

Management Console Top Tasks

Date published: 2022-08-30

Date modified:

The Cloudera logo is displayed in a bold, orange, sans-serif font. The word "CLOUDERA" is written in all caps, with the letter "E" stylized as a horizontal bar with a small triangle in the center.

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

- Verify your credential prerequisites.....4**
 - Cross-account access IAM role.....4
 - Prerequisites for the provisioning credential..... 4
 - Service account for the provisioning credential..... 4

- Register your first environment.....4**
 - Register an AWS environment from CDP UI..... 4
 - Register an AWS environment from CDP CLI..... 9
 - Register an Azure environment from CDP UI..... 9
 - Register an Azure environment from CDP CLI..... 14
 - Register a GCP environment from CDP UI..... 14
 - Register a GCP environment from CDP CLI..... 17

Verify your credential prerequisites

Cross-account access IAM role

To allow CDP to create resources in your AWS account, you create a cross-account access IAM role in your AWS account and grant CDP access to the role as a trusted principal by specifying a specific AWS account and an external ID.

The policy for the cross-account access IAM role must have the permissions enumerated in the documentation linked below. In addition, the IAM role must reference the specific AWS account ID and external ID provided in the Management Console.

Prerequisites for the provisioning credential

To allow CDP to create resources on your Azure account, you must create the app-based credential. The credential allows CDP to access and provision a set of resources in your Azure account.

CDP uses an app-based credential to authenticate your Azure account and obtain authorization to create resources on your behalf. The app-based credential requires that you manually configure the service principal created within your Azure Active Directory. The app-based method requires Owner role to be able to create a service principal, which must be given Contributor role or its equivalent.

To meet Azure prerequisites for CDP:

- Review the provided policies
- Perform the step 1 and step 2 described in the documentation for creating an app-based credential:

Service account for the provisioning credential

The provisioning credential for Google Cloud relies on a service account that can be assumed by CDP.

The following flow describes how the Google Cloud provisioning credential works:

1. Your GCP account administrator creates a service account and assigns the minimum permissions allowing CDP to create and manage resources in your Google Cloud account. Next, the administrator generates a service account access key pair for the service account.
2. The service account is registered as a credential in CDP and its access key is uploaded to CDP.
3. The credential is then used for registering your Google Cloud environment in CDP.
4. Once this is done, CDP uses the credential for provisioning environment-related resources, workload clusters, and resources for other CDP services that you run in CDP.

Review the following to learn about the permissions required for the credential and how to create the service account:

Register your first environment

Register an AWS environment from CDP UI

Once you've met the AWS cloud provider requirements, register your AWS environment.


Before you begin

This assumes that you have already fulfilled the environment prerequisites described in [AWS requirements](#).

Required role: EnvironmentCreator


Steps

1. Navigate to the Management Console > Environments > Register environment:
2. On the Register Environment page, provide the following information:

Parameter	Description
General Information	
Environment Name (Required)	Enter a name for your environment. The name: <ul style="list-style-type: none"> • Must be between 5 and 28 characters long. • Can only include lowercase letters, numbers, and hyphens. • Must start with a lowercase letter.
Description	Enter a description for your environment.
Select Cloud Provider (Required)	Select Amazon.
Credential (Required)	
Select Credential	Select an existing credential or select Create new credential. For instructions on how to create a credential, refer to Creating a role-based credential .  Note: Activate the Enable Permission Verification button if you want CDP to check permissions for your credential. CDP will verify that you have the required permissions for your environment.

