be matched. For instance, if the common name (CN) of the certificate is tetration.cisco.com and the certificate includes a SAN for tetration1.cisco.com, HTTPS requests sent with an SNI-compatible browser to the cluster with tetration1.cisco.com as the hostname will not be served with that certificate. HTTPS requests made to the cluster with a hostname other than the hostname specified in the CN will be served using the default, self signed certificate installed on the cluster. These requests result in browser warnings.

Site Admins and **Customer Support users** can work with SSL Certificates. In the navigation bar on the left, click **Platform** > **SSL Certificate**.

To import the certificate and key, click on the **Import New Certifcate and Key** button.

✎

Note   The first import of SSL certification and the private key should be performed through a trusted network connection to the cluster so that the private key cannot be intercepted by malicious parties who has access to the transport layer.

Enter the following information for your SSL certificate and key:

NAME can be any name for the certificate key pair. This name is for your benefit when looking at which SSL certificate is installed.

X509 Certificate field accepts SSL certificate string in Privacy Enhanced Mail (PEM) format. If your SSL certificate requires intermediary CA bundle, concatenate the CA bundle after your cert so that the SSL certificate for your Secure Workload FQDN is in the beginning of the certificate file.

It should have the following format:

```
-----BEGIN CERTIFICATE-----

< Certificate for Secure Workload FQDN >

-----END CERTIFICATE-----

-----BEGIN CERTIFICATE-----

< Intermediary CA 1 content >

-----END CERTIFICATE-----

-----BEGIN CERTIFICATE-----

< Intermediary CA 2 content >

-----END CERTIFICATE-----

-----BEGIN CERTIFICATE-----

< Root CA content >

-----END CERTIFICATE-----
```

RSA Private Key field should be RSA private key of the public key that is signed in the previous certificate. It should have the following format:

```
-----BEGIN RSA PRIVATE KEY-----

< private key data >

-----END RSA PRIVATE KEY-----
```

Note RSA Private Key is required to be unencrypted. It causes a "500 Internal Server Error" if RSA Private Key is encrypted.

Once the import button is pressed, we run verification steps to ensure that public key that is signed in the certificate and the private key are indeed RSA key pair. If the verification is successful, we display the SHA1 digest (SHA1 signature and creation time) of the certificate bundle.

Reload the browser to see that your SSL connection to the Secure Workload UI is now using the newly imported SSL certificate.

Cluster Configuration

This section displays the running configuration of the Secure Workload cluster about the customer network and administrative contacts. Editable values are indicated with a pencil icon.

Note a. Strong SSL Ciphers for Agent Connections: When this option is enabled, TLS-1.0 and TLS-1.1 protocols and the following ciphers will not be accepted by Secure Workload cluster during the SSL negotiations: DHE-RSA-AES128-GCM-SHA256:DHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES256-SHA:ECDHE-RSA-AES256-SHA:DHE-RSA-AES128-SHA256:DHE-RSA-AES128-SHA:DHE-RSA-AES256-SHA256:DHE-RSA-AES256-SHA:ECDHE-ECDSA-DES-CBC3-SHA:ECDHE-RSA-DES-CBC3-SHA:EDH-RSA-DES-CBC3-SHA:AES128-GCM-SHA256:AES256-GCM-SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:DES-CBC3-SHA

Following connections honor it and use strong ciphers during the TLS handshake:

1. All API and UI connections to Secure Workload.

2. All visibility and enforcement agent connections to Secure Workload.

Note older SSL libraries may not support this option.

Site Admins and **Customer Support users** can access this setting. In the navigation bar on the left, click **Platform** > **Cluster Configuration**.

After the configuration is edited, it takes some time for the new configuration to be applied throughout the cluster and it is indicated by highlighting the particular config.

External IPv6 Cluster Connectivity

Physical Cisco Secure Workload clusters can be configured to connect to both external IPv4 and IPv6 networks. IPv4 connectivity is required but IPv6 connectivity is optional. Once IPv6 connectivity has been configured, it cannot be disabled. Enabling IPv6 connectivity for external networking for the cluster can only be done during deploy or upgrade. See the Cisco Secure Workload Upgrade Guide for more information about enabling external IPv6 cluster connectivity during upgrade or the Cisco Secure Workload Hardware Deployment Guide for more information about enabling external IPv6 cluster connectivity during deployment.

Before you begin

To get agents to operate in dual stack mode (supporting both IPv4 and IPv6)

Prerequisite

- Cluster must have IPv6 enabled.

- Create A and AAAA records (for IPv4 and IPv6) in DNS for an FQDN and wait for the domain names to resolve.

Configure "Sensor VIP FQDN" for agents to operate in dual stack mode

Step 1 Choose **Platform** > **Cluster Configuration** from the navigation bar on the left.

Step 2 Look for the "Sensor IPv6 VIP", "Sensor VIP" and "Sensor VIP FQDN" fields. "Sensor IPv6 VIP" and "Sensor VIP" should already be set.

Step 3 If "Sensor VIP FQDN" is not set, set it to the FQDN created above. The A and AAAA records in DNS for the FQDN must resolve before you do this.

Step 4 If "Sensor VIP FQDN" was already set, make sure there are A and AAAA records in DNS for the FQDN as set in the "Sensor VIP FQDN" field, then click into the "Sensor VIP FQDN" field and save it to the same value so it updates.

Step 5 After the field completes updating (after about 20 minutes, the status is updated automatically), agents will be able to connect to the cluster via both IPv4 and IPv6.

Step 6 Valid "Sensor VIP FQDN" can be set only once.

Note	No IPv6 enforcement support for AIX. For more information on the requirements and limitations for dual-stack mode, see the Cisco Secure Workload Upgrade Guide

Usage Analytics

Site Administrators and **Customer Support users** can enable or disable usage analytics. In the navigation bar, click **Manage** > **Service Settings** > **Usage Analytics**.

Secure Workload collects data, renders anonymously through one-way hashing before sending it to the server. Configure the Privacy settings on a per-appliance basis for an on-premises appliance and a per-tenant basis for Cisco Secure Workload SaaS. You can also enable Data collection and toggle the collection on this page.

Federation

Note	This feature requires all appliances in the Federation to be running release 3.4.x or later.

Federation provides a means of joining multiple Cisco Secure Workload appliances together and consolidating much of their management to a single appliance designated the **leader**.

Setup

To enable federation, find the **Federation** menu item in the sidebar under **Company**.

Step 1 On the designated **leader**, navigate to **Company** > **Federation** and click the **Create New Federation** button.
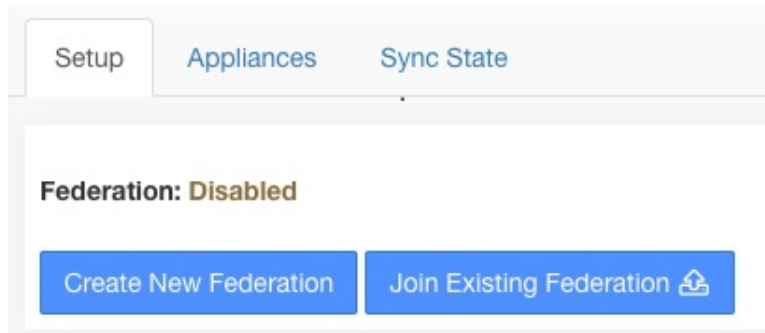
Step 2 To add the first **follower** appliance, enter its name and Fully Qualified Domain Name (FQDN) and click the **Add** button.

