

User Management

Date published: 2019-08-22

Date modified:



Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

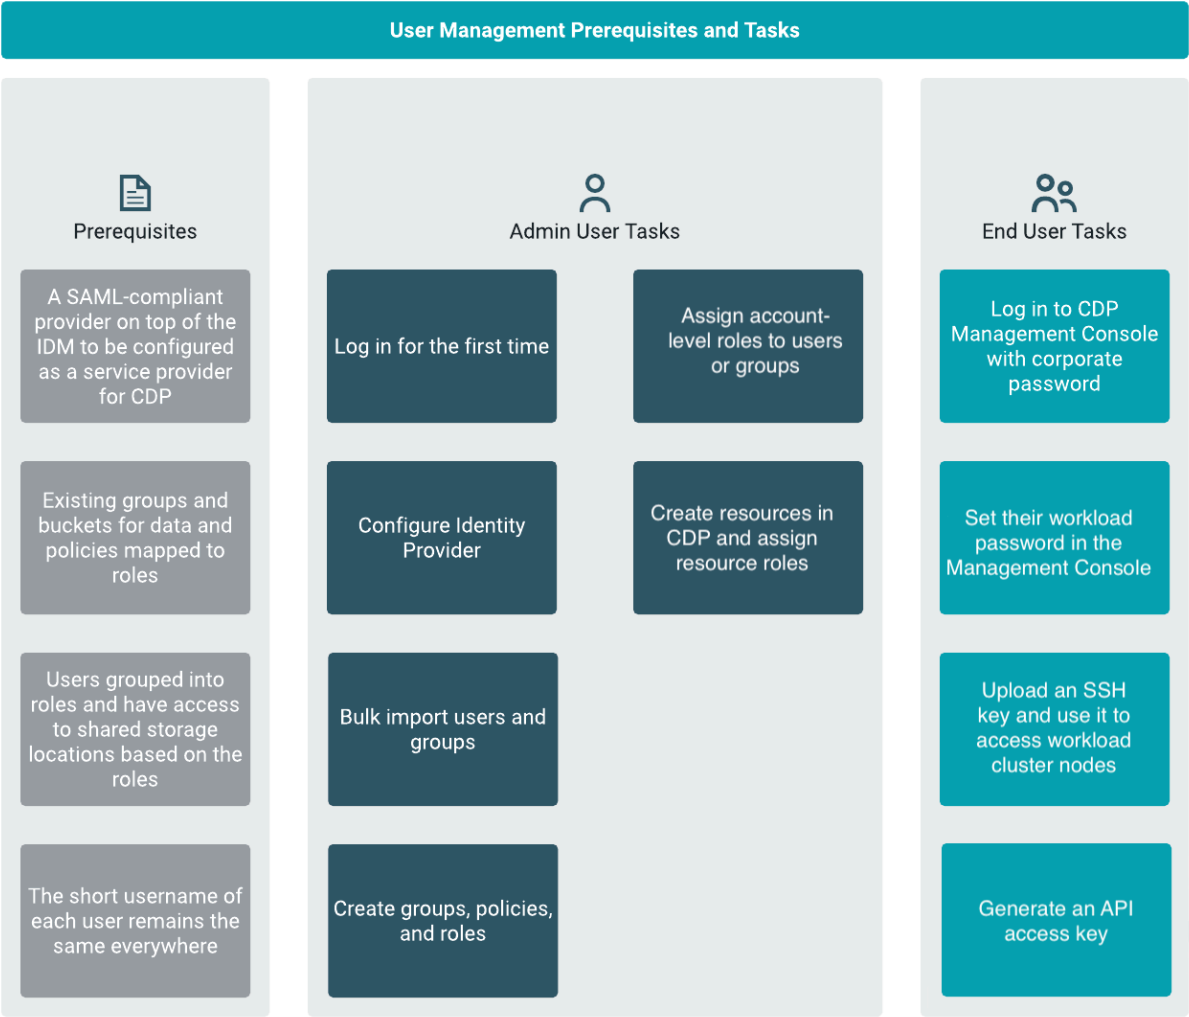
Managing user access.....	5
Onboarding users.....	6
Configuring identity providers in CDP.....	6
Generate IdP metadata.....	7
Set up IdP in CDP.....	7
Set up IdP as service provider.....	8
Synchronizing group membership.....	10
Updating an identity provider.....	11
Disabling the Cloudera SSO login.....	12
Configure AAD in CDP.....	12
Configure SCIM with Azure AD.....	19
Importing or uploading users.....	23
Generating workload usernames based on email.....	23
Known issues and troubleshooting related to IdP setup in CDP.....	24
Understanding CDP user accounts.....	25
CDP account administrator.....	25
CDP user.....	25
CDP workload user.....	26
CDP machine user.....	26
Understanding CDP roles.....	27
Account roles.....	27
Resource roles.....	28
Group membership admin roles.....	34
Example role assignment scenario.....	34
User and group limits.....	35
Managing users and machine users.....	35
Creating a machine user in CDP.....	35
Deleting users and machine users.....	36
Assigning account roles to users.....	37
Assigning resource roles to users.....	38
Assign environment role.....	38
Assign shared resource role.....	39
Assign Data Hub role.....	41
Assign classic cluster role.....	42
Managing groups in CDP.....	43
Reserved group names.....	43
Understanding CDP groups.....	44

Synchronizing group membership.....	44
Creating a group.....	46
Adding or removing a user.....	46
Assigning account roles to groups.....	48
Assigning resource roles to groups.....	49
Assign environment role.....	49
Assign shred resource role.....	50
Assign Data Hub role.....	51
Assign classic cluster role.....	52
Assigning a group admin.....	53
Updating a group.....	54
Removing account roles from a group.....	55
Deleting a group.....	56
 Performing user sync.....	 56
 Access paths to CDP.....	 58
 Setting a default identity provider in CDP.....	 59
 Logging in as workload user.....	 60
 Setting the workload password.....	 60
 Managing SSH keys.....	 63
 Generating an API access key.....	 64
 Retrieve keytabs for workload users.....	 65

Managing user access and authorization

To provide access to resources such as environments and clusters, you must add users and groups and assign roles and resources to them.

Using the CDP Management Console, you can perform the following tasks:



[Generating an API access key](#)

[Retrieve keytabs for workload users](#)

Onboarding users

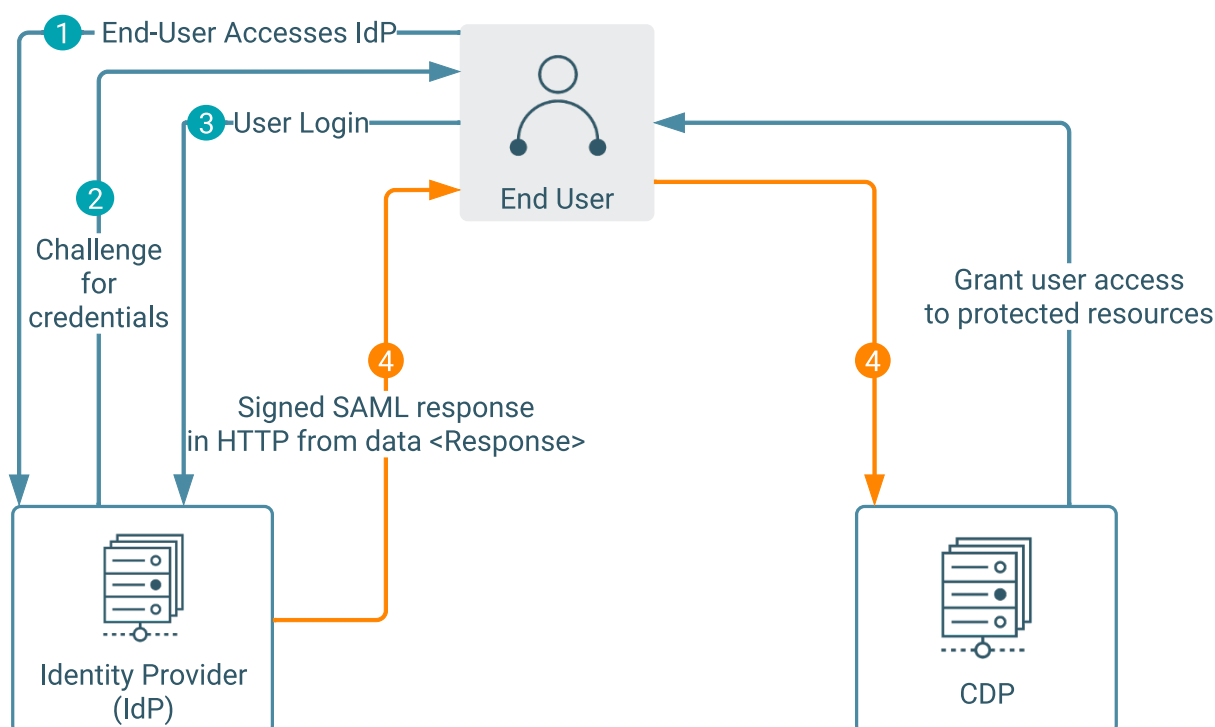
To enable users to work on the various CDP components and services, you can onboard them by configuring identity providers or importing users in bulk.

Configuring identity providers in CDP

An account administrator or PowerUser can onboard users by setting up identity federation with CDP.

If your organization uses an enterprise identity provider (IdP) that is compliant with Security Assertion Markup Language (SAML), you can set up identity federation with CDP. Identity federation allows users within your organization to log in to CDP through the authentication system in your organization without registering with Cloudera or creating a Cloudera account.

The following diagram illustrates how identity federation works with CDP:



Note:

As shown in the diagram, there is no network communication required between CDP and customer IdP, so there is no need to create firewall rules.

CDP supports the following:

- CDP supports the SAML 2.0 standard. You can set up any identity provider for CDP that uses SAML 2.0.
- You can set up a maximum of 10 SAML 2.0-compliant identity providers in CDP.

Setting up an identity provider for CDP involves the following steps:

1. The IdP administrator in your organization generates the SAML metadata that describes your enterprise IdP.
2. The CDP administrator sets up the identity provider in CDP.
3. The IdP administrator configures the enterprise IdP in your organization to work with CDP as a service provider.

Generating the identity provider metadata

Use your enterprise IdP user interface to generate the identity provider SAML metadata file.

CDP has the following requirements for the identity provider SAML metadata file:

- The file must be a valid XML file.
- The metadata must include at least one IDPSSODescriptor element.
- The metadata must contain information about at least one valid x.509 certificate that can be used to verify signed assertions.

The following XML file example shows the elements to include in the identity provider SAML metadata file:

```
<?xml version="1.0" encoding="UTF-8"?>
<md:EntityDescriptor xmlns:md="urn:oasis:names:tc:SAML:2.0:metadata" enti
tyID="http://www.IdP.com/entity_ID">
  <md:IDPSSODescriptor WantAuthnRequestsSigned="false" protocolSupportEnum
eration="urn:oasis:names:tc:SAML:2.0:protocol">
    <md:KeyDescriptor use="signing">
      <ds:KeyInfo xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
        <ds:X509Data><ds:X509Certificate>full_x509-certificate_stri
ng</ds:X509Certificate></ds:X509Data>
      </ds:KeyInfo>
    </md:KeyDescriptor>
    <md:NameIDFormat>urn:oasis:names:tc:SAML:1.1:nameid-format:emailAdd
ress</md:NameIDFormat>
    <md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindin
gs:HTTP-POST"
      Location="https://application.IdP.com/app/.../sso/saml"/>
    <md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:
HTTP-Redirect"
      Location="https://application.IdP.com/app/.../sso/saml"/>
    </md:IDPSSODescriptor>
  </md:EntityDescriptor>
```



Note:

The WantAuthnRequestsSigned=true option is not supported.

Setting up the identity provider in CDP

In CDP, you must create an identity provider to capture the SAML metadata and connection information for your enterprise IdP. To create an identity provider in CDP, you must be a CDP account administrator or have the PowerUser role.

About this task



Note:

Creating and integrating with identity providers should be considered a privileged action and you must think about how it will affect group memberships within CDP. Each identity provider manages their own unique set of group names and memberships. However, different identity providers can define the same group names. When users from different identity providers federate to CDP (with group sync on) the identity providers may provide the same group names. Since groups in CDP are defined by their name - as provided by the identity provider - when this happens, users will be added to the same group in CDP, even though they federated from different identity providers. This may grant unexpected permissions. If your organization needs group names to be unique across all identity providers federating to CDP, our recommended approach is to create different CDP accounts for each identity provider, and only set up one identity provider to federate to a CDP account.

**Note:**

There are certain group names that are reserved and therefore cannot be synchronized to CDP. See [Reserved group names](#).

Required role: Account administrator or PowerUser

Procedure

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Identity Providers.
4. Click Create Identity Provider.
5. On the Create Identity Provider window, enter the name you want to use for the CDP identity provider.
6. Select whether to synchronize the user group membership in CDP with the user group membership in your enterprise IdP.
7. To synchronize the groups, select the Sync Groups on Login option.

For more information about user group synchronization, see [Synchronizing group membership](#).

8. In Provider Metadata, select File Upload to upload a file that contains the identity provider SAML metadata or select Direct Input to paste the identity provider SAML metadata directly.
9. Click Create.

Results

CDP adds the new identity provider to the list of CDP identity providers on the Identity Providers page.

After you create the identity provider in CDP, you can view its properties to get the information you need to configure your enterprise IdP to work with CDP.

On the Identity Providers page, click the name of the new CDP identity provider to see its properties:

Property	Description
Name	Name of the CDP identity provider.
ID	ID generated for the CDP identity provider.
Sync Groups on Login	Indicates whether CDP synchronizes a user's group membership in CDP with the user's group membership in your enterprise IdP when a user logs in. For more information about user group synchronization, see Group Membership Synchronization .
CRN	The Cloudera resource name assigned to the CDP identity provider.
SAML Identity Provider Metadata	The identity provider SAML metadata for your enterprise IdP that you provided when you created the CDP identity provider.
Generate workload username by email	You can optionally check this if you use an opaque ID for SAML NameID and SCIM userName so that the workload username is generated based on the email instead of the default. For more information, see Generating workload usernames based on email .
Enable SCIM	You can optionally check this to enable SCIM for Azure AD. For more information, see Configure SCIM with Azure AD .
CDP SAML Service Provider Metadata	The CDP SAML service provider metadata to configure your enterprise IdP.

Configuring your enterprise IdP to work with CDP as a service provider

CDP provides a service provider SAML metadata file that describes the information that CDP requires to enable users to log in to CDP through your enterprise IdP.

You can get the CDP SAML metadata XML from the Identity Providers page in CDP web interface by navigating to the details of your identity provider configuration.