3. Click Next.
4. On the Data Access and Data Lake Scaling page, provide the following information:

Parameter	Description
Data Lake Settings	
Data Lake Name (Required)	Enter a name for the Data Lake cluster that will be created for this environment. The name: <ul style="list-style-type: none"> • Must be between 5 and 100 characters long • Must contain lowercase letters • Cannot contain uppercase letters • Must start with a letter • Can only include the following accepted characters are: a-z, 0-9, -.
Data Lake Version (Required)	Select Cloudera Runtime version that should be deployed for your Data Lake. The latest stable version is used by default. All Data Hub clusters provisioned within this Data Lake will be using the same Runtime version.
Fine-grained access control on S3	
Enable Ranger authorization for AWS S3 Identity	Enable this if you would like to use Fine-grained access control . Next, from the Select AWS IAM role for Ranger authorizer dropdown, select the DATALAKE_ADMIN_ROLE IAM role created in Minimal setup for cloud storage .
Data Access and Audit	
Assumer Instance Profile (Required)	Select the IDBROKER_ROLE instance profile created in Minimal setup for cloud storage .
Storage Location Base (Required)	Provide the S3 location created for data storage in Minimal setup for cloud storage .
Data Access Role (Required)	Select the DATALAKE_ADMIN_ROLE IAM role created in Minimal setup for cloud storage .
Ranger Audit Role (Required)	Select the RANGER_AUDIT_ROLE IAM role created in Minimal setup for cloud storage .

Parameter	Description
IDBroker Mappings	<p>We recommend that you leave this out and set it up after registering your environment as part of Onboarding CDP users and groups for cloud storage.</p> <p> Note: If you are using Fine-grained access control, this option is disabled, because you should onboard your users and groups via Ranger instead of using IDBroker mappings.</p>
Scale (Required)	<p>Select Data Lake scale. By default, “Light Duty” is used.</p> <p>For more information on Data Lake scale, refer to Data Lake scale.</p>
Enable Compute Cluster	<p>Enable Compute Clusters if you would like to deploy a containerized platform on Kubernetes for data services and shared services.</p>


5. Click on Advanced Options to make additional configurations for your Data Lake. The following options are available:


Parameter	Description
Network and Availability	
Enable Multiple Availability Zones for Data Lake	<p>Click the Enable Multiple Availability Zones for Data Lake toggle button to enable multi-AZ for the Data Lake. This option is disabled by default and is only available when a Medium Duty Data Lake is selected. Refer to Deploying CDP in multiple AWS availability zones.</p>
Hardware and Storage	<p>For each host group you can specify an instance type. For more information on instance types, see Amazon EC2 instance types.</p>
Cluster Extensions	
Recipes	<p>You can optionally select and attach previously registered recipes to run on a specific Data Lake host group. For more information, see Recipes.</p>

6. Click Next.

7. On the Region, Networking and Security page, provide the following information:

Parameter	Description
Region	
Select Region (Required)	<p>Select the region that you would like to use for CDP.</p> <p>If you would like to use a specific existing virtual network, the virtual network must be located in the selected region.</p>
Customer-managed Keys	
Enable Customer-Managed Keys	<p>Enable this if you would like to provide a Customer-Managed Key (CMK) to encrypt environment's disks and databases. Next, under Select Encryption Key, select an existing CMK. For more information, refer to Customer managed encryption keys.</p>
Select Encryption Key	Select an existing CMK.
Network	
Select Network (Required)	<p>You have two options:</p> <ul style="list-style-type: none"> Select the existing virtual network where you would like to provision all CDP resources. Refer to VPC and subnet. Select Create new network to have a new network with three subnets created. One subnet is created for each availability zone assuming three AZs per region; If a region has two AZs instead of three, then still three subnets are created, two in the same AZ.
Select Subnets (Required)	<p>This option is only available if you choose to use an existing network. Multiple subnets must be selected and CDP distributes resources evenly within the subnets.</p>
Network CIDR (Required)	<p>This option is only available if you select to create a new network.</p> <p>If you selected to create a new network, provide Network CIDR that determines the range of private IPs that EC2 instances will use. This must be a valid private IP CIDR IP in IPv4 range.</p> <p>For example 10.10.0.0/16 are valid IPs. /16 is required to allow for enough IP addresses.</p>