Step 3 Click the link to download the join certificate file.

Step 4 On the **follower**, navigate to **Company** > **Federation** and click **Join Existing Federation** and select the join certificate that is created above.

Step 5 Repeat steps 2–4 for each **follower** that will be part of the federation.

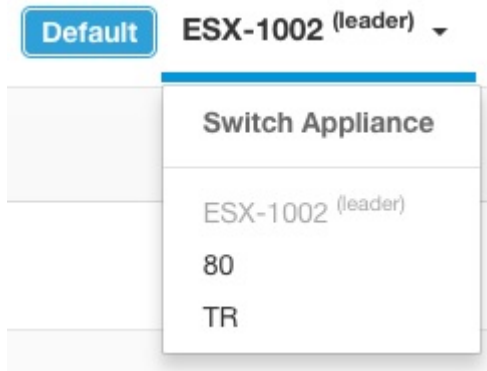*Figure 25: Create or Join Federation*

Figure 26: Federation Add Follower Form

Once federation is enabled, the header includes the name of the appliance and a selector for changing appliances.

Figure 27: Appliance Selector

If one or more of the appliances in the federation have not been seen by the leader in more than 10 minutes, an alert appears in the appliance selector, and the problematic appliances will be flagged. Hovering over them will display the last time that they synced with the leader.

Figure 28: Appliance Selector with Alerts
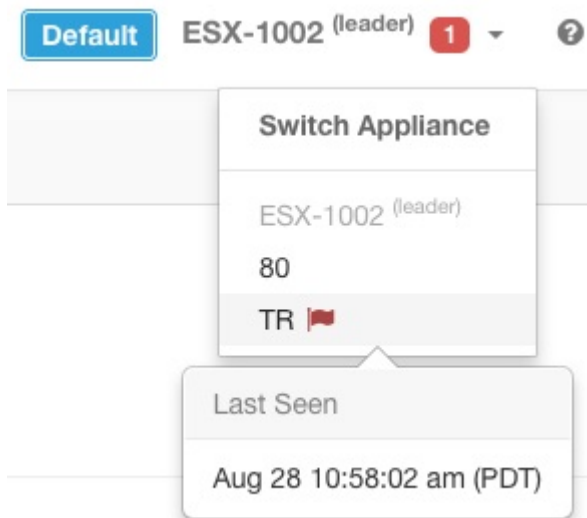
Authentication Setup

Authentication with Federation enabled, is setup using Single Sign On or SSO. SSO must be configured on each appliance part of the federation. The SSO setup is configured on the leader and each of the followers on the Company Page, External Auth Settings as indicated in Configuring Single Sign-on (SSO) on each appliance.

Administrative Tasks

Depending on the administrative task, some must be performed on the **leader** and some on the **followers**. The following table indicates the appliance type for each task.

Task	Appliance
Users	Leader
Scopes	Leader
Roles	Leader
Tenants	Leader
API Keys	Leader
Collection Rules	Leader
Software Agent Config	Leader
Software Agents	Followers
Software Agent Upgrade	Followers
Software Agent Download	Followers
Inventory Filters	Leader
Inventory Upload	Leader
Default Policy Discovery Config	Leader
Policy Order	Leader
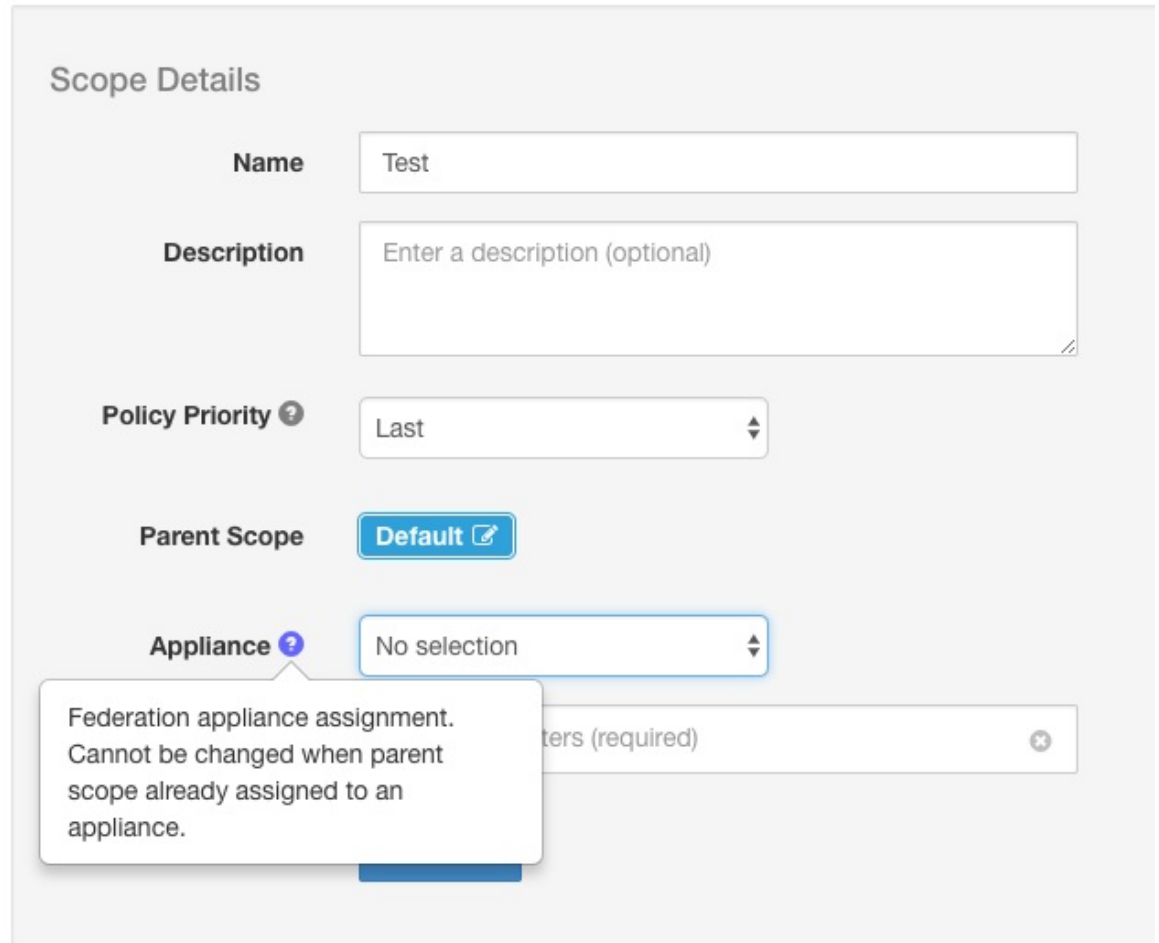
Scopes

When the inventory within a scope is managed by a single appliance, that scope can be assigned to the appliance. This enables automatic policy discovery, policy analysis, and enforcement in workspaces associated with that scope. It will also ensure policies that are created on that scope only apply to agents connected to the appliance.

Applications created on global scopes (not assigned to an appliance) cannot be used for automatic policy discovery or policy analysis. However, they can be used to enforce policies across all appliances in the federation.