The CDP SAML metadata file includes the following information:

Information	Attribute	Description
Name ID formats that CDP supports	NameIDFormat	<p>The metadata includes multiple name ID formats. Use one of the formats in the list for the user ID.</p> <p>CDP supports any type of name ID format other than transient. Cloudera requires that you use name ID formats that are globally unique within your identity provider. The name ID format should also be stable over time. Cloudera does not recommend using email addresses because, although they can be unique, they are typically not stable over time.</p> <p>If your NameID is an opaque ID (such as a UUID), you can Generate workload usernames based on email.</p> <p>The value of NameID is case-sensitive.</p>
CDP SSO URL	Location	<p>It should be the same as the "Location" value of the "<AssertionConsumerService>" in the CDP SAML Service Provider metadata. It has a fixed format.</p> <p>For CDP Control Plane region us-west-1: https://consoleauth.altus.cloudera.com/saml?samlProviderId=CDP-assigned-ID</p> <p>For any other CDP Control Plane region:</p> <p><a href="https://consoleauth.<CONTROL_PLANE_REGION>cdp.cloudera.com/consoleauth/saml?samlProviderId=CDP-assigned-ID">https://consoleauth.<CONTROL_PLANE_REGION>cdp.cloudera.com/consoleauth/saml?samlProviderId=CDP-assigned-ID</p> <p>For more information about the ID that CDP generates and assigns to the CDP identity provider, see Setting Up the Identity Provider in CDP.</p> <p>This attribute is required.</p>
Endpoint for binding	Binding	<p>Use the following URN as the endpoint that your enterprise IdP must bind to:</p> <p>urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST</p> <p>This attribute is required.</p>
User email address	RequestedAttribute: mail	<p>Set the email address attribute to the following URN:</p> <p>urn:oid:0.9.2342.19200300.100.1.3</p> <p>CDP also accepts:</p> <p>mail</p> <p>This attribute is required. Although CDP requires the user email address, it is used for display purposes only.</p>

Information	Attribute	Description
(Optional) List of groups that the user is a member of	RequestedAttribute: groups	Set the group list attribute to the following URN: https://cdp.cloudera.com/SAML/Attributes/groups This attribute is optional. For more information about the group list and how CDP synchronizes group membership, see Synchronizing Group Membership .
(Optional) User first name	RequestedAttribute: firstName	Set the user first name attribute to the following URN: https://cdp.cloudera.com/SAML/Attributes/firstName This attribute is optional; used for display purposes only.
(Optional) User last name	RequestedAttribute: lastName	Set the user last name attribute to the following URN: https://cdp.cloudera.com/SAML/Attributes/lastName This attribute is optional; used for display purposes only.

If your enterprise IdP allows it, you can upload the CDP SAML metadata file to your enterprise IdP. Otherwise, use your enterprise IdP user interface to set up CDP as a service provider.

Synchronizing group membership

CDP can synchronize the user's group membership provided by your enterprise IdP with the user's group membership in CDP.

When a user initially logs in to CDP through the identity management system in your organization, CDP creates a CDP user account for the user. However, without being assigned CDP roles, the user cannot perform tasks in CDP. Cloudera recommends that you create CDP groups with assigned roles and add users to the groups so that the users can take on the roles assigned to the groups.

When you create an identity provider, you can select the Sync Groups on Login option to enable CDP to synchronize the user group membership. By default, the Sync Groups on Login option is disabled. Clear the option selection if you do not want CDP to synchronize the user group membership.

Group names must be alphanumeric, may include dots (.), hyphens (-), and underscores (_), and must be fewer than 64 characters long. Additionally, names can only start with an alphabetic character or an underscore.



Note:

There are certain group names that are reserved and therefore cannot be synchronized to CDP. See [Reserved group names](#).



Note: Cloudera has default limits in place with regard to how many users, machine users, and groups can be added per account. Review [User and group limits](#) and make sure that you do not exceed these limits.

Sync Groups on Login enabled

When the Sync Groups on Login option is enabled, CDP synchronizes a user's group in the following manner:

- The group membership that your enterprise IdP specifies for a user overrides the group membership set up in CDP. Each time a user logs in, CDP updates the user's group membership based on the groups that your enterprise IdP specifies for the user.
- If the group exists in CDP, CDP adds the user to the group. The user takes on all the roles associated with the group.

- If the group does not exist in CDP, CDP creates the group and adds the user to the group. However, no roles are assigned to the new group, so a member of the new group does not take on roles from the group.
- If the user is a member of a group in CDP that is not included in the list provided by your enterprise IdP, CDP removes the user from the group.
- If the list of groups from your enterprise IdP is empty, CDP removes the user from all groups in CDP. After login, the user will not be a member of any CDP group and will not have roles from any group.

To ensure that users can perform tasks in CDP, Cloudera recommends that you set up the groups in CDP with appropriate roles before you assign them to users.

Sync Groups on Login disabled

When the Sync Groups on Login option is disabled, CDP does not synchronize the user's group membership in CDP with the user's group membership provided by the IdP. After login, a user's group membership in CDP is determined by the CDP groups assigned to the user in CDP. The groups assigned to the user in your enterprise IdP are ignored.

Sync Membership option for a newly created group

Additionally, once you have synced your IdP and you create a new group in CDP, you have an option called Sync Membership that determines whether group membership is synced to IdP when a user logs in. By default, Sync Membership is enabled when Sync Groups on Login is enabled.

The following table describes how the global Sync Groups on Login and the per-group Sync Membership options can be used:

	IdP Sync Groups on Login on	IdP Sync Groups on Login off
Group Sync Membership on	Group membership for the specific group is reflected in IdP.	Group membership for the specific group is not reflected in IdP.
Group Sync Membership off	Group membership for the specific group is not reflected in IdP.	Group membership for the specific group is not reflected in IdP.

In other words, if Sync Groups on Login is off at the IdP level, then no groups are getting synced regardless of what the setting for Sync Membership is. But if Sync Groups on Login is turned on at the IDP level, then you have the option to override it for certain groups that you explicitly leave off.

Updating an identity provider

You can update the group synchronization option and the provider metadata in a CDP identity provider. To update an identity provider in CDP, you must be a CDP account administrator or have the PowerUser role.

About this task

You might want to update the CDP identity provider to change the group synchronization option or if you want to update the list of x.509 certificates in the provider metadata.

Required role: Account administrator or PowerUser

Procedure

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Identity Providers.
4. Find the CDP identity provider that you want to update.
5. Click the Actions button and select Update Identity Provider.

6. On the Identity Provider window:

- You can change the Sync Groups on Login option.
- You can add or edit the SAML Identity Provider Metadata.
- You can check the Generate workload username by email box to have the workload username is generated based on the email instead of the default. See [Generating workload usernames based on email](#).
- You can enable SCIM for Azure AD. See [Configure SCIM with Azure AD](#).
- You cannot change the name of the CDP identity provider.

7. Verify the updates and click Update.

CDP updates the information for the CDP identity provider.

Disabling the Cloudera SSO login

After you complete the identity federation setup between Cloudera and your enterprise IdP, you can disable the Cloudera SSO login option if you do not want to allow users in your organization to log in to CDP through the Cloudera registration and login page.

When you disable Cloudera SSO login for non-administrator users, CDP users must log in to CDP through the identity management system in your organization. Only the designated account administrator for your CDP subscription can log in to CDP through the Cloudera registration and login page.

Even after you disable the Cloudera SSO login, the designated account administrator can log in to CDP using their Cloudera SSO credentials. For added security, you can restrict all Cloudera SSO access (including the designated account administrator's access) by contacting Cloudera Support and they can disable or enable the "Cloudera SSO All Login Enabled" setting for your account. You can see "Cloudera SSO All Login Enabled" setting on the UI. When all Cloudera SSO logins are restricted, you will see this on the UI:

Cloudera SSO Login ⓘ Disabled. No users can login via Cloudera SSO, [contact support](#) to change.

Required role: Account administrator or PowerUser

Steps

1. Sign in to the CDP web interface.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Identity Providers.

The Identity Providers page shows the status of the Cloudera SSO Login option.

4. Click Disable to prevent users from logging in through the Cloudera registration and login page.

When the Cloudera SSO Login option is disabled, all CDP users except the CDP account administrators must log in through the identity management system in your organization. To log in to CDP, a user must be among the users included in the identity providers that you set up in CDP.

Configuring Azure Active Directory identity federation in CDP

You can onboard users by configuring Azure Active Directory (Azure AD) identity federation with CDP.

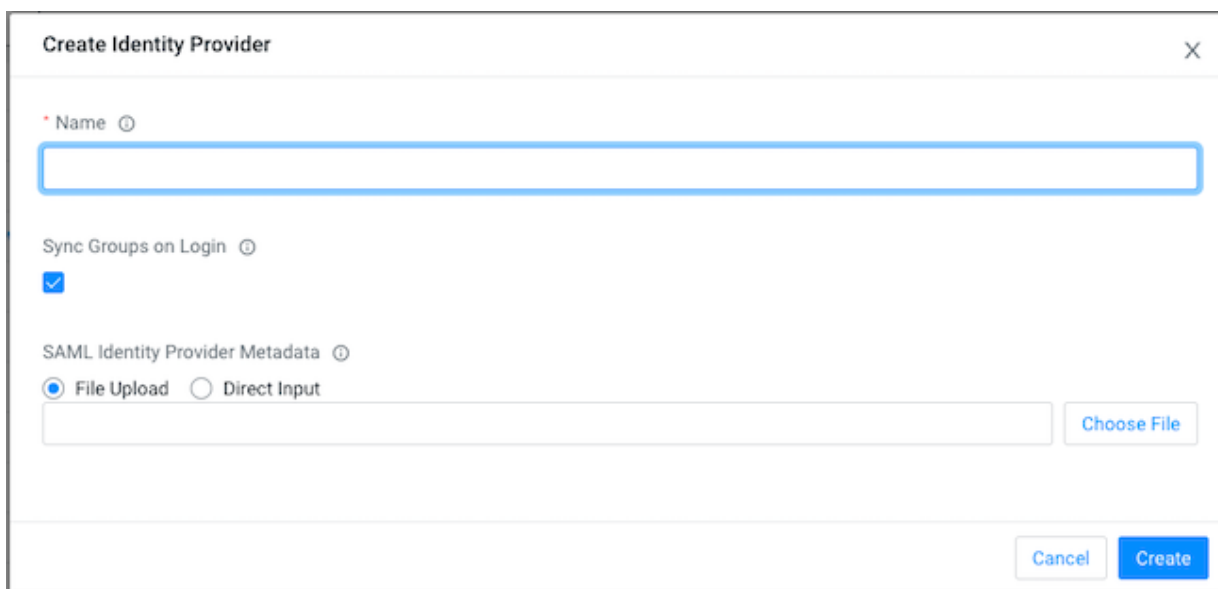
Before you begin

- CDP requires the Azure AD sAMAccountName attribute for the SAML group claim mapping. As per [Configure group claims for applications by using Azure Active Directory](#), sAMAccountName is not available with SAML on groups created in Azure AD; It is available only on groups created with on-premises AD and synced to Azure AD. If you only use Azure AD without on premises AD you can use SCIM to sync group memberships to CDP. See [Configure SCIM with Azure AD](#).
- In order to be compatible with CDP, your Azure AD should be configured to synchronize with on-premises Active Directory via Azure AD Connect sync tool. As per [Azure documentation](#), you must use Azure AD Connect 1.2.70.0 or newer.

Required role: PowerUser

Steps

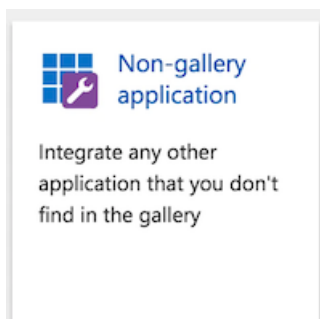
1. Log in to CDP web interface and navigate to Management Console > User Management, select the Identity Providers tab and click on Create Identity Provider to create an identity provider.



The screenshot shows a 'Create Identity Provider' dialog box with the following fields and options:


- Name**: A required text input field with a blue border.
- Sync Groups on Login**: A checkbox that is currently checked.
- SAML Identity Provider Metadata**: A section with two radio buttons: **File Upload** (selected) and **Direct Input**. Below these is a text input field and a **Choose File** button.
- Buttons**: **Cancel** and **Create** buttons at the bottom right.

2. Name your identity provider in CDP, for example MyCompany_AAD.
3. Click Create.
4. Click on MyCompany_AAD on CDP console and copy the CDP SAML Service Provider Metadata to an XML file (for example, saml-metadata.xml). You will need it later.
5. Open another web browser window, navigate to <https://portal.azure.com/>, and log in to your Azure Portal.
6. On your Azure Portal, navigate to the Azure Active Directory.
7. Select the Enterprise applications service.
8. Click on the +New application button.
9. Select the Non-gallery application.

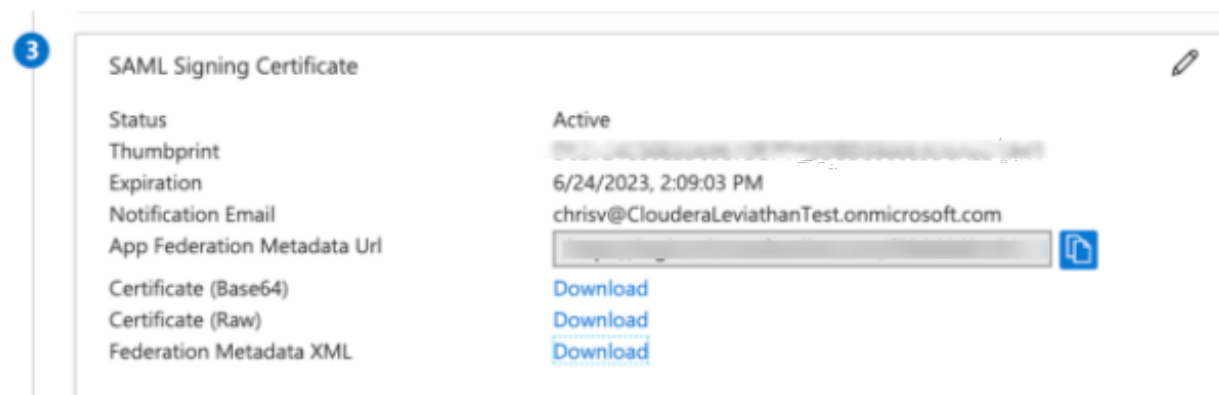


10. Give the application a name, for example. CDP_id_federation.
11. Click on the +Add button.
12. Once the application is added, go to 2. Set up single sign on .

13. Upload the metadata XML file that you saved in the earlier step.



14. Download the Federation Metadata xml for your Azure AD application and save it on your computer.



15. Switch back to CDP web interface and upload the metadata saved from AD and update the identity provider.

- a. Find the identity provider that you just created in CDP.
- b. Click the Actions button and select Update Identity Provider.
- c. On the Identity Provider window, upload the metadata XML file that you saved previously or copy and paste the content of that XML file:



- d. Verify the updates and click Update.

16. Switch back to Azure AD Azure Portal browser window.

17. Edit 1. Basic SAML Configuration:

- a. Make sure that the value for Identifier (Entity ID) is populated, for example “urn:cloudera:cdp:<Identity-Provider-Id>” or “urn:cloudera:altus” for legacy identity provider. Check the Service Provider Metadata to determine which identifier to use.
- b. From CDP SAML Service Provider Metadata you saved earlier, copy the AssertionConsumerService > Location value and paste it into the line Reply URL (Assertion Consumer Service URL).



Basic SAML Configuration		Edit
Identifier (Entity ID)	urn:cloudera:altus	
Reply URL (Assertion Consumer Service URL)	https://consoleauth.altus.cloudera.com/saml?samlProvide rid=f36cc318-60c8-474c-a534-b9289023bc48	
Relay State	Optional	
Logout Url	Optional	

18. Edit 2. User Attributes & Claims:

- a. If the customer is using on-prem Active Directory and Active Directory Connect to sync with Azure AD, you will be able to import Azure AD groups into CDP. Click +Add a group claim.
- b. On the Group Claims blade, do the following:
 1. Select Security groups or Groups assigned to the application.
 2. Select Source Attribute sAMAccountName.
 3. Check the Customize the name of the group claim. Optionally configure Group filtering for fine control of the list of groups that are included as part of the group claim. When a filter is configured, only groups that

match the filter will be included in the group's claim that is sent to CDP. For more information, see [Azure AD group filtering](#).