Parameter	Description
Create Private Subnets	<p>This option is only available if you select to have a new network and subnets created. Is turned on by default so that private subnets are created in addition to public subnets. If you disable it, only public subnets will be created.</p> <p> Important: For production deployments, Cloudera recommends that you use private subnets. Work with your internal IT teams to ensure that users can access the browser interfaces for cluster services.</p>
Create Private Endpoints	<p>This option is only available if you select to have a new network and subnets created. It is disabled by default. Enable this option to use private endpoints instead of public endpoints for the following services:</p> <ul style="list-style-type: none"> • Amazon EC2 • Amazon ECR - api and dkr • Amazon EFS • Amazon RDS for PostgreSQL • AWS Auto Scaling • AWS CloudFormation • AWS ELB • AWS S3 • AWS STS
Enable Public Endpoint Access Gateway	<p>When CCM is enabled, you can optionally enable Public Endpoint Access Gateway to provide secure connectivity to UIs and APIs in Data Lake and Data Hub clusters deployed using private networking.</p> <p>If you are using your existing VPC, under Select Endpoint Access Gateway Subnets, select the public subnets for which you would like to use the gateway. The number of subnets must be the same as under Select Subnets and the availability zones must match. For more information, refer to Public Endpoint Access Gateway documentation.</p>
Proxies	
Select Proxy Configuration	Select a proxy configuration if previously registered. For more information refer to Setting up a proxy server .
Security Access Settings	
Select Security Access Type (Required)	<p>This determines inbound security group settings that allow connections to the Data Lake and Data Hub clusters from your organization's computers. You have two options:</p> <ul style="list-style-type: none"> • Create new security groups - Allows you to provide custom CIDR IP range for all new security groups that will be created for the Data Lake and Data Hub clusters so that users from your organization can access cluster UIs and SSH to the nodes. <p>This must be a valid CIDR IP in IPv4 range. For example: 192.168.27.0/24 allows access from 192.168.27.0 through 192.168.27.255. You can specify multiple CIDR IP ranges separated with a comma. For example: 192.168.27.0/24,192.168.28.0/24</p> <p>If you use this setting, several security groups will get created: one for each Data Lake host group (the Data Lake and one for each host group), one for each FreeIPA host group, and one for RDS; Furthermore, the security group settings specified will be automatically used for Data Hub, Data Warehouse, and Machine Learning clusters created as part of the environment.</p> <ul style="list-style-type: none"> • Provide existing security groups (Only available for an existing VPC) - Allows you to select two existing security groups, one for Knox-installed nodes and another for all other nodes. If you select this option, refer to Security groups to ensure that you open all ports required for your users to access environment resources.
Kubernetes	
Select Private Kubernetes Cluster or provide Authorized IP Ranges	<p>If you have enabled Compute Clusters, you have the following options to configure the necessary networking information for the Kubernetes cluster:</p> <ul style="list-style-type: none"> • Enable Private Kubernetes Cluster to create a private cluster that blocks all access to the API Server endpoint. • Provide the CIDRs to the Kubernetes API Server Authorized IP Ranges field to specify a set of IP ranges that will be allowed to access the Kubernetes API server. <p>You need to provide the advanced configurations only once when creating your environment. The configurations will be applied to all compute clusters in the environment.</p>
Worker Node Subnets	Uses the same set of subnets provided in Network section. You have the option to not use all of the previously provided subnets.

Parameter	Description
SSH Settings	
New or existing SSH public key (Required)	<p>You have two options for providing a public SSH key:</p> <ul style="list-style-type: none"> Select a key that already exists on your AWS account within the specific region that you would like to use. Upload a public key directly from your computer. <p> Note: CDP does not use this SSH key. The matching private key can be used by your CDP administrator for root-level access to the instances provisioned for the Data Lake and Data Hub.</p>
Add tags	You can optionally add tags to be created for your resources on AWS. Refer to Defining custom tags .