An appliance can be assigned to a scope during creation or by editing the scope. All child scopes inherit the parent's appliance and cannot be assigned to a different appliance.

Figure 29: Assign Appliance to Scope

✎

Note Root-level scopes (tenants) are always global and cannot be assigned an appliance.

Workspaces

All workspaces ("applications") must be managed on the **leader**. However, flow-based charts can only be viewed on the corresponding **follower** appliance. These include the charts that are shown under the **Policy Analysis** and **Enforcement** tabs. From the **leader**, click the **View Charts on Local Appliance** to navigate to the corresponding **follower**.

Figure 30: Policy Analysis on Leader
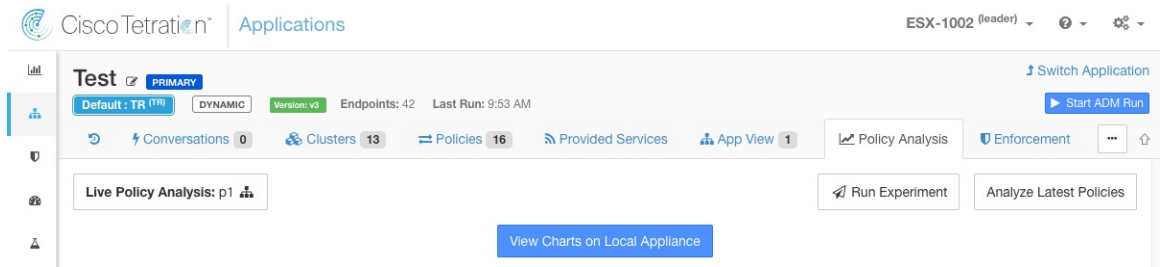
Figure 31: Policy Analysis on Follower
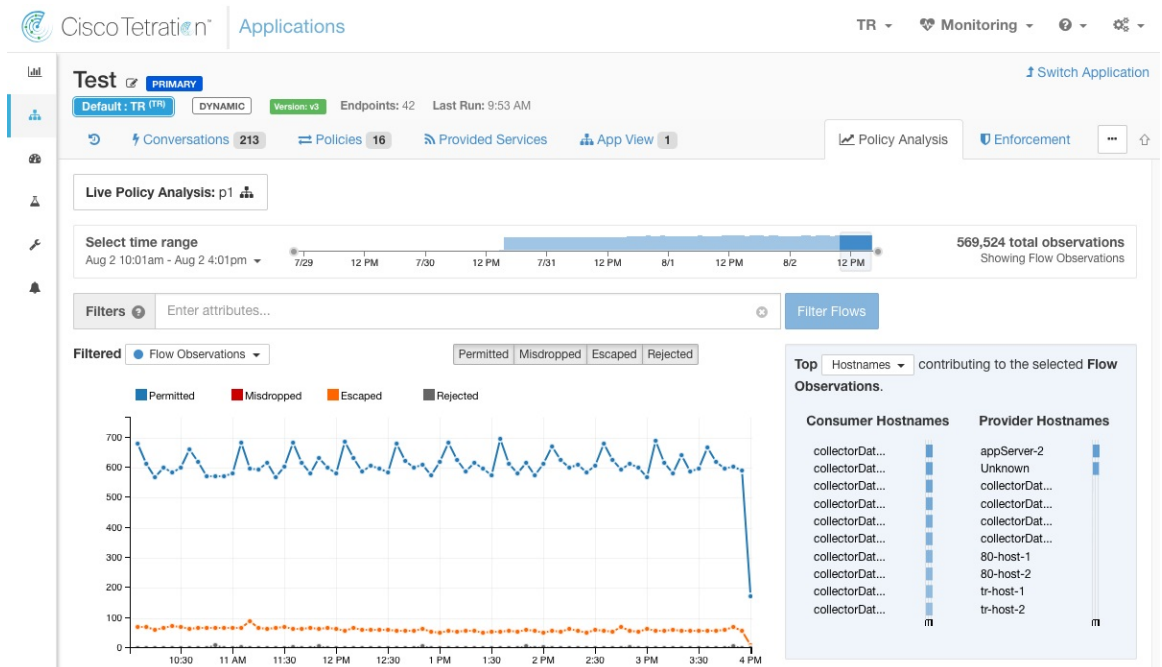
Also, searches against inventory (except for the automatic policy discovery page) are always performed locally. Therefore, it is necessary to navigate to the **follower** to view cluster, filter, and scope endpoints. The same logic applies to viewing cluster, filter, and scope details.

Figure 32: Cluster Sidebar

As explained above, workspaces that are created on global scopes cannot be used for automatic policy discovery or policy analysis. While policies can be enforced, flow-based enforcement charts are not available.

Figure 33: Policy Analysis Disabled on Global Scopes

⚠️ **Warning**
Policies using a scope or restricted inventory filter that is associated with an appliance will only be enforced on that appliance.

Software Agents

All software agents connected to any appliance within the federation are visible on the **leader**.

Step 1   Click the **Settings menu** in the top-right corner.

Step 2   Select **Agent Config**.
The **Agent Config** page displays.

Step 3   Click the **Software Agents** tab.
The **Software Agents** tab opens.

Step 4   Find one or more agents to move and click the checkboxes that are found in their table rows.

Step 5   The **Appliance** column indicates where the agent is connected.

Software Agents	Software Agent Config						

Filters ⓘ | Hostnar *contain* tes | ⊗ | Filter | Download all results

Displaying (**1** to **20**) of **22** matching results ⓘ   First Check-in ▾   ↓≡   Show 20 ▾ Items per page

∇	Hostname	Appliance	Agent Type	IP Addresses	SW Version	Platform	VRF
	test-host-122	follower-2	Enforcement		☑ 1.103.1.5-1	MSServer2008Enterprise	
	test-host-121	follower-1	Enforcement		☑ 1.103.1.5-1	MSServer2008Enterprise	

Moving Software Agents Between Follower Appliances

Software agents can be moved between **follower** appliances. From the appliance the agent or agents are connected to, follow these steps:

Step 1   Click the **Settings menu** in the top-right corner.

Step 2   Select **Agent Config**. The page displays.

Step 3   Click the **Software Agents** tab. The **Software Agents** tab opens.

Step 4   Find one or more agents to move and click the checkboxes that are found in their table rows.

Step 5   From the **–Select Appliance–** dropdown, select the desired appliance for these agents.

Step 6   Click the **Move to Appliance** button.

The table updates to indicate a move are pending. The next time the agent checks-in, it will receive a message to move appliances. Once the move has completed, the agent will no longer be visible on the original appliance. Visit the new appliance **Software Agents** page to verify the move was successful.

Other Tasks

In general, flow and inventory-based queries must be made on **follower** appliances. The following table indicates the appliance type for some common tasks.

Task	Appliance
Visibility > > **Flow Search**	Followers
Visibility > **Inventory Search**	Followers
Visibility > **Inventory Filters**	Followers
Visibility > **External Orchestrators**	Followers
Segmentation > **Automatic Policy Discovery**	Leader
Segmentation > **Policy Analysis**	Leader
Segmentation > **Enforcement History**	Leader
Segmentation > **Conversations**	Followers
Segmentation > **Analysis Results**	Followers
Segmentation > **Enforcement Results**	Followers
Monitoring > **Agents**	Followers
Monitoring > **Enforcement Status**	Followers
Monitoring > **Licenses**	Followers
Software Agents > **Change Appliances**	Followers

Any other tasks not included above should be considered *local* to the appliance. Therefore, any changes that are made or results shown only represent the state of the current appliance, not the federation. On these pages, the following alert will be shown.