4. Enter “groups” in Name (Required).
5. Namespace enter `https://cdp.cloudera.com/SAML/Attributes`.

Group Claims



Manage the group claims used by Azure AD to populate SAML tokens issued to your app


Which groups associated with the user should be returned in the claim?

- ☐ None
- ☐ All groups
- ☒ Security groups
- ☐ Directory roles
- ☐ Groups assigned to the application

Source attribute *

sAMAccountName



 This source attribute only works for groups synchronized from an on-premises Active Directory using AAD Connect Sync 1.2.70.0 or above. [Learn More](#)

^ Advanced options

☐ Filter groups

Attribute to match

Match with

String


☒ Customize the name of the group claim

Name (required)

groups

Namespace (optional)

https://cdp.cloudera.com/SAML/Attributes

☐ Emit groups as role claims 

☐ Apply regex replace to groups claim content

6. Click Save.

c. For the rest of the claims, follow the instructions at [Configuring your enterprise IdP to work with CDP](#).

If your NameID is an opaque ID (such as a UUID), you can [Generate workload usernames based on email](#).

[Dashboard](#) > [CDP_id_federation](#) | [Single sign-on](#) > [SAML-based Sign-on](#) >

Attributes & Claims

[+ Add new claim](#) [+ Add a group claim](#) [Columns](#) [Got feedback?](#)

Required claim

Claim name	Value	
Unique User Identifier (Name ID)	user.objectid	***

Additional claims

Claim name	Value	
https://cdp.cloudera.com/SAML/Attributes/firstName	user.givenname	***
https://cdp.cloudera.com/SAML/Attributes/groups	user.groups	***
https://cdp.cloudera.com/SAML/Attributes/lastName	user.surname	***
mail	user.mail	***

Name	Namespace	Source	Source Attribute
Unique User Identifier (Name ID)		Attribute	user.objectid
firstName	https://cdp.cloudera.com/SAML/Attributes	Attribute	user.givenname
lastName	https://cdp.cloudera.com/SAML/Attributes	Attribute	user.surname
mail		Attribute	user.mail

19. Click on the SAML-based Sign-On on the top.

20. Test this application.

21. Once these steps are completed, a CDP user will be able to log in with their integrated Azure AD identity through their Office 365 applications page (Office.com). A new tile will appear for the CDP application created above:



Once a user signs in, the User and Groups will show up on the CDP Management Console's User Management screen. Navigate to your office 365 applications page, click on the new tile that was created for CDP, and verify that you are able to log in.

22. Once a user signs in, the User and Groups will show up on the CDP Management Console's User Management screen.

What to do next:

- Assign users to groups within your Azure AD that you will map to roles in CDP.
- Assign CDP roles to either the new users or the groups as appropriate.

Related Information

[Configure group claims for applications with Azure Active Directory](#)

[Configure SCIM with Azure AD](#)

Configure SCIM with Azure AD

CDP supports SCIM with Microsoft Azure Active Directory (Azure AD).

SCIM is a common way to get around:

- The Azure AD SAML 150 groups-per-claim limit.
- The Azure AD SAML sAMAccountName not available on groups created in Azure AD limitation (this is where group names are sent as their Object IDs in SAML instead of their human readable name).

For more information on these limitations, see [Configure group claims for applications by using Azure Active Directory](#).

Refer to this documentation if you would like to configure CDP to use SCIM. Prior to configuring CDP to use SCIM, you should be aware of the following limitations:

- You can only configure one identity provider per CDP account to use SCIM.
- Once you start using SCIM, you should not update users and groups in CDP, as they will get out of sync with Azure AD, and you may notice unexpected changes if/when Azure AD realizes the differences and attempts to re-sync the users/groups.
- Updating group names is not supported in CDP.
- Updating userName is not supported in CDP. You must use an Azure AD field that will not change as a userName to map via SCIM to CDP. If that field is an opaque ID (for example, a UUID) then you should generate workload usernames from email as described below.

Once you are aware of the limitations, you can proceed to configuring CDP to use SCIM. The steps include:

- [Prerequisites](#) on page 19
- [Enable SCIM for your identity provider in CDP](#) on page 19
- [Set up SCIM in Azure AD](#) on page 20

Prerequisites

Prior to configuring CDP to use SCIM, ensure that you can meet the following requirements:

- Prior to configuring CDP to use SCIM, you should [Configure Azure AD in CDP](#). SCIM is meant to be used in conjunction with SAML identity federation to synchronize users and groups from your identity provider to CDP.
- Setting up SCIM requires administrative operations in both CDP and Azure AD, and so it requires an Azure AD admin to perform the Azure AD steps.
- Additionally, you need to be a CDP account administrator or have the PowerUser role in CDP.


Enable SCIM for your identity provider in CDP


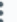
This task involves enabling SCIM for your identity provider and getting the SCIM URL that you will need to configure Azure AD so it can connect to CDP. Next, you create an access token for Azure AD to securely communicate with CDP for SCIM and get the access token secret you will need to configure Azure AD.

You will need to provide the lifetime for the access token used for Azure AD to communicate to CDP. Typical values are 1 year or 3 years; consult your security policies. You will need to rotate the access token before it expires.

In order to perform these steps, you need to be a CDP account administrator or have the PowerUser role in CDP.

Steps

1. Sign in to the CDP console.
2. Navigate to the Management Console.
3. Select User Management from the left pane.
4. Click on Identity Providers.
5. Select your identity provider, and from the  (context menu) select Update Identity Provider.

6. On the Update Identity Provider window, check the box for Enable SCIM. This will update the identity provider to be ready to accept SCIM API calls.
7. (Optional) If you use an opaque ID for SAML NameID and SCIM userName, on the Update Identity Provider window, check the box for Generate workload username by email. Check this box if you use an opaque ID for SAML NameID and SCIM userName. For more information, see [Generating workload usernames based on email](#).
8. Click Update.
9. Select your identity provider, and from the  (context menu) select View Identity Provider.
10. Copy the SCIM URL. You will need it later.
11. Click Close.
12. Select your identity provider and from the  (context menu) select Update SCIM Access Tokens.
13. On the Update SCIM Access Tokens window, click the button to Create SCIM Token.
14. Set your lifetime, in days. This is how long you will have before you need to rotate your SCIM access token. Note that you can always rotate your tokens earlier than their expiration date and you can revoke tokens at any time.
15. Click the Create button. This creates a SCIM access token that is used to authenticate SCIM API calls to this identity provider.
16. Your Access Token Secret will be shown. Copy it somewhere. You will need it later and it will not be shown again.

Set up SCIM in Azure AD

After enabling SCIM for your identity provider in CDP, you should set up SCIM in Azure AD. These steps must be performed by an Azure AD admin.

Steps

1. Sign in to the Azure Portal.
2. Navigate to Azure AD and then click on Enterprise Applications.
3. Select the application you used to configure identity federation to CDP. Or, if you are setting up SCIM in a new Azure AD app, create a new non-gallery enterprise application.
4. Under the Manage menu on the left, click Provisioning.
5. From the Provisioning Mode dropdown, select Automatic.
6. Expand the Admin Credentials section.
7. In the box for Tenant URL, paste in the SCIM URL that you saved earlier.
8. In the box for Secret Token, paste in the Access Token Secret that you saved earlier.
9. Click Test Connection and wait for success.
10. Click Save.
11. Still in the same Provisioning blade, expand the Mappings section. This section is grayed out until the connection has been successfully tested.
12. Click on Provision Azure Active Directory Groups.
13. Update the Attribute Mappings as follows:

Azure Active Directory Attribute	customappsso Attribute	Matching precedence
displayName	displayName	1

members	members	
---------	---------	--

As an outcome, your configuration should look similar to:

Attribute Mapping

Save

Discard

Name

Provision Azure Active Directory Groups

Enabled

Yes

No

Source Object

Group

Source Object Scope

All records

Source Object

umcietf:params:scim:schemas:core:2.0:Group

Target Object Actions

Create

Update

Delete

Attribute Mappings

Attribute mappings define how attributes are synchronized between Azure Active Directory and customappsso

Azure Active Directory Attribute	customappsso Attri...	Matching preceden...	Remove
displayName	displayName	1	Delete
members	members		Delete

Add New Mapping

Show advanced options



Note: Updating group names is not supported in CDP.

14. Click Save to save the group attribute mappings.
15. Close this blade to return to the Provisioning blade.
16. Click on Provision Azure Active Directory Users.
17. Update the Attribute Mappings as follows:

Azure Active Directory Attribute	customappsso Attribute	Matching precedence
objectId	userName	1
mail	emails[type eq "work"].value	
givenName	name.givenName	

surname	name.familyName	
---------	-----------------	--

As an outcome, your configuration should look similar to:

Attribute Mapping ✕

Save Discard

Name
Provision Azure Active Directory Users

Enabled
☒ Yes ☐ No

Source Object
User

Source Object Scope
[All records](#)

Source Object
urn:ietf:params:scim:schemas:extension:enterprise:2.0:User

Target Object Actions
☒ Create
☒ Update
☒ Delete

Attribute Mappings
Attribute mappings define how attributes are synchronized between Azure Active Directory and customappsso

Azure Active Directory Attribute	customappsso Attri...	Matching preceden...	Remove
objectId	userName	1	
mail	emails[type eq "work...		
givenName	name.givenName		
surname	name.familyName		

[Add New Mapping](#)

☐ Show advanced options



Note: The "customappsso Attribute" userName must be the same as your SAML NameID attribute. The value of the field mapped to NameID/userName is used to uniquely identify a user in CDP and must also be immutable, or you will run into known issues. See [Known issues and troubleshooting related to IdP setup in CDP](#).

18. Click Save to save the user attribute mappings.
19. Close this blade to return to the Provisioning blade.
20. Click Save in the Provisioning blade.
21. Still in the Provisioning blade, expand the Settings section.
22. Set Scope to Sync only assigned users and groups.
23. Click Save.
24. Now in the blade for the application, under the Manage menu click Users and groups.
25. If there are no users or groups, add a few. You will use this to test SCIM later.
26. Back in the blade for the application, under the Manage menu click Provisioning.
27. Click Start provisioning.
28. Wait for the sync cycle to run.

The initial provisioning cycle may take a few minutes to start but subsequent provisioning cycles run using a fixed cadence. You can see the cadence under the View provisioning details expandable section.



Note: Provisioning cycles are set by Azure AD and are not customizable.

29. Once the provisioning cycle is complete, check your users and groups in CDP. If there are any errors, contact your Cloudera representative.



Attention:

- CDP does not treat PATCH requests as atomic. PATCH operations are idempotent on CDP.
- Querying resources using HTTP POST is not supported.

Related Information

[Configuring Azure Active Directory identity federation in CDP](#)

[Known issues and troubleshooting related to IdP setup in CDP](#)

Importing or uploading users

You can bulk import users so as to assign policies to users and groups without requiring the users to log in at least once.

Before you begin

- Make sure the identity provider user ID matches with the NameID attribute value that is passed for the user in the SAML response using the associated SAML provider.
- Cloudera has default limits in place with regard to how many users, machine users, and groups can be added per account. Review [User and group limits](#) and make sure that you do not exceed these limits.

Required role: PowerUser

Steps

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. Click User Management in the left navigation panel.

The Users page displays the list of all CDP users.

4. Click Actions.
5. From the dropdown list that appears, click Upload Users.

The Upload Users page appears.

6. Select an identity provider from the dropdown list.
7. Select an option for adding user details:
 - File Upload - If you want to upload a CSV file with the details of the users that you like to add.
 - Direct Input - Enter the user details in the format specified. Make sure you add the header row as specified in the sample.
8. Click Next.
9. In the Preview Users screen that appears, verify if all the details are uploaded or added accurately.
10. Click Upload.
11. Confirmation of success is shown. Click Finish.

You can view newly added users from the Users tab.

Generating workload usernames based on email

CDP offers an option to generate workload usernames for CDP users based on user email addresses.

By default, workload usernames are generated using the identity provider user ID. For SAML logins that is the SAML NameID, for SCIM that is the SCIM userName, and when using the CDP APIs that is the identity-provider-user-id. Sometimes the identity provider user ID is an opaque ID, like a uuid or employee ID, which gives equally opaque workload usernames.

Alternatively, you can generate workload usernames based on users' email addresses instead of using the default workload usernames. For example, if your identity-provider-user-id is 8d16a2ea, and your email is bob@example.com, by default your workload username will be "8d16a2ea". If you choose to generate workload usernames by email, your workload username will instead be "bob".



Note: Once your CDP users are created and have the default or email-based workload usernames assigned, you cannot change the workload usernames already generated before the setting was changed. At that point, changing this setting does not change or regenerate any existing workload usernames. Specifically, for SAML logins users are created in CDP when they log in for the first time, and for SCIM users are created when the identity provider runs a SCIM sync cycle.

Steps

For CDP UI

When creating or updating an identity provider in CDP, you can check the Generate workload username by email box to have workload usernames generated based on email addresses.

For CDP CLI

From the CDP CLI, you can change how workload usernames are generated when you create (`iam create-saml-provider`) or update (`iam update-saml-provider`) a SAML provider by using the `--generate-workload-username-by-email` or `--no-generate-workload-username-by-email` flags. See:

```
iam create-saml-provider --help
iam update-saml-provider --help
```

Related Information

[Setting up the identity provider in CDP](#)

[Updating an identity provider](#)

Known issues and troubleshooting related to IdP setup in CDP

This topic covers known issues that you may encounter when setting up an identity provider in CDP and steps to troubleshoot them.

Known issues with mutable SAML NameID or SCIM userName

Issue (If using SAML without SCIM):

If you use SAML without SCIM, and you set the SAML NameID field to a mutable Azure AD field (such as an email), then you will end up with duplicate users in CDP when the Azure AD value changes. The duplicate user is a new, different user that has a different CDP workload username. This is because SAML has no way to differentiate between two different users and a user whose SAML NameID has changed.

Issue (If using SAML with SCIM):

If you are using both SAML and SCIM, you must set the SAML NameID and the SCIM userName fields to the same Azure AD field. If that Azure AD field is a mutable field, when it changes you will end up with duplicate users in CDP. This will cause errors in your Azure AD Enterprise Application and SCIM updates to the user will fail.

Workaround:

There is no automatic way to recover and you will likely have to delete and recreate the affected user. This includes having to delete and recreate the user's permissions, passwords, keys, and so on. If this happens, contact Cloudera support.

Known issues with Azure AD User Principal Name (UPN)

While uncommon in practice, the Azure AD UPN is actually a mutable field. Ask your Azure AD team if your organization mutates UPN before using it.

By default, when users are deleted in Azure AD they are moved to a "recently deleted" list for 30 days, before they are permanently deleted and removed from Azure AD. When a user is moved to the "recently deleted" list, their UPN is automatically changed by Azure AD. This will cause SCIM errors to show up in your Azure AD Enterprise

Application: The error is Azure AD trying to update the userName field, and CDP returning an error saying that operation is not supported. These errors can be ignored.

Note that even though users in Azure AD are in the "recently deleted" list, they will persist in CDP as active users until after the 30 day wait time when they are fully removed from Azure AD (deleted from the Azure AD system). At that point SCIM will delete them from CDP as well. If you permanently delete the user from Azure AD before the 30 day period (that is, delete the user from the "recently deleted" list), then that user will also be deleted in CDP. If neither of these options work for you, then you must delete the user manually in CDP.