8. Click on Advanced Options to make additional configurations for the FreeIPA cluster. The following options are available:

Parameter	Description
Network and Availability	
Enable Multiple Availability Zones for Data Lake	Click the Enable Multiple Availability Zones for Data Lake toggle button to enable multi-AZ for the FreeIPA cluster. Refer to Deploying CDP in multiple AWS availability zones .
Hardware and Storage	For each host group you can specify an instance type. For more information on instance types, see Amazon EC2 instance types .
Cluster Extensions	
Recipes	You can optionally select and attach previously registered recipes to run on a specific FreeIPA host group. For more information, see Recipes .

9. Click Next.

10. On the Storage page, provide the following information:

Parameter	Description
Logs	
Logger Instance Profile (Required)	Select the LOG_ROLE instance profile created in Minimal setup for cloud storage .
Logs Location Base (Required)	Provide the S3 location created for log storage in Minimal setup for cloud storage .
Backup Location Base	Provide the S3 location created for FreeIPA and Data Lake backups in Minimal setup for cloud storage . If not provided, the default Backup Location Base uses the Logs Location Base.
Telemetry	
Enable Workload Analytics	Enables Cloudera Observability support for workload clusters created within this environment. When this setting is enabled, diagnostic information about job and query execution is sent to Cloudera Observability. For more information, refer to Enabling workload analytics and logs collection .
Enable Deployment Cluster Logs Collection	When this option is enabled, the logs generated during deployments will be automatically sent to Cloudera. For more information, refer to Enabling workload analytics and logs collection .

11. Click on Register Environment to trigger environment registration.

12. The environment creation takes about 60 minutes. The creation of the FreeIPA server and Data Lake cluster is triggered. You can monitor the progress from the web UI. Once the environment creation has been completed, its status will change to “Running”.

After you finish

After your environment is running, perform the following steps:

- You must assign roles to specific users and groups for the environment so that selected users or user groups can access the environment. Next, you need to perform user sync. For steps, refer to [Enabling admin and user access to environments](#).
- You must onboard your users and/or groups for cloud storage. For steps, refer to [Onboarding CDP users and groups for cloud storage](#).

- You must create Ranger policies for your users. For instructions on how to access your Data Lake, refer to [Accessing Data Lake services](#). Once you've accessed Ranger, [create Ranger policies](#) to determine which users have access to which databases and tables.

Register an AWS environment from CDP CLI

Once you've met the AWS cloud provider requirements, register your AWS environment.

Before you begin

This assumes that you have already fulfilled the environment prerequisites described in [AWS requirements](#).

Required role: EnvironmentCreator

Steps

Unlike in the CDP web interface, in CDP CLI environment creation is a three-step process with environment creation, setting IDBroker mappings and Data Lake creation being three separate steps. The easiest way to obtain the correct commands is to provide all parameters in CDP web interface and then generate the CDP CLI commands on the last page of the wizard. For detailed steps, refer to [Obtain CLI commands for registering an environment](#).

To learn more about how to create Compute Cluster enabled environments with CLI, see [Enabling default Compute Cluster for new environments](#).

After you finish

After your environment is running, perform the following steps:

- You must assign roles to specific users and groups for the environment so that selected users or user groups can access the environment. Next, you need to perform user sync. For steps, refer to [Enabling admin and user access to environments](#).
- You must onboard your users and/or groups for cloud storage. For steps, refer to [Onboarding CDP users and groups for cloud storage](#).
- You must create Ranger policies for your users. For instructions on how to access your Data Lake, refer to [Accessing Data Lake services](#). Once you've accessed Ranger, [create Ranger policies](#) to determine which users have access to which databases and tables.

Register an Azure environment from CDP UI

Once you've met the Azure cloud provider requirements, register your Azure environment.

Before you begin

This assumes that you have already fulfilled the environment prerequisites described in [Azure requirements](#).