Figure 34: Local Appliance Alert

The contents of this page are local to this federation appliance.

See the user guide for more information.

Existing Deployment

The following sections provide a set of guidelines for preserving data on appliances joining a Federation.

Which Data Can Be Preserved?

The user is responsible for copying users, roles, collection rules, forensic profiles, user uploaded labels and agent configs from the follower to the leader before adding it to a Federation. Data on the follower not copied to the leader is wiped and replaced with data from the leader.

We provide a mechanism for preserving scopes, filters, and policies on **followers** by exporting them to files that can then be imported on the **leader**. This is done as follows.

Step 1   On the follower appliance, navigate to **Company** > **Federation** and click the **Join New Federation** button.

Download the scopes, filters, and workspaces local to the appliance.

Figure 35: Existing Deployment Export Workflow on the Follower

Step 2   On the **leader**, navigate to **Company** > **Federation**. Add the follower by entering its name and Fully Qualified Domain Name (FQDN) and clicking the **Add** button. Then switch to the appliances view and **click the import button** to the right of the appliance FQDN.

Figure 36: Existing Deployment Import Icon on the Follower

This brings up the import wizard which prompts the user to **upload** scopes, filters, and workspaces downloaded from the follower. At every stage, the wizard waits for the user to resolve conflicts before proceeding to the next.

Figure 37: Existing Deployment Import Wizard on the Follower

Conflicts between entries on the leader and the follower are detected by comparing the names of these entries on the two appliances. For example, consider a scope **Default:host** which exists on both the leader and follower. On the leader, the query for this scope is set to **Hostname eq foo** and on the follower, it is **Hostname eq bar**. The import wizard warns the user that a conflict exists for this scope and picks the query from the leader (that is, **Hostname eq foo**).

Step 3   Finally, the user must **disable enforcement** on all existing workspaces on the follower before adding it to the Federation.

Step 4 Steps 1–3 must be repeated for each follower joining the Federation.

Which Data Cannot Be Preserved?

1. Virtual appliances (including those used with connectors) must be reprovisioned on a follower after it joins the Federation.

2. Flow data on a follower up to the point when it joins the Federation is inaccessible for scopes that are common with the leader.

Disconnected Mode of Operation

Note Applicable to followers.

Under certain circumstances such as a network partition, it makes sense to disable Federation on one or more **followers** allowing them to operate in standalone mode. To do this, navigate to **Company** > **Federation** and click the **Disable** button. A follower that is disconnected from the Federation continues operating as a standalone cluster.

New scopes, inventory filters, and workspaces on the follower can be preserved by exporting them to files that are then imported on the leader before adding it back to the Federation. This preserves changes to policies for existing workspaces. However, changes to scopes and inventory filters already present on the leader are lost once the follower rejoins the Federation.

Alerts

To enable alerts, click **Alert** > **Configuration** and update the alert configuration for Federation.

Figure 38: Federation Alerts
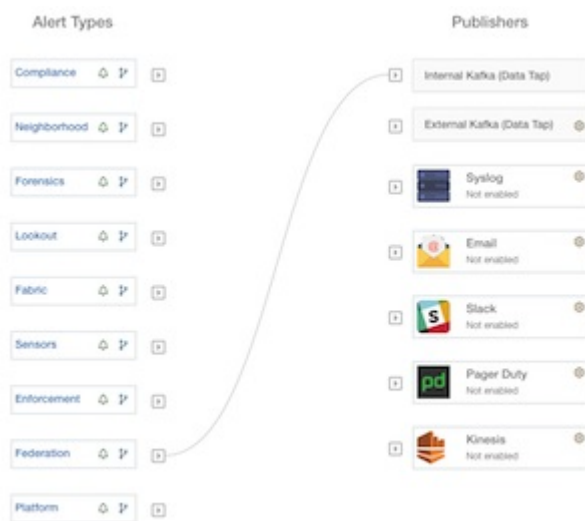
Alerts are generated for the following events:

- Alert is generated on the leader at MEDIUM severity when one or more appliances in the federation hasn't contacted it in more than 10 minutes.

- Alert is generated on the follower at MEDIUM severity when it is unable to contact the leader for over 10 minutes.

Alert Details

See Common Alert Structure for general alert structure and information about fields. The *alert_details* field is structured and contains the following subfields for federation alerts

Note Appliance is the appliance that triggered the alert.

Field	Alert Type	Format	Explanation
id	*all*	string	Appliance ID
name	*all*	string	Appliance name
fqdn	*all*	string	FQDN of the appliance
is_leader	*all*	boolean	True if the appliance is the leader
status	*all*	string	Status of the appliance
current_sw_version	*all*	string	Software version on the appliance
last_seen_at	*all*	integer	Unix timestamp of when the appliance was last seen
created_at	*all*	integer	Unix timestamp of when the appliance was created
updated_at	*all*	integer	Unix timestamp of when the appliance was updated
created_at	*all*	integer	Unix timestamp of when the appliance was created
deleted_at	*all*	integer	Unix timestamp of when the appliance was deleted
disconnected	*all*	boolean	Set to true when the follower has disconnected from the leader. Always set to false for the leader

Example of alert_details for follower-down alert

```
{
  "id": "5f219ad8755f024b46c2524a",
  "name": "esx-3018",
  "fqdn": "esx-3018.tetrationanalytics.com",
  "is_leader": false,
  "status": "Ready",
  "current_sw_version": "3.4.0.39.devel",
  "last_seen_at": 1596140582,
  "created_at": 1596037848,
  "updated_at": 1596140582,
  "deleted_at": 0,
  "disconnected": true
}
```

Example of alert_details for leader-down alert

```
{
  "id": "5f219acc755f024b46c25248",
  "name": "sherekhan",
  "fqdn": "sherekhan.tetrationanalytics.com",
  "is_leader": true,
  "status": "Ready",
  "current_sw_version": "3.4.0.39.devel",
  "last_seen_at": 1596140582,
  "created_at": 1596037848,
  "updated_at": 1596140582,
  "deleted_at": 0,
  "disconnected": false
}
```

APIs

> **Note** Credentials for a Federation cluster must be generated on the leader and can be used to query followers.

This section lists APIs added or updated for Federation:

Appliances

The appliances endpoint allows the user to retrieve the state of an appliance in a Federation.

Appliance Object

The appliance object's attributes are described in the following table:

Attribute	Type	Description
id	string	A unique identifier for the appliance.
name	string	User specified name of the appliance.
fqdn	string	User specified FQDN of the appliance.

Attribute	Type	Description
is_leader	boolean	Indicates if the appliance is a leader.
status	string	Status of the appliance.
current_sw_version	string	Cisco Secure Workload software version on the appliance.
last_seen_at	integer	Unix timestamp of when the follower was last seen by the leader. It's always null for the leader.
deleted_at	integer	Unix timestamp of when the appliance was deleted.
disconnected	boolean	Indicates if the follower has lost contact with the leader. It is set to false for the leader.