Known issues with updating group names

Updating group names is rare - most organizations do not do this. CDP does not support updating group names. To fix the issue, create a new group instead.

Understanding CDP user accounts

User accounts identify the users who can access services, applications, and components in the Cloudera Data Platform.

Roles assigned to a user account determine the actions that the user can perform in CDP.

There are four types of user accounts in CDP:

- CDP account administrator
- CDP user
- CDP workload user
- CDP machine user

CDP account administrator

During the initial setup of the CDP subscription for a customer, Cloudera designates a user account as a CDP account administrator.

A CDP account administrator has administrator privileges in CDP. The CDP account administrator user account cannot be managed within CDP. You must contact Cloudera support to add or remove an account administrator from your CDP account.

As an account administrator, you have all the privileges in CDP and can perform any task in CDP. You can set up users and assign roles, services, and environments to users in CDP according to the tasks that they need to perform. You can set up another user as a CDP administrator by assigning the PowerUser role to the user. However, you cannot set up another user as a CDP account administrator.

A CDP account administrator requires a Cloudera user account. To be designated as a CDP account administrator, you must register for a Cloudera user account. To register for a Cloudera user account, go to the Cloudera Account Registration page and create an account.

CDP user

To perform tasks using CDP and its services, you must be a CDP user and roles and resources need to be assigned to this user.

CDP allows users within your organization to log in to CDP through the authentication system in your organization without registering with Cloudera or creating a Cloudera account. During the initial process of configuring the environment, the account administrator must set up identity federation and thus automatically add users.

When a CDP user who is not an account administrator logs in to CDP for the first time, the user has limited privileges. A CDP administrator must assign the appropriate roles to the user after the initial user login.

The CDP account administrator can delete the user accounts. Deleting a user removes all access keys and SSH keys associated with the user, and unassigns all roles and resource roles assigned to the user. The user is also removed from all groups that they belong to.

A user who has a valid account in CDP but is not assigned any role can perform a limited number of tasks. A user who logs in to the CDP console without an assigned role or environment can perform only the following tasks:

- Download the CDP client.
- View the CDP documentation.

CDP workload user

You need the workload username to access Data Hub clusters and non-SSO interfaces, and to SSH to clusters.

CDP assigns workload username to all CDP users when the user is created in CDP. The workload username is generated from either the identity provider user ID (default) or the email (configurable via [Generate workload usernames](#)). If a workload user name already exists in CDP, which can occur when multiple identity providers are mapped, the workload username will contain a numeric suffix. Workload usernames are immutable - once set they never change.

To identify your workload username, navigate to Home > in the bottom left corner click on your user name > Profile > Workload User Name.

Alternatively, you can obtain your workload username by following these steps:

1. Sign in to the CDP web interface.
2. From the CDP home page, click Management Console.
3. Click User Management in the left navigation panel. The Users page displays the list of all CDP users.
4. Search for your username.
5. In the list of users that appears, you can see the Workload User Name column. Your workload username will appear under this column.

You can set your workload password to get access to non-SSO interfaces.

CDP machine user

A machine user account provides programmatic access to CDP. Create a machine user account if you have an application that needs to access the CDP services with the CLI or the CDP SDK for Java.

You can define the machine user account in your application to create and manage clusters and run jobs in CDP using the CLI or API commands.

You can create and manage a machine user account within CDP. You must assign an API access key to a machine user account to enable it to access the CDP service with the CDP CLI or CDP SDK. You must assign roles to a machine user account to authorize it to perform tasks in CDP.

A machine user account does not have an associated Cloudera user account. You cannot use a machine user to log in to the CDP console.

Use the following guidelines when you manage user accounts in CDP:

- When you create a machine user account, you assign roles and environments to the machine user account in the same way that you assign roles and environments to other user accounts.
- You can revoke permissions for a CDP machine user account by removing the machine user from groups or unassigning account level or resource roles.
- Deleting a machine user removes all access keys and SSH keys associated with the machine user, and unassigns all roles and resource roles assigned to the machine user. The machine user is also removed from all groups that they belong to.

Understanding account roles and resource roles

To access resources and perform tasks in CDP, each user requires permissions. As a CDP administrator, you can assign a role to a user or a machine user to give the user permission to perform the tasks either on the whole account or on a specific resource.

Each role has an attached policy that defines the permissions associated with the role. The policy attached to a role determines the operations that the role allows the user to perform. When users attempt to perform operations that are not permitted in their assigned role, they get a permission denied error message.

CDP has predefined roles for your use. You can assign a role or a combination of roles to give the user the appropriate permissions to complete tasks in CDP. You cannot modify the predefined CDP roles or the policies associated with the predefined roles.

The scope of predefined roles and resource roles can vary. For example, a role might grant view access only to Data Hub clusters but not to environments in which these clusters are running. You might need to assign multiple roles to ensure that a user can perform all required tasks in CDP.

CDP provides the following types of roles:

- Account roles - An account role grants a user, machine user, or group permissions to access or perform tasks on all resources within the CDP tenant.
- Resource roles - A resource role grants a user, machine user, or group permissions to access or perform tasks on a specific resource (such as a specific environment or a specific Data Hub cluster).
- Group membership administration roles - The `IamGroupAdmin` role can be assigned to a user to manage group membership for a specific group.

Review the following documentation to learn more about these role types:

Account roles

A CDP role grants permissions to perform tasks in CDP that are not associated with a specific resource. You explicitly assign a role to a user, machine user, or a group.

When assigning roles to users and groups, consider the following:

- Only `PowerUser` can assign account roles.
- A user needs the following two types of roles in order to assign access to resources to other users:
 - One of the roles that allow role assignment: `EnvironmentCreator`, `EnvironmentAdmin`, `DataSteward`, `DatahubAdmin` or another admin role for a CDP service.
 - One of the roles that allow listing users within the organization: `IamUser` or `IamViewer`.
- All users who need to access CDP CLI need the `IamUser` role.

Each role is identified by a CRN, which uses the following format:

```
crn:altus:iam:<CONTROL_PLANE_REGION>:altus:role:<ROLE_NAME>
```

For example, the following is the `IamViewer` role CRN:

```
crn:altus:iam:us-west-1:altus:role:IamViewer
```

You can view all available roles and their CRNs by using the `cdp iam list-roles` command.

Account roles can be assigned from the Management Console > User Management > Roles tab or from CDP CLI by using the `cdp iam assign-user-role` or `cdp iam assign-group-role` commands.

The predefined account roles available in CDP are as follows:

Table 1: Account roles

Account role	Description	Important considerations
PowerUser	Grants permission to perform all tasks on all resources.	Unlike other users (who only see the resources that they are authorized to list), Power Users can list all resources. By default, Power Users don't not have full access to all resources but can assign themselves a resource role that grants them access to these resources.
EnvironmentCreator	Grants permission to create environments and shared resources (cluster templates, recipes, image catalogs, credentials, proxies), and sync users.	Since shared resources are managed separately from environments, in order for a user with the EnvironmentCreator role to be able to use a provisioning credential for creating an environment, that user needs to be Owner or SharedResourceUser for that credential. EnvironmentCreator is the only role that allows you to manage access to proxies that have been registered in CDP.
IamUser	Grants permission to create access keys and upload SSH keys for the user (but not for other users). Moreover, this role includes all permissions of the IamViewer role. It grants permission to view all users in the account and their assigned roles and access keys.	Either the IamUser or IamViewer role is required to list other users, therefore any user who needs to assign roles, such as EnvironmentCreator, EnvironmentAdmin, DataHubAdmin, and so on, should be assigned either IamUser or IamViewer.
IamViewer	Grants permission to view all users in the account and their assigned roles and access keys.	Either the IamUser or IamViewer role is required to list other users, therefore any user who needs to assign roles, such as EnvironmentCreator, EnvironmentAdmin, DataHubAdmin, and so on, should be assigned either IamUser or IamViewer.
ClassicClustersCreator	This role is required to register a new cluster. If this role is not present then the "Add Cluster" button is not visible to the user.	
DataCatalogCspRuleManager	Grants permission to perform all tasks on CSP rules in Data Catalog.	
DataCatalogCspRuleViewer	Grants permission to list and view CSP rules in Data Catalog.	
DFCatalogAdmin	Grants permission to perform all tasks on objects stored in the DataFlow Catalog. This includes importing and deleting flow definitions, as well as uploading new versions of existing flow definitions.	
DFCatalogViewer	Grants permission to browse the DataFlow Catalog and view flow definitions.	
BillingAdmin	Grants permission to monitor CDP credit consumption.	

Related Information[Assigning account roles to users](#)[Assigning account roles to groups](#)**Resource roles**

A role that is associated with a specific resource is called a resource role. This type of role gives permission to perform tasks on a specific resource, such as a specific CDP environment, shared resource, or Data Hub cluster.

When you assign a resource role, you must specify the resource on which to grant the resource role permissions. For example, you can assign a user a resource role that grants permission on an environment. The user assigned the resource role can list, access, and perform tasks only on that environment, but not on other environments.

A resource role determines a specific set of tasks that the user can perform on the resources. For example, the EnvironmentUser resource role assigned to a user allows the user the rights contained in the resource role only on that particular environment.

The predefined resource roles available in CDP that you can assign to CDP users, machine users, and groups are as follows:

- Environment resource roles
- Shared resource resource roles
- Data Hub resource roles
- Classic cluster resource roles
- The Owner resource role (available on all resources)

Each role is identified by a CRN, which uses the following format:

```
crn:altus:iam:<CONTROL_PLANE_REGION>:altus:resourceRole:<RESOURCE_ROLE_NAME>
```

For example, the following is the DataHubAdmin role CRN:

```
crn:altus:iam:us-west-1:altus:resourceRole:DataHubAdmin
```

You can view all available roles and their CRNs by using the `cdp iam list-resource-roles` command.

Learn more about different resource role types:

- [Environment resource roles](#) on page 29
- [Shared resource resource roles](#) on page 32
- [Data Hub resource roles](#) on page 32
- [Classic cluster resource roles](#) on page 33
- [The Owner resource role](#) on page 33

Environment resource roles

Environment resource roles can be assigned on the scope of a specific environment.

These resource roles can be assigned from the Management Console > Environments > navigate to a specific environment > Actions > Manage Access > Access or from CDP CLI using the `cdp iam assign-user-resource-role` command.

Table 2: Environment resource roles

Resource role	Description	Important considerations
EnvironmentAdmin	Grants all rights to the environment and Data Hub clusters running in it, except the ability to delete the environment.	<p>The user who created the environment automatically gets the EnvironmentAdmin role on the scope of that environment.</p> <p>The EnvironmentAdmin resource role is assigned the Limited Cluster Administrator role in Cloudera Manager. Users with this role can manage the cluster lifecycle, change configurations, and manage parcels. For more information on CM roles, see the topic Default User Roles.</p> <p>The Cloudera Manager Limited Cluster Administrator role is assigned to the EnvironmentAdmin because the CDP Control Plane is responsible for certain tasks historically done in Cloudera Manager, for example: adding or removing hosts as part of up/down-scaling and repair operations, executing upgrades of clusters in coordination with upgrading the OS images used by clusters, and creating new clusters based on templates preconfigured to work in a CDP environment. In addition, only selected services and workload types are currently supported in Data Hub, represented by the built-in cluster definitions. Finally, certain CDP services like encryption-at-rest infrastructure are explicitly not designed for use in the public cloud, where the cloud provider's object store encryption capabilities should be used. Because of this, Data Hub is prescriptive in its choice of workload types and the CDP Control Plane is best suited to manage most cluster life cycle operations. Doing so directly in Cloudera Manager could lead to unexpected operational issues. Data Hub does, however, support fully customizable cluster templates.</p> <p>EnvironmentAdmin can manage access to the environment by assigning a user EnvironmentAdmin, DataSteward, or EnvironmentUser role.</p>
EnvironmentUser	<p>Grants permission to view Data Hub clusters and set the workload password for the environment.</p> <p>The EnvironmentUser resource role is assigned the Read-Only role in Cloudera Manager. For more information on CM roles, see the topic Default User Roles.</p>	This role should be used in conjunction with service-specific roles such as DataHubAdmin, DWAdmin, DWUser, MLAdmin, MLUser, and so on. When assigning one of these service-specific roles to users, make sure to also assign the EnvironmentUser role.
DataSteward	Grants permission to perform user/group management functions in Ranger and Atlas Admin, manage ID Broker mappings, and start user sync for the environment.	DataSteward can manage access to the environment by assigning a user DataSteward or EnvironmentUser role.
DataHubCreator	Grants permission to create Data Hub clusters in the environment.	
DEAdmin	Grants permission to create, delete and administer Data Engineering services for the environment.	When assigning this role, you should also assign the EnvironmentUser role.
DEUser	Grants permission to list and use Data Engineering services for the environment.	When assigning this role, you should also assign the EnvironmentUser role.

Resource role	Description	Important considerations
DFAdmin	Grants permission to enable, disable and administer the CDP environment for DataFlow. This includes granting and revoking the ability to access the DataFlow Kubernetes API server.	When assigning this role, you should also assign the EnvironmentUser role.
DFFlowAdmin	Grants permission to create, terminate, administer and monitor running deployments for the environment.	When assigning this role, you should also assign the EnvironmentUser role.
DFFlowDeveloper	Grants permission to view, create, modify, or delete flow drafts; start and end test sessions in an environment.	When assigning this role, you should also assign the EnvironmentUser role.
DFFlowUser	Grants permission to view and monitor deployments for the environment.	When assigning this role, you should also assign the EnvironmentUser role.
DFProjectCreator	Grants permission to create a DataFlow Project within a given CDP environment.	When assigning this role, you should also assign the EnvironmentUser role.
DWAdmin	Grants permission to activate/terminate or launch/stop/update services in Database Catalogs and Virtual Warehouses.	When assigning this role, you should also assign the EnvironmentUser role.
DWUser	Grants permission to view and use Cloudera Data Warehouse clusters within the environment.	When assigning this role, you should also assign the EnvironmentUser role.
MLAdmin	Grants permission to create and delete Cloudera Machine Learning workspaces within the environment. MLAdmins will also have Site Administrator access to all the workspaces provisioned within this environment. They can run workloads, monitor, and manage all user activity on these workspaces.	When assigning this role, you should also assign the EnvironmentUser role.
MLBusinessUser	Grants permission to view Cloudera Machine Learning workspaces for the environment. MLBusinessUsers are granted view-only access to applications that have been shared with them through projects inside a workspace.	When assigning this role, you should also assign the EnvironmentUser role.
MLUser	Grants permission to view Cloudera Machine Learning workspaces provisioned within the environment. MLUsers are also able to run workloads on all the workspaces provisioned within this environment.	When assigning this role, you should also assign the EnvironmentUser role. MLUsers currently require the SharedResourceUser role on the cloud credential used for the environment.
ODAdmin	Grants permission to create, drop and administer the Cloudera Operational Databases for the environment.	When assigning this role, also assign the DataSteward or EnvironmentAdmin role.
ODUser	Grants permission to list and use Cloudera Operational Databases for the environment.	
Owner	Grants all permissions required to manage the environment in CDP including the ability to delete it.	The user who created the environment automatically gets the Owner role on the scope of that environment. The Owner role on the scope of an environment allows you to delete that environment, but to access the environment's clusters (Data Lakes, Data Hubs), you need EnvironmentAdmin or EnvironmentUser.