Required role: EnvironmentCreator

Steps


- Navigate to the Management Console > Environments > Register environment:
- On the Register Environment page, provide the following information:

Parameter	Description
General Information	
Environment Name (Required)	Enter a name for your environment. The name: <ul style="list-style-type: none"> Must be between 5 and 28 characters long. Can only include lowercase letters, numbers, and hyphens. Must start with a lowercase letter.
Description	Enter a description for your environment.

Parameter	Description
Select Cloud Provider (Required)	Select Azure.
Microsoft Azure Credential (Required)	
Select Credential	Select an existing credential or select Create new credential. For instructions on how to create a credential, refer to Create an app-based credential .

3. Click Next.

4. On the Data Access and Data Lake Scaling page, provide the following information:

Parameter	Description
Data Lake Settings	
Data Lake Name (Required)	Enter a name for the Data Lake cluster that will be created for this environment. The name: <ul style="list-style-type: none"> Must be between 5 and 100 characters long Must contain lowercase letters Cannot contain uppercase letters Must start with a letter Can only include the following accepted characters are: a-z, 0-9, -, .
Data Lake Version (Required)	Select Cloudera Runtime version that should be deployed for your Data Lake. The latest stable version is used by default. All Data Hub clusters provisioned within this Data Lake will be using the same Runtime version.
Fine-grained access control on ADLS Gen2	
Enable Ranger authorization for ADLS Gen2 Identity	If you would like to use Fine-grained access control , enable this option and then select the Ranger RAZ managed identity created in the Minimal setup for cloud storage .
Data Access and Audit	
Assumer Identity (Required)	Select the Assumer managed identity created in Minimal setup for cloud storage .
Storage Location Base (Required)	Provide the ADLS Gen2 location created for data storage in Minimal setup for cloud storage .
Data Access Identity (Required)	Select the Data Lake Admin managed identity created in Minimal setup for cloud storage .
Ranger Audit Identity (Required)	Select the Ranger Audit managed identity created in Minimal setup for cloud storage .
IDBroker Mappings	We recommend that you leave this out and set it up after registering your environment as part of Onboarding CDP users and groups for cloud storage .  Note: If you are using Fine-grained access control , this option is disabled, because you should onboard your users and groups via Ranger instead of using IDBroker mappings.
Scale (Required)	Select Data Lake scale. By default, "Light Duty" is used. For more information on data lake scale, refer to Data Lake scale .
Enable Compute Cluster	Enable Compute Clusters if you would like to deploy a containerized platform on Kubernetes for data services and shared services.

5. Click on Advanced Options to make additional configurations for your Data Lake. The following options are available:



Parameter	Description
Hardware and Storage	For each host group you can specify an instance type. For more information on instance types, see Sizes for virtual machines in Azure .


Parameter	Description
Cluster Extensions	
Recipes	You can optionally select and attach previously registered recipes to run on a specific Data Lake host group. For more information, see Recipes .