List Appliances

This endpoint returns an array of appliances in the Federation.

```
GET /openapi/v1/appliances
```

Parameters: None

Response object: Returns an array of appliance objects.

Sample python code

```
restclient.get('/appliances')
```

Scopes

The Scope object now includes the ID of the appliance that is associated with a scope. It is set to null for global scopes.

The following APIs now accept an appliance ID when creating or updating scopes.

Create a Scope

An appliance ID provided when creating a scope associates it with a specific appliance.

```
POST /openapi/v1/app_scopes
```

Parameters:

Name	Type	Description
short_name	string	User specified name of the scope.
description	string	User specified description of the scope.

Name	Type	Description
short_query	JSON	Filter (or match criteria) associated with the scope.
parent_app_scope_id	string	ID of the parent scope.
policy_priority	integer	Default is 'last'. Used to sort workspace priorities. See Policy Ordering under Review Automatically Discovered Policies.
appliance_id	string	A unique identifier for the appliance.

Sample python code

```
req_payload = {
    "short_name": "App Scope Name",
    "short_query": {
        "type":"eq",
        "field":"ip",
        "value": <....>
    },
    "parent_app_scope_id": <parent_app_scope_id>,
    "appliance_id": <appliance_id>,
}
resp = restclient.post('/app_scopes', json_body=json.dumps(req_payload))
```

Update a Scope

This API allows existing scopes to be associated with appliances using their appliance IDs.

```
PUT /openapi/v1/app_scopes/{app_scope_id}
```

Parameters:

Name	Type	Description
short_name	string	User specified name of the scope.
description	string	User specified description of the scope.
short_query	JSON	Filter (or match criteria) associated with the scope.
appliance_id	string	A unique identifier for the appliance.

Returns the modified scope object that is associated with the specified ID.

Sample python code

```
req_payload = {
    "short_name": "App Scope Name",
    "short_query": {
        "type":"eq",
        "field":"ip",
```

```
            "value": <....>
        },
        "appliance_id": <appliance_id>,
}
resp = restclient.put('/app_scopes/%s' % <app_scope_id>,
                        json_body=json.dumps(req_payload))
```

Idle Session

For those who are authenticating using a local database, this section explains how failed login attempts may lock the user account:

Step 1 Five failed login attempts using email and password result in locking the account.

Note As a security measure against probing, no specific message indicating the lock will be provided in the login interface when trying to sign in a locked account.

Step 2 Lock out interval is set at 30 minutes. After the account is unlocked, use the correct password to log in or initiate password recovery by clicking *Forgot password?*

Note Once a user is successfully signed in, one hour of inactivity logs out the user. This timeout is configured from **Manage** > **Service Settings** > **Session Configuration**.

Preferences

The **Preferences** page displays your account details and enables you to update your display preferences, change your landing page, change your password, and configure two-factor authentication.

Change Your Landing Page Preference

To change the page you see when you sign in:

Step 1 On the top-right corner of the window, click the user icon and choose **User Preferences**.

Step 2 Choose a landing page from the drop-down menu. Your preference is saved as the default or home page when you log in. To see the change, click the Secure Workload logo at the top-left corner of the page.

Changing a Password

Step 1 Click on the user icon in the top-right corner.

Step 2 Select **User Preferences**.

Step 3 In the **Change Password** pane, enter your current password in the **Old Password** field.

Step 4   Enter your new password in the **Password** field.

Step 5   Re-enter your new password in the **Confirm Password** field.

Step 6   Click **Change Password** to submit the change.

> **Note**   **Password must be 8–128 characters and contain at least one of the each following:**
>
> - Lower case letters (a b c d . . .)
> - Upper case letters (A B C D . . .)
> - Numbers (0 1 2 3 4 5 6 7 8 9)
> - Special characters (! " # $ % & ' ( ) * + , - . / : ; < = > ? @ [\ ] ^ _ ' { | } ~), space included

Recovering Passwords

This section explains how to recover your password.

Before you begin

To reset a password, you must first have an account. A new account can be added by **Site Admins** and **Customer Support users**.

Step 1   Point your browser to the Cisco Secure Workload URL and click the **Forgot Password** link. The **Forgot your password?** dialog box is displayed.

Step 2   Enter your email address in the **Email** field.

Step 3   Click **Reset Password**.

Password reset instructions are sent to your email.

> **Note**   The password recovery procedure for two-factor authentication requires contacting Secure Workload Customer Support because the email-based password recovery cannot contain the one-time password.

Enabling Two-Factor Authentication

This section explains how to enable two-factor authentication.

Step 1   Click on the user icon in the top-right corner.

Step 2   Select **User Preferences**.

Step 3   In the **Two-Factor Authentication** pane, click the **Enable** button. A new **Two-Factor Authentication** pane appears.

Step 4   Enter your password.

Step 5   Scan the QR code that is displayed under the **Current Password** field using any time-based one-time password (TOTP) app, such as Google Authenticator (for Android or iOS) or Authenticator (for Windows Phone).

Step 6   Enter the validation code that is shown by your chosen TOTP app.

Step 7   Click **Enable**.

Figure 39: Two-Factor Authentication Pane

Two-Factor Authentication

Two-factor authentication is disabled.

Current Password:

Current Password

Scan QR Code:

Scan this code using any Time-based One-Time Password (TOTP) app, such as:

- Google Authenticator for Android ↗ and iOS ↗
- Authenticator for Windows Phone ↗

Verify:

Validation Code

Enable   Cancel

The next time that you log into the system, you must select the **Use two-factor authentication** check box and enter the verification code that is shown in your TOTP app to sign in.

Note   The password recovery procedure for two-factor authentication requires contacting Secure Workload Customer Support because the email-based password recovery cannot contain the one-time password.

Disabling Two-Factor Authentication

This section explains how to disable two-factor authentication.

Step 1   Click on the user icon in the top-right corner.

Step 2   Select **User Preferences**.

Step 3   Under two-factor authentication, click the **Disable** button. The **Two-Factor Authentication** pane appears.

Step 4   Enter your password.

Step 5   Click the **Disable** button again.

You will no longer be required to enter a two-factor verification code during login.

Roles

You can restrict access to features and data using role-based access control (RBAC) model.

- User - someone with login access to Cisco Secure Workload.

- Role - user created set of capabilities that is assigned to a user.