Shared resource resource roles

Shared resources resource roles can be assigned on the scope of a specific shared resource such as a credential, cluster template, image catalog, proxy, or recipe. This does not include default shared resources (such as default cluster templates), which can be seen by everyone who is able to access the account.

These resource roles can be assigned from the Management Console > Environments > Shared Resources > select a shared resource > navigate to a specific shared resource > Manage Access, or from CDP CLI using the `cdp iam assign-user-resource-role` command.

You can view all available resource roles and their CRNs by using the `cdp iam list-resource-roles` command.

Table 3: Shared resource resource roles

Resource role	Description	Important considerations
SharedResourceUser	This role enables shared resource sharing with other users. It grants permission to access and use the specific shared resource such as a specific cluster template, credential, image catalog, proxy, or recipe.	In order for a user to be able to use a provisioning credential for creating an environment, that user needs to be Owner or SharedResourceUser for that credential.
Owner	Grants all permissions required to manage the shared resource in CDP including the ability to delete it.	The user who created the shared resource automatically gets the Owner role on the scope of that shared resource. In order for a user to be able to use a provisioning credential for creating an environment, that user needs to be Owner or SharedResourceUser for that credential.

Data Hub resource roles

Data Hub resource roles can be assigned on the scope of a specific Data hub cluster.



Note: While full access to manage a Data Hub via the Management Console can be granted via assigning the Owner role on the scope of the Data Hub, access to the underlying cluster can only be granted by assigning the EnvironmentAdmin or EnvironmentUser role on the scope of the environment where the Data Hub is running.

These resource roles can be assigned from the Management Console > Data Hub clusters > click on a cluster > Actions > Manage Access, or from CDP CLI using the `cdp iam assign-user-resource-role` command.

You can view all available roles and their CRNs by using the `cdp iam list-resource-roles` command.

Table 4: Data Hub resource roles

Resource role	Description	Important considerations
DataHubAdmin (Technical Preview)	Grants administrative rights over the Data Hub cluster, such as start, stop, scale, repair and grant or revoke access.	When assigning this role, you should also assign the EnvironmentUser role. Granting DataHubAdmin role does not grant Cloudera Manager admin rights or Runtime service admin rights (for example NiFi Registry Admin).

Resource role	Description	Important considerations
Owner	Grants all permissions required to manage the Data Hub in CDP including the ability to delete it.	<p>The user who created the Data Hub automatically gets the Owner role on the scope of that Data Hub.</p> <p>The Owner role does not grant any cluster-level permissions such as the ability to access or manage a cluster via Cloudera Manager. In order to access Data Hubs running within an environment, you need EnvironmentAdmin or EnvironmentUser on the scope of that environment.</p>

Classic cluster resource roles

Classic cluster resource roles can be assigned on the scope of a specific classic cluster.

These resource roles can be assigned from the Management Console > Classic clusters > context menu > Manage Access, or from CDP CLI using the `cdp iam assign-user-resource-role` command.

You can view all available roles and their CRNs by using the `cdp iam list-resource-roles` command.

Table 5: Classic cluster resource roles

Resource role	Description	Important considerations
ClassicClusterAdmin	<p>Grants permission to perform any operation on the cluster, except deleting it.</p> <p>Grants permission to assign access to the cluster to other users.</p>	
ClassicClusterUser	Grants permission to access details of the cluster.	
Owner	Grants all permissions required to manage the classic cluster in CDP including the ability to delete it.	<p>The user who created the classic cluster automatically gets the Owner role on the scope of that classic cluster.</p> <p>The Owner role does not grant any cluster-level permissions such as the ability to access or manage a cluster via Cloudera Manager.</p>

The Owner resource role

In addition to the aforementioned resource roles, CDP includes the Owner resource role.

The Owner role:

- Grants full permissions on a specific resource in the Management Console, including the ability to delete the resource. It does not grant any cluster-level permissions such as the ability to access or manage a cluster via Cloudera Manager.
- Is assigned automatically on a resource to the user who created the resource. For example, if a user creates an environment called “test”, the user is assigned the Owner role for that environment.
- Allows a user to grant a set of rights (including the Owner role) on the resource to other users and groups. This is possible only if the user also has the `IamUser` or `IamViewer` role allowing to list users within the organization.
- Can be assigned at the scope of the following resources: an environment, Data Lake, shared resource (cluster template, recipe, image catalog, credential, proxy), Data Hub cluster, or classic cluster
- Can be assigned using the same steps as other resource roles.

Related Information

[Assigning resource roles to users](#)

[Assigning resource roles to groups](#)

[Default User Roles \(Cloudera Manager\)](#)

Group membership administration roles

The `IamGroupAdmin` role can be assigned to a user or a group on the scope of a group to allow them to manage membership of that group.

Note that:

- The `IamGroupAdmin` role grants a user or a group the permission to add users to or remove users from a group. The role does not grant permission to manage roles and resources for the group.
- In order for a user with the `IamGroupAdmin` to add or remove users from a group, the user must also have the `IamUser` or `IamViewer` role that allows listing IAM users and groups within the organization.

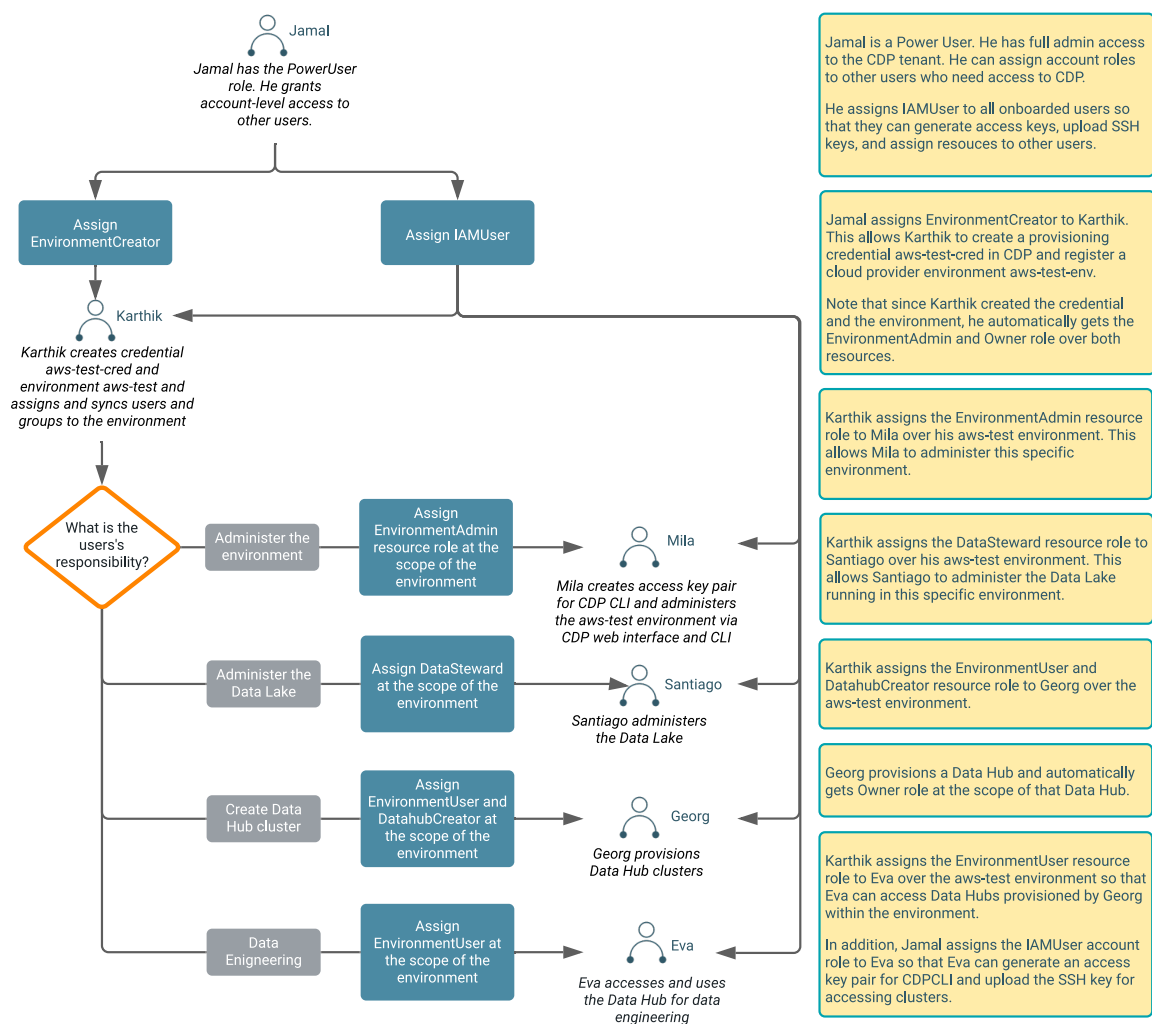
Related Information

[Assigning a group membership administrator](#)

Example role assignment scenario

This section outlines an example role assignment scenario.

The following illustration presents an example scenario where account and resources roles are assigned to multiple users within an organization:



User and group limits

Cloudera has default limits in place with regard to how many users, machine users, and groups can be added per account. Customers can contact Cloudera Support to increase these limits.

The following table describes the default limit (second column) and the maximum limit that can be requested by contacting Cloudera Support (third column):

Limit name	Default limit	Maximum limit that can be requested
Maximum number of users	1,000 users	10,000 users Note: The user limit and machine user limit must total 10,000 or less. Approved limit increases may require upgrades of one or more services.
Maximum number of groups	50 groups	3,000 groups
Maximum number of machine users	100 machine users	500 machine users Note: The user limit and machine user limit must total 10,000 or less.
Maximum number of access keys	200 access keys	2,000 access keys
Maximum number of access tokens	20 access tokens	200 access tokens

Managing users and machine users in CDP

A PowerUser can manage CDP users on the CDP web interface or via CDP CLI.

Creating a machine user in CDP

You can create a machine user for programmatic access to CDP

Before you begin

Machine user names cannot start with a double underscore ("__").

Required role: PowerUser

Steps

For CDP UI

1. Navigate to the Management Console > User Management > Users.
2. From the Actions menu, select Create Machine User.
3. Provide a name and click on Create.

For CDP CLI

Use the following command:

```
cdp iam create-machine-user \
--machine-user-name MACHINE_USER_NAME
```

Next, generate an API access key for your machine user by using the following command:

```
cdp iam create-machine-user-access-key \  
  --machine-user-name MACHINE_USER_NAME
```

What to do next

- You need to perform user sync for the change to take effect. See [Performing user sync](#).
- To generate an API access key for your machine user, see [Generating an API access key](#).
- You should grant the newly created machine user access to resources. To do this, follow the steps for [Assigning resources to users](#).



Note: When searching for the machine user in the "Select group or user" text box, make sure to enter machine user name, not workload user name.

Deleting users and machine users


CDP administrators have the ability to delete users and machine users in CDP through both the CDP user interface and the CDP CLI.

If a CDP user or machine user is deleted from an integrated identity provider system (for example, if a user leaves the company), the user is not automatically deleted in CDP. CDP administrators have the ability to delete a user in CDP through both the user interface and the CLI.

Required role: PowerUser

Steps

For CDP UI

1. From CDP user interface, navigate to the Management Console > User Management.
2. Search for the user or machine user that you want to delete and click the  (context menu) at the end of the user entry row.
3. Perform one of the following:
 - Click Delete User (for a user) or Delete Machine User (for a machine user) and then OK on the confirmation screen.
 - Alternatively, you can click on the user or machine user name to enter the user's detail page. From there, click Actions > Delete User (for a user) or Delete Machine User (for a machine user).

For CDP CLI

In the CDP CLI, the command to delete a user is `delete-user` and the command to delete a machine user is `delete-machine-user`.

Run the `delete-user` command as shown in the following example:

```
cdp iam delete-user --user-id <value>
```

The `delete-user` command requires a `user-id` value, which can be either the user ID or the CRN (Cloudera resource name) of the given user. You can obtain the user ID from the web interface from the user's details, or from the CDP CLI from the output of the `cdp iam list-users` command.

Run the `delete-machine-user` command as shown in the following example:

```
cdp iam delete-machine-user --machine-user-name <value>
```

The `delete-machine-user` command requires a `machine-user-name` value, which can be either the user ID or the CRN (Cloudera resource name) of the given machine user. You can obtain the user ID from the web interface from the machine user's details, or from the CDP CLI from the output of the `cdp iam list-machine-users` command.

For a detailed description of the command properties, call the CDP help for the command:

```
cdp iam delete-user --help
cdp iam delete-machine-user --help
```

What to do next

Deleting a user or machine user removes all access keys and SSH keys associated with the user, and unassigns all roles and resource roles assigned to the user. The user is also removed from all groups that they belong to. Once the deletion process is completed and synced, the user will not be able to use any access keys to access CDP.

You must trigger user sync to ensure that the deleted user loses access to all environments. See [Performing user sync](#). Only once the user sync is complete, the deleted user loses access to CDP.

It takes around 2 minutes to fully delete a user or machine user in CDP. During this time you will not be able to recreate the user (that is, for 2 minutes you will not be able to create a user in the same Identity Provider with the same NameID), but you can proceed to trigger user sync right away.

Assigning account roles to users

Assign account roles to a CDP user to manage the tasks that the user can perform in CDP. You can assign multiple roles to users or machine users to provide them with the permissions they need to perform their required tasks.

Required role: PowerUser

Steps

For CDP UI

1. Sign in to CDP.
2. From the CDP home page, click Management Console.
3. Click User Management.
The Users page displays the list of all CDP users.
4. Click the name of the user to whom you want to assign a role.
The user details page displays information about the user.
5. Click the Roles tab.
6. Click Update Roles.
7. On the Update Roles window, select the roles you want to assign to the user.
To remove a role from the user account, clear the selected role.
8. Click Update.
The roles that you select displays in the list of roles assigned to the user.
To remove a role from a user account, click check box next to the assigned role that you want to remove. Click Update to confirm that you want to revoke the role permissions.

For CDP CLI

You can use the following command to assign a role to a user or a machine user:

```
cdp iam assign-user-role \
--user-name <value> \
--role <value>
```

To remove a role from a user or a machine user:

```
cdp iam unassign-user-role \  
--user-name <value> \  
--role <value>
```

```
cdp iam unassign-machine-user-role \  
--machine-user-name <value> \  
--role <value>
```

The `--role` parameter requires the CRN of the CDP role. You can use the `cdp iam list-roles` command to list resource roles with role CRNs.

To get a list of the roles assigned to a group:

```
cdp iam list-user-assigned-roles \  
--user-name <value>
```

```
cdp iam list-machine-user-assigned-roles \  
--machine-user-name <value>
```

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Related Information

[Account roles](#)

Assigning resource roles to users

To grant a user or a machine user access to a resource (such as an environment, a shared resource, or a Data Hub cluster), assign a resource role to the user on the scope of that resource or, in some cases (Data Hub clusters), on the scope of the parent resource.

In general, resource roles can be assigned from CDP user interface using the Manage Access option available from the resource details page or from CDP CLI using the `cdp iam assign-user-resource-role` or `cdp iam assign-machine-user-resource-role` commands. For detailed instructions, see the following sections:

Assign an environment resource role to a user

To assign an environment to a user or a machine user, assign a specific resource role on the scope of the specific environment.