6. Click Next.

7. On the Region, Networking and Security page, provide the following information:

Parameter	Description
Region	
Select Region (Required)	<p>Select the region that you would like to use for accessing and provisioning resources from CDP.</p> <p>If you would like to use a specific existing virtual network, the virtual network must be located in the selected region.</p>
Resource Group	
Select Resource Group (Required)	<p>You have two options:</p> <ul style="list-style-type: none"> Select one existing resource group. If you select this, all CDP resources will be provisioned into that resource group. Select Create new resource groups to have CDP create multiple resource groups.
Customer Managed Encryption Keys	
Enable Customer-Managed Keys	Enable this if you would like to provide a Customer-Managed Key (CMK) to encrypt environment's disks and databases. For more information, refer to Customer managed encryption keys .
Select Encryption Key Resource Group	Select the resource group where the CMK is located.
Encryption key URL	Provide the URL of the key value where the CMK resides. This is the same as the key identifier that you can copy directly from Azure Portal.
Managed identity for encryption	If using Azure Database for PostgreSQL Flexible Server, you can optionally select a managed identity created for encrypting it. For more information, refer to Managed identity for encrypting Azure Database for PostgreSQL Flexible Server .
Network	
Select Network (Required)	<p>You have two options:</p> <ul style="list-style-type: none"> Select the existing virtual network where you would like to provision all CDP resources. Refer to VNet and subnets. Select Create new network to have a new network with three subnets created.
Select Subnets (Required)	This option is only available if you choose to use an existing network. Multiple subnets must be selected and CDP distributes resources evenly within the subnets.
Network CIDR (Required)	<p>This option is only available if you select to create a new network.</p> <p>If you selected to create a new network, provide Network CIDR that determines the range of private IPs that VMs will use. This must be a valid private IP CIDR IP in IPv4 range.</p> <p>For example 10.10.0.0/16 are valid IPs. /16 is required to allow for enough IP addresses.</p>
Create Private Subnets	<p>This option is only available if you select to have a new network and subnets created. Is is turned on by default so that private subnets are created in addition to public subnets. If you disable it, only public subnets will be created.</p> <p> Important: For production deployments, Cloudera recommends that you use private subnets. Work with your internal IT teams to ensure that users can access the browser interfaces for cluster services.</p>

Parameter	Description
Enable Public Endpoint Access Gateway	<p>When CCM is enabled, you can optionally enable Public Endpoint Access Gateway to provide secure connectivity to UIs and APIs in Data Lake and Data Hub clusters deployed using private networking.</p> <p>If you are using your existing VPC, under Select Endpoint Access Gateway Subnets, select the public subnets for which you would like to use the gateway. The number of subnets must be the same as under Select Subnets and the availability zones must match. For more information, refer to Public Endpoint Access Gateway documentation.</p>
Create Private Endpoints	<p>By default, the PostgreSQL Azure database provisioned for your Data Lake is reachable via a service endpoint (public IP address). To increase security, you can optionally select to have it reachable via a private endpoint instead of a service endpoint.</p> <p> Note: This option is only available if an existing resource group is selected.</p> <p> Note: Only the subnets that have Azure private endpoint network policies turned off are eligible for private endpoint creation. At least one such subnet is required.</p> <p>If you select to create a private endpoint and you are using your own VNet, you have two options:</p> <ul style="list-style-type: none"> Select “Create new private DNS zone” and CDP creates and manages a private DNS zone for you in the provided existing resource group. Select your existing private DNS zone. <p>If you select to create a private endpoint and you would like for CDP to create a new VNet, CDP creates a private DNS zone for you.</p> <p>For more information, refer to Private endpoint for Azure Postgres.</p>
Create Public IPs	This option is disabled by default when CCM is enabled and enabled by default when CCM is disabled.
Flexible Server	During environment registration in CDP, the Flexible Server in public service mode is used by default, but you can specify to use the Flexible Server in private service mode (“Flexible Server with Private Link” or “Flexible Server with Delegated Subnet (deprecated)”). For more information, refer to Using Azure Database for PostgreSQL Flexible Server .
Proxies	
Select Proxy Configuration	Select a proxy configuration if previously registered. For more information refer to Setting up a proxy server .
Security Access Settings	
Select Security Access Type (Required)	<p>This determines inbound security group settings that allow connections to the Data Lake and Data Hub clusters from your organization’s computers. You have two options:</p> <ul style="list-style-type: none"> Create new security groups - Allows you to provide custom CIDR IP range for all new security groups that will be created for the Data Lake and Data Hub clusters so that users from your organization can access cluster UIs and SSH to the nodes. <p>This must be a valid CIDR IP in IPv4 range. For example: 192.168.27.0/24 allows access from 192.168.27.0 through 192.168.27.255. You can specify multiple CIDR IP ranges separated with a comma. For example: 192.168.27.0/24,192.168.28.0/24.</p> <p>If you use this setting, several security groups will get created: one for each Data Lake host group the Data Lake and one for each host group), one for each FreeIPA host group, and one for RDS; Furthermore, the security group settings specified will be automatically used for Data Hub, Data Warehouse, and Machine Learning clusters created as part of the environment.</p> Provide existing security groups (Only available for an existing VPC) - Allows you to select two existing security groups, one for Knox-installed nodes and another for all other nodes. If you select this option, refer to Security groups to ensure that you open all ports required for your users to access environment resources.
Kubernetes	