- Capability - scope + ability pair

- Ability - collections of actions

- Action - low-level user action such as "change workspace name"

Figure 40: Role Model
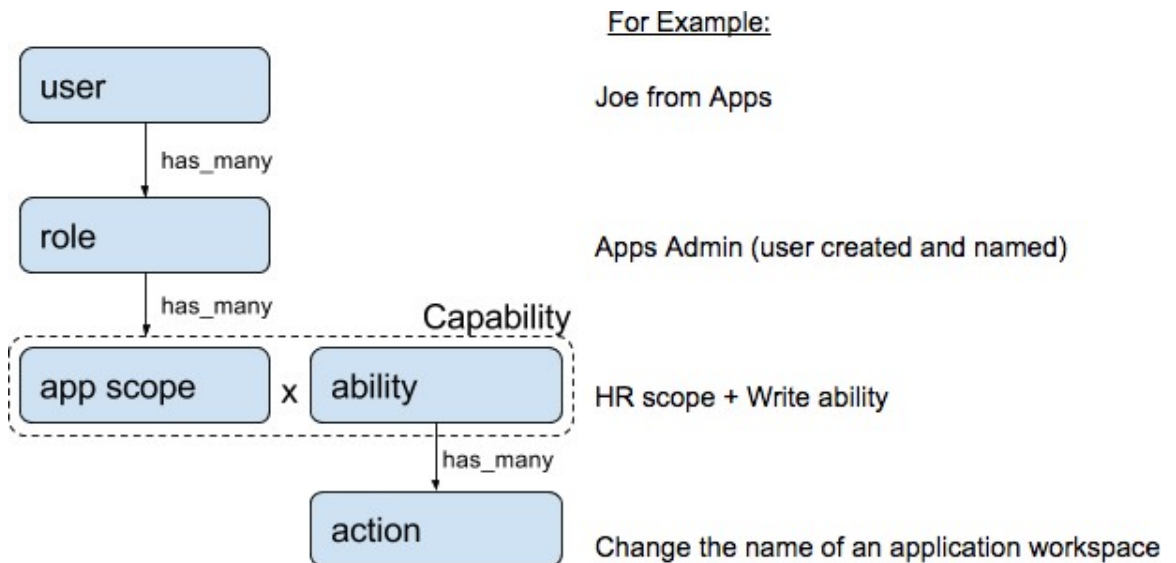
A user can have any number of roles. Roles can have any number of capabilities. For example, the "HR Search Engineer" role could have two capabilities: "Read on the HR Scope" to give visibility and context and "Execute on "HR:Search" capability to allow the engineers assigned this role to make specific changes that are related to their applications.

Roles contain sets of Capabilities and are assigned to users on the **Users** page. A user can have any number of roles. Roles can have any number of capabilities.

Role	Description
Agent Installer	Provide the ability to manage agents life cycle including install, monitor, upgrade, and convert, but **cannot** delete agents and access agent config profile.
Customer Support	For Technical Support or Advanced Services. Provides access to cluster maintenance features. Allows the same access as Site Admin, but **cannot** modify users.
Site Admin	Provides the ability to manage users, agents, and so on. Can view and edit all features and data. There must be at least one site admin.
Global Application Enforcement	Provides the Enforce ability on every scope.
Global Application Management	Provides the Execute ability on every scope.
Global Read Only	Provides the Read ability on every scope.

Abilities and Capabilities

Roles are made up of capabilities which include a scope and an ability. These define the allowed actions and the set of data that they apply to. For example, the (HR, Read) capability should be read and interpreted as "Read ability on the HR scope". This capability would allow access to the HR scope and all its children.

Ability	Description
Installer	Install, monitor, and upgrade software agents.
Audit	Global appliance data read support and access to change logs.
Read	Read all data including flows, application, and inventory filters.
Write	Make changes to applications and inventory filters.
Execute	Perform Automatically discover policies run and publish policies for analysis.
Enforce	Enforce policies that are defined in application workspaces that are associated with the given scope.
Owner	Required to toggle an application workspace from secondary to primary. Access to Data Tap Admin abilities, such as managing User App sessions, adding Data Taps, and creating Visualization Data Sources.

☞

Important   Abilities are inherited, for example, the Execute ability allows all the Read, Write, and Execute actions.

☞

Important   Abilities apply to the scope and all the scope's children.

Menu Access by Role

The menus a user can see and use depend on the user's assigned role:

Table 2: Overview Menu

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Overview	Overview	Yes	Yes	Yes	Yes	Yes	No

Table 3: Organize Menu

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Organize	Scopes and Inventory	Yes	Yes	Yes	Yes	Yes	No
Organize	Use Uploaded Labels	Yes	Yes	No	No	No	No
Organize	Inventory Filters	Yes	Yes	Yes	Yes	Yes	No

Table 4: Defend Menu

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Defend	Segmentation	Yes	Yes	Yes	Yes	No	No
Defend	Enforcement Status	Yes	Yes	No	No	No	No
Defend	Policy Templates	Yes	Yes	No	No	No	No
Defend	Forensic Rules	Yes	Yes	No	No	No	No

Table 5: Investigate Menu

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Investigate	Traffic	Yes	Yes	Yes	Yes	Yes	No
Investigate	Alerts	Yes	Yes	Yes	Yes	Yes	No
Investigate	Vulnerabilities	Yes	Yes	Yes	Yes	Yes	No
Investigate	Forensics	Yes	Yes	Yes	Yes	Yes	No
Investigate	Neighborhood	Yes	Yes	Yes	Yes	Yes	No

Table 6: Manage Menu

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Manage	Agents	Yes	Yes	No	No	No	Yes
Manage	Alerts Configs	Yes	Yes	Yes	Yes	Yes	No
Manage	Change Logs	Yes	No	No	No	No	No
Manage	Connectors	Yes	Yes	No	No	No	No
Manage	External Orchestrators	Yes	Yes	No	No	No	No
Manage	Secure Connector	Yes	Yes	No	No	No	No
Manage	Virtual Appliances	Yes	Yes	No	No	No	No
Manage	Users	Yes	Yes	No	No	No	No
Manage	Roles	Yes	Yes	No	No	No	No
Manage	Threat Intelligence	Yes	Yes	No	No	No	No
Manage	Licenses	Yes	No	No	No	No	No
Manage	Collection Rules	Yes	Yes	No	No	No	No
Manage	Session Configuration	Yes	Yes	No	No	No	No

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Manage	Usage Analytics	Yes	Yes	No	No	No	No
Manage	Data Tap Admin	Yes	No	No	No	No	No

Table 7: Platform Menu

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Platform	Tenants	Yes	Yes	No	No	No	No
Platform	Cluster Configuration	Yes	Yes	No	No	No	No
Platform	Outbound HTTP	Yes	Yes	No	No	No	No
Platform	Collectors	Yes	Yes	No	No	No	No
Platform	External Authentication	Yes	Yes	No	No	No	No
Platform	SSL Certificate	Yes	Yes	No	No	No	No
Platform	Login Page Message	Yes	Yes	No	No	No	No
Platform	Federation	See below	See below	No	No	No	No
Platform	Data Backup	See below	See below	No	No	No	No
Platform	Data Restore	See below	See below	No	No	No	No
Platform	Upgrade/ Reboot/ Shutdown	Yes	Yes	No	No	No	No

Note
- The Federation option is available to Site Admin and Customer Support roles if Federation is enabled.
- Data Backup and Restore options are available to Site Admin and Customer Support roles if data backup and restore are enabled.

Table 8: Troubleshoot Menu

Menu	Option	Site Admin	Customer Support	Global Application Management	Global Read Only	Global Application Enforcement	Agent Installer
Troubleshoot	Service Status	Yes	Yes	No	No	No	No
Troubleshoot	Cluster Status	See below	See below	No	No	No	No
Troubleshoot	Virtual Machine	Yes	Yes	No	No	No	No
Troubleshoot	Snapshots	Yes	Yes	No	No	No	No
Troubleshoot	Maintenance Explorer	Yes	Yes	No	No	No	No
Troubleshoot	Resque	Yes	Yes	No	No	No	No
Troubleshoot	Hawkeye (Charts)	Yes	Yes	No	No	No	No
Troubleshoot	Abyss (Pipeline)	Yes	Yes	No	No	No	No

Note The Cluster Status option is available to Site Admin and Customer Support roles if the cluster type is 'Physical' or 'OCI'.