Required roles:

- Owner or a role that allows administering the environment AND
- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. Navigate to the Environments page.
4. In the list of environments that appear, select an environment by clicking on it.
5. From the Actions menu select Manage Access.

6. In the Access tab, enter the name of the user in the text box.
7. In the Update Resource Roles window, select the required resource role.
8. Click Update Roles.

For CDP CLI

Use the following commands to assign a resource to a user or a machine user:

```
cdp iam assign-user-resource-role \  
--user-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

```
cdp iam assign-machine-user-resource-role \  
--machine-user-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

To remove a resource role from a user or a machine user:

```
cdp iam unassign-user-resource-role \  
--user-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

```
cdp iam unassign-machine-user-resource-role \  
--machine-user-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

- The resource-role-crn parameter requires the CRN of the resource role you want to assign to the user. You can use the `cdp iam list-resource-roles` command to list resource roles with role CRNs.
- The resource-crn parameter requires the CRN of the resource on which you want to grant the resource role permissions. You can obtain it from the details of the resource.

To get a list of the resource roles assigned to a user or a machine user:

```
cdp iam list-user-assigned-resource-role \  
--user-name <value>
```

```
cdp iam list-machine-user-assigned-resource-role \  
--machine-user-name <value>
```

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Related Information

[Resource roles](#)

Assign a shared resource resource role to a user

You can assign shared resources such as credentials, clusters templates, recipes, image catalogs, or proxies to users and machine users. To assign a shared resource to a user or a machine user, assign a specific resource role on the scope of the specific shared resource.

Required roles:

- Owner or a role that allows administering the environment AND

- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. From the left pane, select Shared Resources, and then select a resource type (for example Cluster Templates) to view a summary of all resources of that type.
4. Find the specific resource (for example a specific cluster template) and click on it to navigate to its details page.
5. Click on Manage Access.
6. Enter the name of the user in the text box.
7. In the Update Resource Roles window, select the required resource role.
8. Click Update Roles.

For CDP CLI

Use the following commands to assign a resource to a user or a machine user:

```
cdp iam assign-user-resource-role \
--user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

```
cdp iam assign-machine-user-resource-role \
--machine-user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

To remove a resource role from a user or a machine user:

```
cdp iam unassign-user-resource-role \
--user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

```
cdp iam unassign-machine-user-resource-role \
--machine-user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

- The `resource-role-crn` parameter requires the CRN of the resource role you want to assign to the user. You can use the `cdp iam list-resource-roles` command to list resource roles with role CRNs.
- The `resource-crn` parameter requires the CRN of the resource on which you want to grant the resource role permissions. You can obtain it from the details of the resource.

To get a list of the resource roles assigned to a user or a machine user:

```
cdp iam list-user-assigned-resource-role \
--user-name <value>
```

```
cdp iam list-machine-user-assigned-resource-role \
--machine-user-name <value>
```


Related Information

Resource roles

Assign a Data Hub resource role to a user

You can assign a Data Hub resource role to a user or a machine user to allow them to manage a specific Data Hub.



Note: While full access to manage a Data Hub via the Management Console can be granted via assigning the Owner role on the scope of the Data Hub, access to the underlying cluster can only be granted by assigning the EnvironmentUser role on the scope of the environment where the Data Hub is running.

Required roles:

- Owner or a role that allows administering the environment AND
- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. Sign in to the CDP console.
2. Navigate to the details page of your Data Hub cluster. This can be done in a few ways. For example:
 - a. From the CDP home page, click Data Hub Clusters and then click on the specific cluster.
 - b. From the CDP home page, click on Management Console, navigate to the Data Hub Clusters page, and then click on the specific cluster.
3. From the Actions menu select Manage Access.
4. Enter the name of the user in the text box.
5. In the Update Resource Roles window, select the required resource role.
6. Click Update Roles.

For CDP CLI

Use the following commands to assign a resource to a user or a machine user:

```
cdp iam assign-user-resource-role \
--user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

```
cdp iam assign-machine-user-resource-role \
--machine-user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

To remove a resource role from a user or a machine user:

```
cdp iam unassign-user-resource-role \
--user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

```
cdp iam unassign-machine-user-resource-role \
--machine-user-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

- The `resource-role-crn` parameter requires the CRN of the resource role you want to assign to the user. You can use the `cdp iam list-resource-roles` command to list resource roles with role CRNs.

- The resource-crn parameter requires the CRN of the resource on which you want to grant the resource role permissions. You can obtain it from the details of the resource.

To get a list of the resource roles assigned to a user or a machine user:

```
cdp iam list-user-assigned-resource-role \  
--user-name <value>
```

```
cdp iam list-machine-user-assigned-resource-role \  
--machine-user-name <value>
```

Related Information

[Resource roles](#)

Assign a classic cluster resource role to a user


You can assign a specific resource role to a user or a machine user on the scope of a specific classic cluster to allow them to manage a specific classic cluster.

Required roles:

- Owner or a role that allows administering the environment AND
- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. In the Management Console navigate to the Classic Clusters dashboard.
2. Click on the  (context menu) next to the cluster that you want to update and select Manage Access.
3. Find the user that you want to update and click on Update Roles.
4. Select or deselect the roles and then click on Update Roles.

For CDP CLI

Use the following commands to assign a resource to a user or a machine user:

```
cdp iam assign-user-resource-role \  
--user-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

```
cdp iam assign-machine-user-resource-role \  
--machine-user-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

To remove a resource role from a user or a machine user:

```
cdp iam unassign-user-resource-role \  
--user-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

```
cdp iam unassign-machine-user-resource-role \  
--machine-user-name <value> \  
--resource-role-crn <value> \
```

```
--resource-crn <value>
```

- The resource-role-crn parameter requires the CRN of the resource role you want to assign to the user. You can use the `cdp iam list-resource-roles` command to list resource roles with role CRNs.
- The resource-crn parameter requires the CRN of the resource on which you want to grant the resource role permissions. You can obtain it from the details of the resource.

To get a list of the resource roles assigned to a user or a machine user:

```
cdp iam list-user-assigned-resource-role \  
--user-name <value>
```

```
cdp iam list-machine-user-assigned-resource-role \  
--machine-user-name <value>
```

Managing groups in CDP

A PowerUser can create and manage CDP groups on the CDP web interface or via CDP CLI.

Reserved group names

There are certain group names that are reserved and therefore cannot be used in CDP. This applies to groups synchronized from your identity provider as well as groups created directly from CDP.

If you attempt to synchronize or register a group with a reserved name, you will get an error including the following message:

```
Invalid group name  
Name cannot be a reserved group name
```

To avoid problems, review the following list and avoid synchronising or creating groups with the following names.

The following group names are reserved:

- accumulo
- admins
- atlas
- cruisecontrol
- dpprofiler
- druid
- editors
- flink
- flume
- h2o
- hbase
- hdfs
- hive
- httpfs
- hue
- impala
- ipausers
- kafka

- keytrustee
- kms
- knox
- kudu
- livy
- mapred
- nifi
- nifiregistry
- oozie
- phoenix
- ranger
- rangerraz
- schemaregistry
- sentry
- solr
- spark
- sqoop
- sqoop2
- streamsmgmr
- streamsrepmgr
- tez
- trust admins
- yarn
- yarn-ats
- zeppelin
- zookeeper

Understanding CDP groups

A CDP group is a collection of user accounts that have the same roles and resource roles. A group can include CDP user accounts and machine user accounts. A group cannot include other groups. All users in a group inherit the roles and resource roles assigned to the group.

As a CDP administrator, you can create a group and manage the group membership. You can also manage the roles and resources assigned to the group. If you are not a CDP administrator, you can add users to and remove users from a group if you have the PowerUser role.

When you create a group, you do not automatically become a member of the group. To become a member of the group, you must add your user account to the group.

You can use groups to manage user access more efficiently. If multiple users require the same roles, you can create a group, add the user accounts to the group, and assign the required roles to the group. All user accounts in the group are assigned the roles assigned to the group.

If you delete a group, users in the group lose the roles that they inherit from the group. To allow a user to retain the group roles, assign the same roles to the user separately.

Synchronizing group membership

CDP can synchronize the user's group membership provided by your enterprise IdP with the user's group membership in CDP.

When a user initially logs in to CDP through the identity management system in your organization, CDP creates a CDP user account for the user. However, without being assigned CDP roles, the user cannot perform tasks in CDP.

Cloudera recommends that you create CDP groups with assigned roles and add users to the groups so that the users can take on the roles assigned to the groups.

When you create an identity provider, you can select the Sync Groups on Login option to enable CDP to synchronize the user group membership. By default, the Sync Groups on Login option is disabled. Clear the option selection if you do not want CDP to synchronize the user group membership.

Group names must be alphanumeric, may include dots (.), hyphens (-), and underscores (_), and must be fewer than 64 characters long. Additionally, names can only start with an alphabetic character or an underscore.



Note:

There are certain group names that are reserved and therefore cannot be synchronized to CDP. See [Reserved group names](#).



Note: Cloudera has default limits in place with regard to how many users, machine users, and groups can be added per account. Review [User and group limits](#) and make sure that you do not exceed these limits.

Sync Groups on Login enabled

When the Sync Groups on Login option is enabled, CDP synchronizes a user's group in the following manner:

- The group membership that your enterprise IdP specifies for a user overrides the group membership set up in CDP. Each time a user logs in, CDP updates the user's group membership based on the groups that your enterprise IdP specifies for the user.
- If the group exists in CDP, CDP adds the user to the group. The user takes on all the roles associated with the group.
- If the group does not exist in CDP, CDP creates the group and adds the user to the group. However, no roles are assigned to the new group, so a member of the new group does not take on roles from the group.
- If the user is a member of a group in CDP that is not included in the list provided by your enterprise IdP, CDP removes the user from the group.
- If the list of groups from your enterprise IdP is empty, CDP removes the user from all groups in CDP. After login, the user will not be a member of any CDP group and will not have roles from any group.

To ensure that users can perform tasks in CDP, Cloudera recommends that you set up the groups in CDP with appropriate roles before you assign them to users.

Sync Groups on Login disabled

When the Sync Groups on Login option is disabled, CDP does not synchronize the user's group membership in CDP with the user's group membership provided by the IdP. After login, a user's group membership in CDP is determined by the CDP groups assigned to the user in CDP. The groups assigned to the user in your enterprise IdP are ignored.

Sync Membership option for a newly created group

Additionally, once you have synced your IdP and you create a new group in CDP, you have an option called Sync Membership that determines whether group membership is synced to IdP when a user logs in. By default, Sync Membership is enabled when Sync Groups on Login is enabled.

The following table describes how the global Sync Groups on Login and the per-group Sync Membership options can be used:

	IdP Sync Groups on Login on	IdP Sync Groups on Login off
Group Sync Membership on	Group membership for the specific group is reflected in IdP.	Group membership for the specific group is not reflected in IdP.
Group Sync Membership off	Group membership for the specific group is not reflected in IdP.	Group membership for the specific group is not reflected in IdP.

In other words, if Sync Groups on Login is off at the IdP level, then no groups are getting synced regardless of what the setting for Sync Membership is. But if Sync Groups on Login is turned on at the IDP level, then you have the option to override it for certain groups that you explicitly leave off.

Creating a group

Create CDP groups based on the tasks performed by CDP users in your organization.

Before you begin

Consider the following when selecting a name for your group:

- The group name must be unique. Note that there are certain group names that are reserved and therefore cannot be used in CDP. See [Reserved group names](#).
- The group name can be up to 64 characters and can include only alphanumeric characters, dots (.), hyphens (-), and underscores (_). The first character in the name must be an alphabetic character or underscore.
- The group name is not case sensitive. For example, the group name AAa is equivalent to the group name aaa.
- Depending on your IdP setup in CDP, you may be able to manipulate the Sync Membership option. To learn more about this option, refer to [Synchronizing group membership](#).

Required role: PowerUser

Steps

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Groups.

The Groups page displays the list of all CDP groups.

4. Click Create Group.
5. On the Create Group window, enter the name of the group to create.
6. Click Create.

CDP creates the group and adds it to the list of CDP groups on the Groups page.

For CDP CLI

You can use the following command to create a group:

```
cdp iam create-group \  
--group-name <value>
```

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Adding or removing a user from a group

You can add or remove a CDP user or a machine user account from a group.

Note that:

- SAML login is required to propagate group membership changes from your IdP to CDP. That is, the user who was added or removed from a group must log in to CDP in order for the group membership change to take effect in CDP.
- You cannot add a group to another group.
- All members of the group inherit the roles and resources assigned to the group.

- Cloudera has default limits in place with regard to how many users, machine users, and groups can be added per account. Review [User and group limits](#) and make sure that you do not exceed these limits.

Required roles: `IamGroupAdmin` is the minimum role required for adding or removing users from a group. In addition, in order for a user with the `IamGroupAdmin` to add or remove users from a group via CDP web interface, the user must have either the `IamUser` or the `IamViewer` role that allows listing IAM users and groups. This is not required when adding or removing users from a group via CDP CLI, as long as the admin has the CRN of the user that needs to be added or removed.

Steps

For CDP UI

To add a user to a group:

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Groups.

The Groups page displays the list of all CDP groups.

4. Click the name of the group to which you want to add a user.

The group details page displays information about the group.

5. Click the Members tab.

6. To add a user:

- If the group does not yet have members, click Add Member. Select the name of the user that you want to add to the group.
- If the group already has a list of members, click in the Add a member dropdown box. Select the name of the user that you want to add to the group.

To remove a user from a group, click Remove from Group next to the user that you want to remove. Click OK to confirm that you want to remove the user from the group.

For CDP CLI

The `user-id` parameter requires the CRN of the CDP user or machine user.

You can use the following command to add a user to a group:

```
cdp iam add-user-to-group \  
--group-name <value> \  
--user-id <value>
```

To remove a user from a group:

```
cdp iam remove-user-from-group \  
--group-name <value> \  
--user-id <value>
```

You can use the following command to add a machine user to a group:

```
cdp iam add-machine-user-to-group \  
--group-name <value> \  
--user-id <value>
```

To remove a machine user from a group:

```
cdp iam remove-machine-user-from-group \  
--group-name <value> \  
--machine-user-name <value>
```

To get a list of the users in a group:

```
cdp iam list-group-members \  
--group-name <value>
```

To get the list of groups that a user or machine user is a member of:

```
cdp iam list-groups-for-user \  
--user-id <value>
```

```
cdp iam list-groups-for-machine-user \  
--machine-user-name <value>
```

What to do next

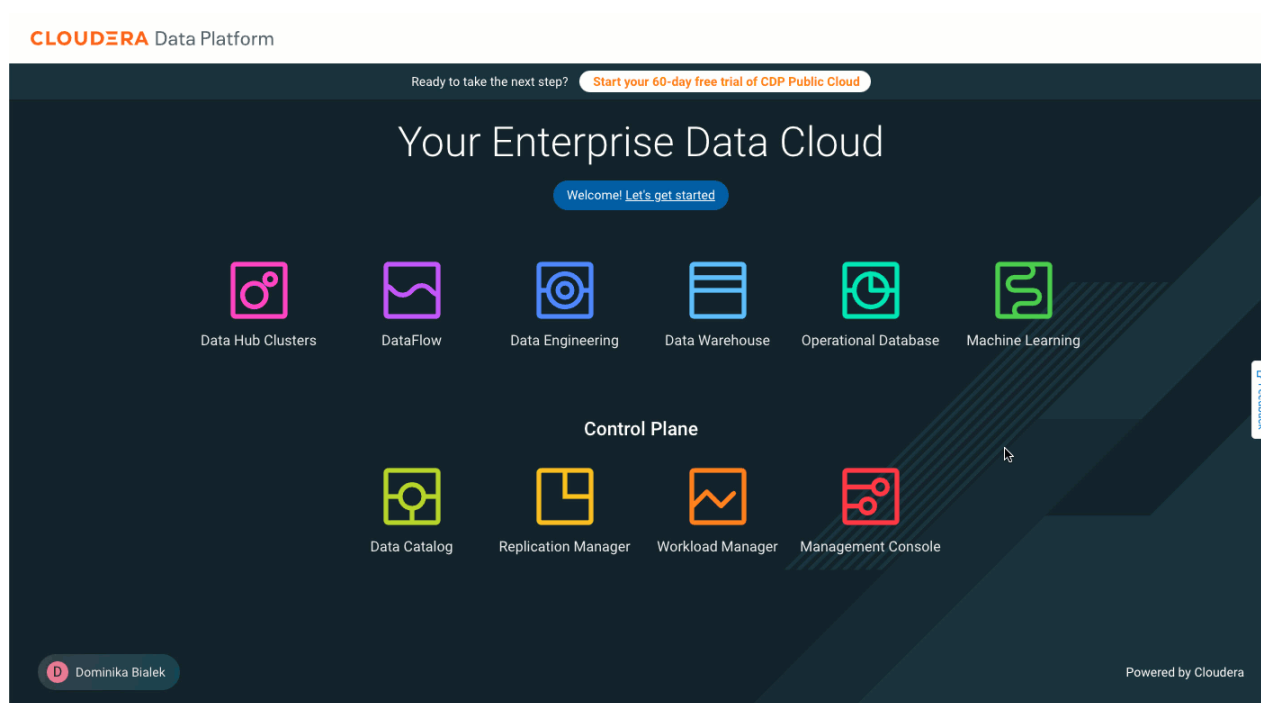
You need to perform user sync for the change to take effect. See [Performing user sync](#).

Assigning account roles to groups

When you assign a role to a group, the role is also assigned to all user and machine user accounts in the group.

Required role: PowerUser

Steps



For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Groups.

The Groups page displays the list of all CDP groups.

4. Click the name of the group to which you want to assign a role.

The group details page displays information about the group.

5. Click the Roles tab.
6. Click Update Roles.
7. On the Update Roles window, select the roles you want to assign to the group.
8. To view the permissions that the role grants to the group, click Policies. To remove a role from the group, clear the selected role.
9. Click Update.

The roles that you select displays in the list of group roles.

To remove a role from a group, click Unassign Role next to the role that you want to remove. Click OK to confirm that you want to remove the role permissions from the group.

For CDP CLI

You can use the following command to assign a role to a group:

```
cdp iam assign-group-role \  
--group-name <value> \  
--role <value>
```

The --role parameter requires the CRN of the CDP role. You can use the `cdp iam list-roles` command to list resource roles with role CRNs.

To get a list of the roles assigned to a group:

```
cdp iam list-group-assigned-roles \  
--group-name <value>
```

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Related Information

[Account roles](#)

Assigning resource roles to groups

When you assign a resource role to a group, the resource role is also assigned to all user and machine user accounts in the group.

To grant a group access to a resource (such as an environment, a shared resource, or a Data Hub cluster), assign a resource role to the group on the scope of the resource and the resource role is also assigned to all user and machine user accounts in the group.

In general, resource roles can be assigned from CDP web interface using the Manage Access option available from the resource details page, or from CDP CLI using the `cdp iam assign-group-resource-role` command. Next, you need to perform user sync. For detailed instructions, see the following sections:

Assign an environment resource role to a group

To assign an environment to a group, assign a specific resource role on the scope of the specific environment.

Required roles:

- Owner or a role that allows administering the environment AND
- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. Navigate to the Environments page.
4. In the list of environments that appear, select an environment by clicking on it.
5. From the Actions menu select Manage Access.
6. In the Access tab, enter the name of the group in the text box.
7. In the Update Resource Roles window, select the required resource role.
8. Click Update Roles.

For CDP CLI

To assign a resource role to a group:

```
cdp iam assign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

To remove a resource role from a group:

```
cdp iam unassign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

- The resource-role-crn parameter requires the CRN of the resource role you want to assign to the group.
- The resource-crn parameter requires the CRN of the resource on which you want to grant the resource role permissions.

To get a list of the resource roles assigned to a group:

```
cdp iam list-group-assigned-resource-role \  
--group-name <value>
```

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Related Information

[Resource roles](#)

Assign a shared resource resource role to a group

You can assign shared resources such as credentials, clusters templates, recipes, image catalogs, or proxies to groups. To assign a shared resource to a group, assign a specific resource role on the scope of the specific shared resource.

Required roles:

- Owner or a role that allows administering the environment AND
- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.

3. From the left pane, select Shared Resources, and then select a resource type (for example Cluster Templates) to view a summary of all resources of that type.
4. Find the specific resource (for example a specific cluster template) and click on it to navigate to its details page.
5. Click on Manage Access.
6. Enter the name of the group in the text box.
7. In the Update Resource Roles window, select the required resource role.
8. Click Update Roles.

For CDP CLI

To assign a resource role to a group:

```
cdp iam assign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

To remove a resource role from a group:

```
cdp iam unassign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

- The resource-role-crn parameter requires the CRN of the resource role you want to assign to the group.
- The resource-crn parameter requires the CRN of the resource on which you want to grant the resource role permissions.

To get a list of the resource roles assigned to a group:

```
cdp iam list-group-assigned-resource-role \  
--group-name <value>
```

Related Information

[Resource roles](#)

Assign a Data Hub resource role to a group

You can assign a Data Hub resource role to a group to allow them to manage a specific Data Hub.



Note: While full access to manage a Data Hub via the Management Console can be granted via assigning the Owner role on the scope of the Data Hub, access to the underlying cluster can only be granted by assigning the EnvironmentUser role on the scope of the environment where the Data Hub is running.

Required roles:

- Owner or a role that allows administering the environment AND
- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. Sign in to the CDP console.

2. Navigate to the details page of your Data Hub cluster. This can be done in a few ways. For example:
 - From the CDP home page, click Data Hub Clusters and then click on the specific cluster.
 - From the CDP home page, click on Management Console, navigate to the Data Hub Clusters page, and then click on the specific cluster.
3. From the Actions menu select Manage Access.
4. Enter the name of the group in the text box.
5. In the Update Resource Roles window, select the required resource role.
6. Click Update Roles.

For CDP CLI

To assign a resource role to a group:

```
cdp iam assign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

To remove a resource role from a group:

```
cdp iam unassign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

- The resource-role-crn parameter requires the CRN of the resource role you want to assign to the group.
- The resource-crn parameter requires the CRN of the resource on which you want to grant the resource role permissions.

To get a list of the resource roles assigned to a group:

```
cdp iam list-group-assigned-resource-role \  
--group-name <value>
```

Related Information

[Resource roles](#)

Assign a classic cluster resource role to a group


You can assign a classic cluster resource role to a group to allow them to manage a specific classic cluster.

Required roles:

- Owner or a role that allows administering the environment AND
- One of the following: `IamViewer` or `IamUser` (required for listing users).

In order to assign a role, a user must have all rights from the role that they are planning to assign to another user; That is, a user can only assign a role higher than his own.

For CDP UI

1. In the Management Console navigate to the Classic Clusters dashboard.
2. Click on the  (context menu) next to the cluster that you want to update and select Manage Access.
3. Find the group that you want to update and click on Update Roles.
4. Select or deselect the roles and then click on Update Roles.

For CDP CLI

To assign a resource role to a group:

```
cdp iam assign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

To remove a resource role from a group:

```
cdp iam unassign-group-resource-role \  
--group-name <value> \  
--resource-role-crn <value> \  
--resource-crn <value>
```

- The resource-role-crn parameter requires the CRN of the resource role you want to assign to the group.
- The resource-crn parameter requires the CRN of the resource on which you want to grant the resource role permissions.

To get a list of the resource roles assigned to a group:

```
cdp iam list-group-assigned-resource-role \  
--group-name <value>
```

Assigning a group membership administrator

As a CDP administrator, you can create a CDP group and manage the users, roles, and resources assigned to the group. You can also assign other users and groups the `IamGroupAdmin` role to allow them to manage the users in the group.

Note that:

- The `IamGroupAdmin` role grants a user or a group the permission to add users to or remove users from a group. The role does not grant permission to manage roles and resources for the group.
- In order for a user with the `IamGroupAdmin` to add or remove users from a group via CDP web interface, the user must have either the `IamUser` or the `IamViewer` role that allows listing IAM users and groups. This is not required when adding or removing users from a group via CDP CLI, as long as the admin has the CRN of the user that needs to be added or removed.

Required role: `PowerUser`

Steps

For CDP UI

To assign a group membership administrator:

1. Sign in to the Cloudera CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Groups.

The Groups page displays the list of all CDP groups.

4. Click the name of the group to which you want to assign a group membership administrator.

The group details page displays information about the group.

5. Click the Admins tab.
6. Click in the Select group or user dropdown box.

CDP displays the list of CDP groups and users that you can give group membership administrator permissions.

7. Select the name of a group or user.

The name of the group or user you select displays in the list of group membership administrators.

To remove group membership administrator permissions from a user or group, click **Remove Resource Role** next to the user or group for whom you want to revoke membership administrator permissions.

For CDP CLI

You assign the `IamGroupAdmin` resource role to users and groups to allow them to manage the users in a specified group.

You can use the following command to assign the `IamGroupAdmin` role to a user:

```
cdp iam assign-user-resource-role \
--user <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

The `user` parameter requires the CRN of the user to whom you want to assign the `IamGroupAdmin` resource role.

The `resource-role-crn` parameter requires the CRN of the `IamGroupAdmin` role.

The `resource-crn` parameter requires the CRN of the group on which the user will have administrator permission.

To assign the `IamGroupAdmin` role to a group:

```
cdp iam assign-group-resource-role \
--group-name <value> \
--resource-role-crn <value> \
--resource-crn <value>
```

The `group-name` parameter requires the name of the group to which you want to assign the `IamGroupAdmin` resource role.

The `resource-role-crn` parameter requires the CRN of the `IamGroupAdmin` role.

The `resource-crn` parameter requires the CRN of the group on which the group specified in the `group-name` parameter will have administrator permission.

For example, to assign the `IamGroupAdmin` to `GroupABC` so that `GroupABC` can manage the users in `GroupXYZ`, run a command similar to the following command:

```
cdp iam assign-group-resource-role \
--group-name groupABC \
--resource-role-crn crn:cdp:iam:us-west-1:cdp:resourceRole:IamGroupAdmin \
--resource-crn crn:cdp:iam:us-west-1:4e9d74e5-1cad-47d8-b645-7ccf9edbb73d:group:GroupXYZ/54218ac1-187b-40f7-aadb-5ghm96c35xy4
```

To assign the users in a group to be the administrators of their own group, set the values of the `group-name` parameter and the `resource-crn` parameter to refer to the same group.

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Updating a group

In some cases, you can enable or disable `SyncMembership` for a group.

Depending on your IdP setup in CDP, you may be able to manipulate the `Sync Membership` option. To learn more about this option, refer to [Synchronizing group membership](#).

Required role: PowerUser

Steps

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Groups.
4. From the context menu to the right of the desired group, click Update Group.
5. Select or deselect the Sync Membership checkbox.
6. Click Update.

For CDP CLI

To update a group:

```
cdp iam update-group \  
--group-name<value> \  
--sync-membership-on-user-login
```

or

```
cdp iam update-group \  
--group-name <value> \  
--no-sync-membership-on-user-login
```

Removing account roles from a group

When you unassign a role to a group, the role is also unassigned to all user and machine user accounts in the group.

Required role: PowerUser

Steps

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Groups.

The Groups page displays the list of all CDP groups.
4. Click the name of the group to which you want to assign a role.

The group details page displays information about the group.
5. Click the Roles tab.
6. From the context menu to the right of a role, click Unassign role.
7. Click OK to confirm that you want to remove the role permissions from the group.

For CDP CLI

To remove a role from a group:

```
cdp iam unassign-group-role \  
--group-name <value> \  
--role <value>
```

The role parameter requires the CRN of the CDP role.

To get a list of the roles assigned to a group:

```
cdp iam list-group-assigned-roles \  
--group-name <value>
```

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Deleting a group

You can delete a group from CDP.

Required role: PowerUser

Steps

For CDP UI

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. In the User Management section of the side navigation panel, click Groups.
4. From the context menu to the right of the desired group, click Delete Group.
5. Click OK to confirm removal.
6. CDP removes the group and removes it from the list of CDP groups on the Groups page.

For CDP CLI

To delete a group:

```
cdp iam delete-group \  
--group-name <value>
```

What to do next

You need to perform user sync for the change to take effect. See [Performing user sync](#).

Performing user sync

When making any kind of user or group-related changes, you need to perform user sync in order for the changes to be synced to FreeIPA.

All user and group related changes with one exception of generating API access keys require user sync. For example all of the following require user sync:

- Creating machine users
- Deleting users and machine users
- Assigning CDP roles or environment-level resource roles to users
- Creating a group, assigning group membership, and deleting a group
- Setting workload password
- Managing user SSH keys

During user sync:

- All control plane actors (users and machine users) with the environments/accessEnvironment right are synced to the FreeIPA.

- All groups are synced to the FreeIPA.
- All users with PowerUser role are synchronized to all environments.

Syncing users to all environments

Required role: PowerUser

Steps

For CDP UI

1. From the CDP web interface, you can perform user sync from the Management Console > User Management.
2. Click Actions > Synchronize Users.
3. Click Synchronize Users.
4. Status shown will be Running, then Completed.

For CDP CLI

From the CDP CLI, you can use the following commands:

```
cdp environments sync-all-users
```

This command synchronizes all users and groups with all CDP environments.