Create a Role

Before you begin

You must already have a **Site Admin** or **Customer Support** user role.

1. In the navigation bar on the left, click **Manage** > **User Access** > **Roles**.

2. Click **Create New Role**. The **Roles** panel appears.
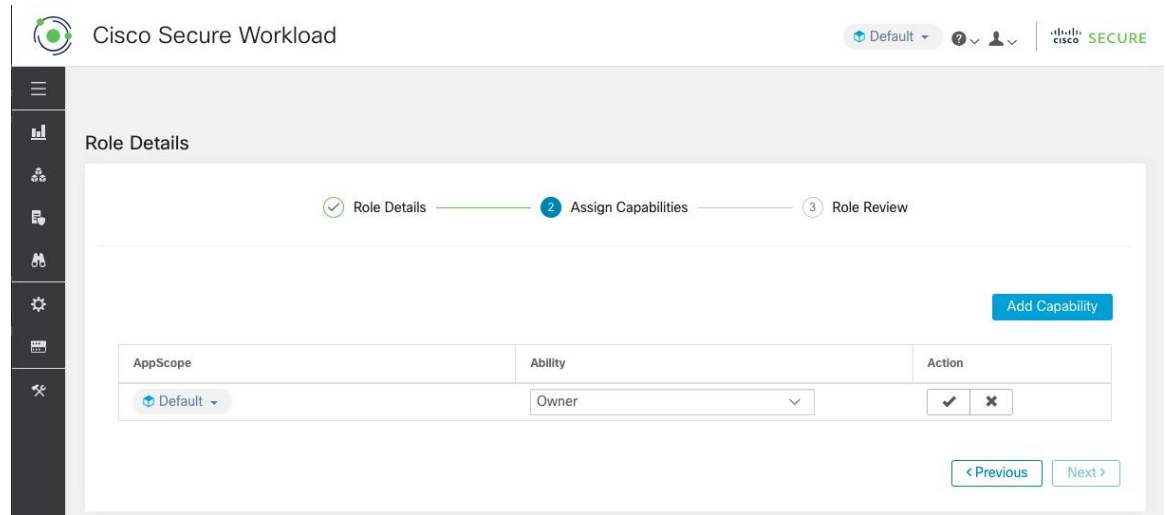
Creating a role using the Create Role Wizard is three-step process.

Step 1   a)   Enter the appropriate values in the following fields:

Field	Description
Name	The name to identify the role.
Description	A short description to add context about the role.

　　　　b)   Click the **Next** button to move to the next step or **Back to Roles Page** to go back to Roles Page.

Step 2   a)   Click the **Add Capability** button to show the creation form in the top row.
　　　　b)   Select scope and ability.
　　　　c)   Click the **Checkmark** button to create a new capability or **Cancel** button to cancel.
　　　　d)   Click **Next** to review role details or **Previous** to go back and edit.

Figure 41: Capability Assignment

Step 3   a)   Review the role details and capabilities.

b)   Click **Create** to create role.

Figure 42: Role Review

Edit a Role

This section explains how **Site Admins** and **Customer Support users** can edit roles.

Before you begin

You must be Site Admin or Customer Support User.

1. In the navigation bar on the left, click **Manage** > **User Access** > **Roles**.

2. In the row of the role to edit, click the **Edit** button in the right-hand column. The **Roles** panel appears.

Editing a role using the Edit Role Wizard is three-step process.

Step 1
   a) Update the name or description if desired.
   b) Click the **Next** button to move to the next step or **Back to Roles Page** to go back to Roles Page.

Step 2
   a) Remove any capability as needed. In the row of the capability to delete, click the **Delete** icon in the right-hand column.
   b) To add, click the **Add Capability** button to show the creation form in the top row.
   c) Select scope and ability.
   d) Click **Next** to review role details or **Previous** to go back and edit.

Step 3
   a) Review the role details and capabilities.
   b) Click **Update** to create the role or **Previous** to go back and edit. Changes to role details and capability assignment are saved after **Update**.

Note   Capabilities cannot be edited, they must be deleted and recreated.

Change Log – Roles

Site Admins and users with the SCOPE_OWNER ability on the root scope can view the change logs for each role by clicking on the icon in the **Action** column as shown in the following figure.

Figure 43: Change Log

These users can also view a list of deleted roles by clicking on the **View Deleted Roles** link below the table.

For more information on the **Change Log**, see Change Log. Root scope owners are restricted to viewing change log entries for entities belonging to their scope.

Scopes

Note   The **Scopes** page is merged with **Inventory Search**. For more information, see the Scopes and Inventory page.

Tenants

Site Admins and **Customer Support users** can access the **Tenants** page under the **Platform** > **Tenants** menu in the left navigation pane. The Tenants page displays the currently configured Tenants and VRFs. Secure Workload is preconfigured with one or more Tenants and VRFs, and you can add, edit, and delete tenants.

✎

Note   These values affect the results of the cluster output. We recommend consulting Cisco TAC before changing these values to understand the system impact.

Figure 44: Tenants Page

VRF ID ↑↓	Name ↑	Description	Switch VRF Count	Tenant ID ↑↓	Action
1	Default		0	0	✏ 🗑
676767	Tetration		0	676767	✏ 🗑
0	Unknown		0	0	✏ 🗑

Add a Tenant

Before you begin

You must be a **Site Admin** or **Customer Support** user.

Step 1   In the left navigation pane, click **Platform** > **Tenants**.

Step 2   Click **Create New Tenant**.

Step 3   Enter the appropriate values in the following fields:

Field	Description
Name	Enter a desired name for the tenant.
Description	(Optional) The description field contains additional information about the tenant.

Step 4   Click **Create**.

Edit a Tenant

Before you begin

You must be a **Site Admin** or **Customer Support** user.

Step 1 In the left navigation pane, click **Platform** > **Tenants**.

Step 2 Find the tenant that you want to edit and click the **pencil** icon in the column on the right.

Field	Description
Name	Update a name for the tenant.
Description	(Optional) Update the description field contains additional information about the tenant.
VRF ID	Displays the ID for this particular Tenant or VRF.
Change log	Clicking the change log icons displays a new page which displays the change log for the Tenant or VRF.

Step 3 Click **Update**.

Users

Site Admins and Root Scope Owners can access the **Users** page under the **Manage** > **User Access** menu in the navigation bar at the left side of the window.

This page shows all Service Provider users and those associated with the scope that is selected on the page header.

Multitenancy

To support multitenancy, users can be assigned to a root scope. These users can be managed by users with the 'Owner' ability on the root scope and can only be assigned roles that are associated with the same scope.

Users without a scope are called 'Service Providers' and they can be assigned any role allowing them to perform actions across root scopes.

Adding a New User Account

This section explains how **Site Admins** and Users with the **SCOPE_OWNER** ability on the root scope can add new user accounts.

If a user is assigned a scope for multitenancy, only roles that are assigned to the same scope may be selected.

Note This page is filtered by the scope preference that is selected on the page header.

Before you begin

1. In the left navigation pane, click **Manage** > **User Access** > **Users**.

2. If applicable, select the appropriate root scope from the top right of the page.

3. Click the **Add New User** button. The **Users** wizard appears.

   User creation is a three-step process.

Step 1 **a.** Enter the appropriate values in the following fields:

Field	Description
Email	Enter the new user's email address, it is non case-sensitive. We use the lower cased version of your email if it contains letters.
First Name	Enter the new user's first name.
Last Name	Enter the new user's last name.
Scope	Root Scope that is assigned to the user for multitenancy.

You can also optionally import an SSH Public Key now, or do so later.
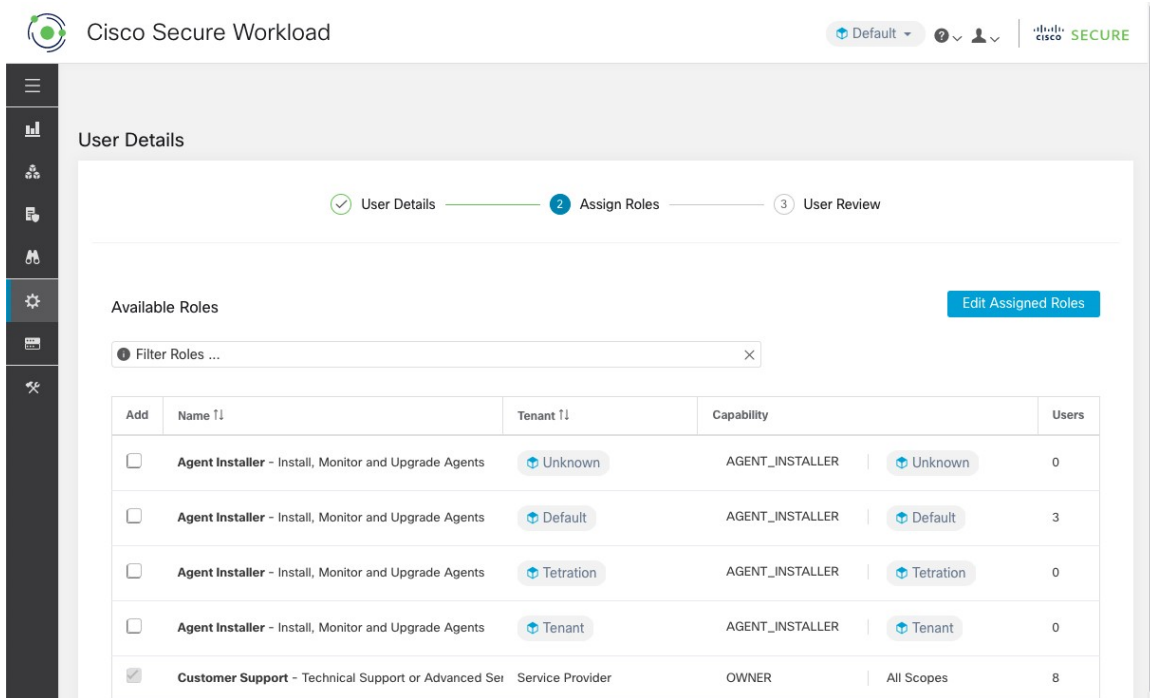
b. Click the **Next** button to move to the next step or **Back to Users List** to go back to the Users Page.

Step 2 In this view you can 'Add Roles', 'Delete Roles' or 'Select Roles':

a. Click on **Add Roles** to assign Roles.

Figure 45: Available Roles

b. Click on **Edit Assigned Roles** to delete them.

c. Filter roles by Name and Tenant.

Figure 46: Filter Roles

Step 3   Review selections and click **Create**.

If external authentication is enabled, authentication details are displayed.

Note   After user creation, the user will receive an email to set up password.

Editing a User Account

Before you begin

You must be a **Site Admin** or **Root Scope Owner** user.

Note   This page is filtered by the scope preference that is selected on the page header.

1. In the navigation bar on the left, click **Manage** > **User Access** > **Users**.

2. If applicable, select the appropriate root scope from the top right of the page.

3. In the row of the account you want to edit, click **Edit** button in the right-hand column. The **Users** Wizard appears.

Editing user using the wizard is a three-step process.

Step 1   **a.**   Update the following fields, if desired:

Field	Description
Email	Update the new user's email address.
First Name	Update the new user's first name.
Last Name	Update the new user's last name.
Scope	Root Scope that is assigned to the user for multitenancy. (Available to site admins)

b.   Click **Next** to go to Role Assignment.

Step 2   **a.**   In this view, assigned roles can be removed.

b.   Click on **Add Roles** to assign new roles.

c.   Click **Next** button to review the user details and role assignment or **Previous** button to go back and edit details.

Step 3   **a.**   Review the user details and role assignment.

> **b.** Click **Update** to update the user or **Previous** to go back and edit roles. Changes to user details and role assignment are saved.

If external auth is enabled, authentication details are displayed.

Deactivating a User Account

![Note icon]

Note　To maintain consistency of change log audits, users can only be deactivated, they are not deleted from database.

Before you begin

You must be a **Site Admin** or **Root Scope Owner** user.

![Note icon]

Note　This page is filtered by the scope preference that is selected on the page header.

Step 1　In the navigation bar on the left, click **Manage** > **User Access** > **Users**.

Step 2　If applicable, select the appropriate root scope from the top right of the page.

Step 3　In the row of the account you want to deactivate, click **Deactivate** button in the right-hand column.

To view deactivated users, toggle **Hide Deleted Users** button.

Reactivating a User Account

If a user has been deactivated, you can reactivate the user.

Before you begin

You must be a **Site Admin** or **Root Scope Owner** user.

![Note icon]

Note　This page is filtered by the scope preference that is selected on the page header.

Step 1　In the navigation bar on the left, click **Manage** > **User Access** > **Users**.

Step 2　If applicable, select the appropriate root scope from the top right of the page.

Step 3　Toggle **Hide Deleted Users** to display all users, including deactivated users.

Step 4　For the required deactivated account, click **Restore** in the right-hand column to reactivate the account.

Import SSH Public Key

To enable SSH access as **ta_guest** user via one of the collector IP addresses, SSH public key can be imported for each user. This menu will only be available to **Site Admins** and users with the `SCOPE_OWNER` ability on the root scope. The SSH Public Key automatically expires in 7 days.

Site Configuration in Secure Workload Setup

This section explains how **Site Admins** can set up a site during the Secure Workload Setup process.

Field	Description
UI Admin Email	The email address of the individual who will be responsible for administering Secure Workload within your organization.
UI Primary Customer Support Email	The email address of primary support. Must be different from UI Admin Email.
Admiral Alert Email	This email address receives alerts that are related to the cluster health. Must be different from UI Admin Email and UI Primary Customer Support Email.

The email addresses are non case-sensitive. We use the lower cased version of your email if it contains letters.

Figure 47: Configure UI Admin, Primary Customer Support, and Admiral Admin Alert Emails

Change Log – Users

Site Admins and users with the SCOPE_OWNER ability on the root scope can view the change logs for each user by clicking on the icon in the **Actions** column as shown in the following figure.

Figure 48: Change Log

For more information on the **Change Log**, see Change Log. Root scope owners are restricted to viewing change log entries for entities belonging to their scope.