```
cdp environments sync-user
```

This commands only syncs the current user with all their CDP environments. You can use it if you are making changes to your own user, but you can't use it for syncing other users.

What to do next

Depending on how many users you have in CDP, it may take a few minutes for the user sync to complete. The sync operation times out after 30 minutes.

Syncing users to a selected environment

Required role: EnvironmentAdmin and DataSteward can perform user sync for a single environment. PowerUser can perform user sync for all environments.

Steps

For CDP UI

1. From the CDP web interface, you can perform user sync from the Management Console > Environments.
2. Navigate to a specific environment.
3. Do one of the following:
 - Click on Actions > Synchronize Users to FreeIPA
 - Navigate to Summary > FreeIPA > Actions > Synchronize Users to FreeIPA
 - Click on Actions > Manage Access and then click the Synchronize Users button in the top right corner.
4. Click Synchronize Users.
5. Status shown will be Running, then Completed.

For CDP CLI

From the CDP CLI, you can use the following command:

```
cdp environments sync-all-users --environment-names <value>
```

This command synchronizes all users and groups with one or more CDP environments specified in --environment-names <value>.

What to do next

Depending on how many users you have in CDP, it may take a few minutes for the user sync to complete. The sync operation times out after 30 minutes.

Access paths to CDP and its components

To access the various CDP components, you must understand the access paths unique to the entry points that are specific to users and situations.






The typical access entry methods and their details are as follows:

- SSO access through Management Console - After the initial identity provider configuration, users can access CDP and its SSO-interfaces through the Management Console. You can access various CDP services including Management Console, Machine Learning Workspaces, Data Catalog, Replication Manager, Data Hub, and Data Lakes.
- Accessing non-SSO interfaces using workload password - You must set your workload password to add the user credentials to the IPA server. Using the workload username and workload password, you can access non-SSO interfaces such as SSH to clusters, JDBC connections, REST APIs, Data Warehouses, and so on.



Note: To access Machine Learning workspaces, you can enter the kerberos principal and the workload password, and gain kerberized access to components such as the HMS.

- Accessing CDP CLI using access keys - On the User Management section of the Management Console, you can generate your access keys. Use these access keys for CDP CLI and for running jobs.
- Machine user access - To get programmatic access to CDP and its services, you can create and use a machine user. The process to set up the machine user for access is as follows:
 - Create a machine user in the User Management section of Management Console.
 - Get access keys in the Management Console for this machine user.
 - Use APIs and CDP CLI to set the workload password for the machine user.
- SSH access - There are two types of SSH access:
 - Admin users who create Environments can access the environments directly using the SSH key access.
 - All users can access workload clusters by using their SSH keys previously uploaded in CDP.

User Access Paths to CDP			
Access Point	Initial Configuration	Accessible Components	Limitations
 SSO through Management Console	Using SSO configured as part of the initial IDP setup	Management Console, Data Lakes, Data Hub, Data Catalog, Replication Manager, Machine Learning	Exceptions: Data Analytics Studio (DAS) and Grafana
 Access through workload password	Find your workload user ID and set workload password	SSH to cluster, DWX, JDBC, REST APIs	None
 Access using access keys	Generate access keys in the User Management panels	CLI and for running jobs	None
 Machine user access	Create machine user in the User Management panels and set IPA password through CLI	To use APIs to create applications such as ETL	Need to create access keys. Machine user password can be set from the UI by Power Users
 SSH key access	SSH key for power user access is uploaded during environment creation. All users can upload their SSH keys	To access the environment and clusters directly	None

Related Information

[Setting the workload password](#)

[Accessing non-SSO interfaces using workload user and password](#)

[Managing SSH keys](#)

[Generating an API access key](#)

Setting a default identity provider in CDP

You can set a default identity provider (IdP) in CDP for workload-initiated SSO using CDP user interface or CDP CLI.

By default, the oldest configured identity provider is used workload-initiated SSO, but you can optionally set a default IdP using CDP user interface or CDP CLI.

Required roles: Account administrator or PowerUser

Steps

For CDP UI

1. In the CDP user interface, navigate to the Management Console.
2. Select User Management from the navigation pane and then navigate to Identity Providers.
3. Click on the context menu next to the entry for a previously registered identity provider and select Set As Default Identity Provider from the menu.

Once the default identity provider has been updated, you will see the “Default” label next to the idP name.

For CDP CLI

To set a default IdP, use the following command:

```
cdp iam set-default-identity-provider --name-or-crn <IDP-NAME>
```

To print the CRN of the default IdP, use the following command:

```
cdp iam get-default-identity-provider
```

This returns a CRN of the default identity provider.

Accessing non-SSO interfaces using workload user and password

You can access all SSO-based interfaces that can be accessed using the browser with your CDP credentials, but you cannot access SSH-based or other non-SSO connections with these credentials. Instead, you must use your workload user and set the workload password to access such interfaces.

You must import and trust the IPA's root certificate as the IPA CA will be a self-signed CA.



Note:

If you would like to access cluster nodes, you have two options: you can use your workload password or an SSH key. The advantage of the workload password is that it allows a user to interface with Kerberos. Therefore, if you would like to interface with Kerberos, use your workload password to SSH to cluster nodes.

Related Information

[Setting the workload password](#)

Setting the workload password

To access non-SSO interfaces, each user and machine user must set a workload password (also known as "FreeIPA password"). An administrator can set other users' workload passwords.

Required roles: All users can manage their workload passwords from the account management page. All users can manage their workload password from CDP CLI, but this action requires an API access key, which can only be generated by users with the IAMUser role. As a CDP administrator or PowerUser, you can manage the workload password for all user accounts.

Workload password requirements

Your CDP administrator may set a custom workload password policy for your organization. If your CDP administrator did not set a custom workload password policy default, CDP has the following workload password requirements:

- A minimum password length of 8 characters
- Must include at least 1 upper case character, lowercase character, number and special character. Supported special characters are: "#", "&", "*", "\$", "%", "@", "^", ".", "_", and "!".
- All previous passwords can be reused
- The password can be changed at any time
- The password never expires

Set your own workload password

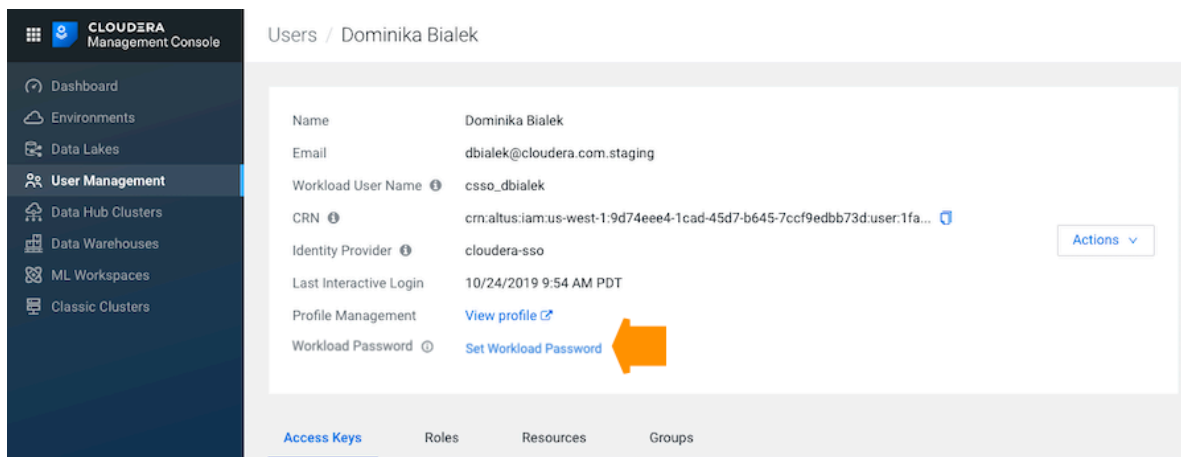
As a CDP user, you can see on your profile page if you have previously set your workload password and if the password is about to expire. There are two cases when you may want to set your workload password:

1. When you first start using CDP.
2. When your password expires. This may or may not happen depending on your company's policies. If your password does expire, you will see a banner notification on the CDP web interface 10 days before the expiry date. You can also see on your user's profile page the state of your workload password (if it expires soon or cannot yet be changed).

Steps

For CDP UI

1. Sign in to the CDP web interface.
2. Click on your user name in the bottom left corner and then select Profile.
3. Click Set Workload Password:



4. In the dialog box that appears, enter the new workload password twice.
5. In the Environments text box, All is pre-selected so that the workload password is synced to all environments by default.
6. Click Set Workload Password. A message appears saying that the password is set successfully.
7. Click Close.

For CDP CLI

Use the following command to set workload password:

```
cdp iam set-workload-password --password <value>
```

Set workload password for another user or machine user (admin only)

There are two cases when you may want to set workload password for a machine user:


1. When you are first onboarding the machine user to CDP.

- When the machine user's password expires. This may or may not happen depending on your company's policies. A CDP administrator or PowerUser is able to navigate to the list of all users to see for which machine user passwords are about to expire. In the "Password expiring" column, any password that is about to expire in 10 days or less is flagged as "Expires in X days". Only machine users (and not human users) are flagged in this manner. A CDP administrator or PowerUser can then reset the password for each machine user whose password is about to expire.


Steps

For CDP UI

- Sign in to the CDP web interface.
- From the CDP home page, click Management Console.
- On the side navigation panel, click User Management.
- On the Users page, enter your name in the search bar and then click on your user name:



- Click Set Workload Password:



- In the dialog box that appears, enter the new workload password twice.
- Click Set Workload Password. A message appears saying that the password is set successfully.
- Click Close.

For CDP CLI

Use the following command to set workload password for other users:

```
cdp iam set-workload-password --actor-crn <value> --password <value>
```

The CRN can be obtained from CDP web interface from the user profile.

What to do next

Each time you reset your workload password, you must regenerate your keytab. See [Retrieve keytab](#).

Related Information

[Accessing non-SSO interfaces using workload user and password](#)

Managing SSH keys

A Power User can add and delete SSH keys for all users and machine users, and users can add and delete their own SSH keys. Once these SSH keys are uploaded and synced, they can be used to access workload cluster nodes. RSA or ED25519 keys are supported.

Required roles: All users can manage their SSH keys from the account management page. All users can manage their SSH keys from CDP CLI, but this action requires an API access key, which can only be generated by users with the IAMUser role. As a CDP administrator or PowerUser, you can manage the SSH keys for all user accounts.

Manage your own SSH keys

Steps

For CDP UI

To add/delete an SSH keys via CDP web interface, click on your user name in the bottom left corner and then select Profile. Next, click on the SSH Keys tab.

- To add an SSH key, click on Add SSH key, then provide a description, paste your SSH public key and click Save.
- To delete click on Delete next to the SSH key that you would like to delete and then click Yes to confirm.

Once the SSH public SSH key is added and synced, the user to which the key is assigned can use a matching private SSH key to access workload cluster nodes.

For CDP CLI

You can manage your SSH keys by using the following CDP CLI commands:

```
cdp iam add-ssh-public-key
cdp iam list-ssh-publi-keys
cdp iam describe-ssh-public-key
cdp iam delete-ssh-public-key
```

Manage SSH keys for another user or machine user (admin only)

Steps

For CDP UI

To add/delete an SSH keys via CDP web interface, navigate to the Management Console > User Management > Users > search for a user name > click on a user name > SSH Keys.

- To add an SSH key, click on Add SSH key, then provide a description, paste your SSH public key and click Save.
- To delete click on Delete next to the SSH key that you would like to delete and then click Yes to confirm.

Once the SSH public SSH key is added and synced, the user to which the key is assigned can use a matching private SSH key to access workload cluster nodes.

For CDP CLI

You can manage SSH keys for other users by using the following CDP CLI commands:

```
cdp iam add-ssh-public-key --actor-crn <value>
cdp iam list-ssh-publi-keys --actor-crn <value>
cdp iam describe-ssh-public-key --actor-crn <value>
cdp iam delete-ssh-public-key --actor-crn <value>
```

Generating an API access key

A CDP user account (a user or a machine user) must have API access credentials to access CDP services through the CDP CLI or API.

Required roles: Users who have the IAMUser role can generate an API access key from their own account page. As a CDP administrator or PowerUser, you can generate an API access key for all user accounts.

Generate your own access key

Steps

1. Sign in to the CDP console.
2. Click on your user name in the bottom left corner and then select Profile.
3. On the user profile page that appears, click Generate Access Key.
4. From the CDP home page, click Management Console.
5. On the side navigation panel, click Users.
6. On the Users page, click the name of the user or machine user account for which you want to generate an access key.
7. On the user account page, go to the Access Keys section and click Generate Access Key.
8. CDP creates the key and displays the information on the screen.

Copy the access key and private key to a text file and send it to the CDP user who requires it. The private key is a very long string of characters. Make sure that you copy the full string. You can optionally download the credentials file containing the access key information.

9. Click OK to exit the access key window.



Note:

The CDP console displays the API access key immediately after you create it. You must copy or download the access key ID and private key information when it is displayed. Do not exit the console without copying the private key. After you exit the console, there is no other way to view or copy the private key.

Once you've generated the access key, you can configure CDP CLI, SDK, or other utilities that require it.

Generate an access key for another user or machine user (admin only)

Steps

1. Sign in to the CDP console.
2. From the CDP home page, click Management Console.
3. On the side navigation panel, click Users.
4. On the Users page, click the name of the user or machine user account for which you want to generate an access key.
5. On the user account page, go to the Access Keys section and click Generate Access Key.

6. CDP creates the key and displays the information on the screen.

Copy the access key and private key to a text file and send it to the CDP user who requires it. The private key is a very long string of characters. Make sure that you copy the full string. You can optionally download the credentials file containing the access key information.

7. Click OK to exit the access key window.

**Note:**

The CDP console displays the API access key immediately after you create it. You must copy or download the access key ID and private key information when it is displayed. Do not exit the console without copying the private key. After you exit the console, there is no other way to view or copy the private key.

Once you've generated the access key, you can configure CDP CLI, SDK, or other utilities that require it.

Retrieve keytabs for workload users

A keytab file stores long-term keys for a principal in Kerberos. Retrieve your keytab for a specific environment.

You may need to generate a keytab for a workload user in certain Data Hub use cases, for example long-running Spark streaming jobs, which require a keytab as a long-lived credential; or NiFi flows requiring a keytab to write data into HBase.

Note that:

- CDP users can retrieve their keytabs. A PowerUser can retrieve keytabs for other users.
- Each time you reset your workload password, you must regenerate your keytab.
- Keytabs are scoped to an environment, whereas workload passwords are the same for every environment. A keytab is, however, tied to the workload password. If you change the workload password, you must retrieve a new keytab. When you change a workload password, retrieve the keytab only after the user sync operation is complete.
- There are ways to generate keytabs with utilities outside of CDP, such as `ipa-getkeytab` or `ktutil`. You should not use these methods as they may not work as expected. For example, `ipa-getkeytab` creates a keytab that may work but only temporarily.

Required roles: All users can retrieve their keytabs from the account management page. All users can retrieve their keytabs from CDP CLI, but this action requires an API access key, which can only be generated by users with the IAMUser role. As a CDP administrator or PowerUser, you can retrieve the keytab for all user accounts.

Before you begin

In order to retrieve a keytab for an environment, you must set workload password for that environment. See [Setting the workload password](#).

Retrieve your own keytab

Steps

For CDP UI

1. Log in to CDP web interface.
2. Click on your user name in the bottom left corner and then select Profile.
3. Click on Actions > Get keytab.
4. In the pop-up window that appears, select the environment for which you would like to get the keytab.
5. Click Download.
6. Save the keytab file in a location of your choice.

For CDP CLI

To retrieve keytab for yourself, use the following command:

```
cdp environments get-keytab --environment-name <VALUE>
```

The output of the command is a base64-encoded representation of a keytab. The contents of the output must be base64 decoded and saved to a file for it to work as a keytab.

What to do next

Once you have downloaded the keytab file, you can copy it to the machine on which the cluster runs and use the keytab to authenticate as the workload user principal, or point to the keytab file when running a Spark job or other job that requires a keytab.

Retrieve keytab for another user or machine user (admin only)

Steps

For CDP UI

1. Log in to CDP web interface.
2. Navigate to the Management Console > User Management, find and click on the user name of the user that you would like to retrieve a keytab for.
3. Click on Actions > Get keytab.
4. In the pop-up window that appears, select the environment for which you would like to get the keytab.
5. Click Download.
6. Save the keytab file in a location of your choice.

For CDP CLI

To retrieve keytab for another user or machine user, use the following command:

```
cdp environments get-keytab --environment-name <VALUE> --actor-crn <CRN>
```

The output of the command is a base64-encoded representation of a keytab. The contents of the output must be base64 decoded and saved to a file for it to work as a keytab.

What to do next

Once you have downloaded the keytab file, you can copy it to the machine on which the cluster runs and use the keytab to authenticate as the workload user principal, or point to the keytab file when running a Spark job or other job that requires a keytab.