Parameter	Description
Enable Private Kubernetes Cluster or provide Authorized IP Ranges	<p>If you have enabled Compute Clusters, you have the following options to configure the necessary networking information for the Kubernetes cluster:</p> <ul style="list-style-type: none"> • Enable Private Kubernetes Cluster to create a private cluster that blocks all access to the API Server endpoint. • Provide the CIDRs to the Kubernetes API Server Authorized IP Ranges field to specify a set of IP ranges that will be allowed to access the Kubernetes API server. <p>You need to provide the advanced configurations only once when creating your environment. The configurations will be applied to all compute clusters in the environment.</p>
Enable User Defined Routing	Enable User Defined Routing (UDR) in case public IPs are blocked for egress. In case you enable UDR, you must select the specific worker node subnet where the UDR is configured.
AKS Private DNS Zone ID	When selecting Private Kubernetes Cluster , you also need to select an existing private DNS zone or select creating a new private DNS zone by CDP on your Azure account for the database.
Worker Node Subnets	Uses the same set of subnets provided in Network section. You have the option to not use all of the previously provided subnets.
SSH Settings	
New SSH public key (Required)	<p>Upload a public key directly from your computer.</p>  <p>Note: CDP does not use this SSH key. The matching private key can be used by your CDP administrator for root-level access to the instances provisioned for the Data Lake and Data Hub.</p>
Add tags	You can optionally add tags to be created for your resources on Azure. Refer to Defining custom tags .

8. Click on Advanced Options to make additional configurations for FreeIPA. The following options are available:

Parameter	Description
Hardware and Storage	For each host group you can specify an instance type. For more information on instance types, see Sizes for virtual machines in Azure .
Cluster Extensions	
Recipes	You can optionally select and attach previously registered recipes to run on FreeIPA nodes. For more information, see Recipes .

9. Click Next.

10. On the Storage page, provide the following information:

Parameter	Description
Logs	
Logger Identity (Required)	Select the Logger managed identity created in Minimal setup for cloud storage .
Logs Location Base (Required)	Provide the ADLS Gen2 location created for log storage in Minimal setup for cloud storage .
Backup Location Base	Provide the ADLS Gen2 location created for FreeIPA and Data Lake backups in Minimal setup for cloud storage . If not provided, the default Backup Location Base uses the Logs Location Base.
Telemetry	
Enable Workload Analytics	Enables Cloudera Observability support for workload clusters created within this environment. When this setting is enabled, diagnostic information about job and query execution is sent to Cloudera Observability. For more information, refer to Enabling workload analytics and logs collection .
Enable Deployment Cluster Logs Collection	When this option is enabled, the logs generated during deployments will be automatically sent to Cloudera. For more information, refer to Enabling workload analytics and logs collection .

11. Click on Register Environment to trigger environment registration.

12. The environment creation takes about 60 minutes. The creation of the FreeIPA server and Data Lake cluster is triggered. You can monitor the progress from the web UI. Once the environment creation has been completed, its status will change to “Running”.

After you finish

After your environment is running, perform the following steps:

- You must assign roles to specific users and groups for the environment so that selected users or user groups can access the environment. Next, you need to perform user sync. For steps, refer to [Enabling admin and user access to environments](#).
- You must onboard your users and/or groups for cloud storage. For steps, refer to [Onboarding CDP users and groups for cloud storage](#).
- You must create Ranger policies for your users. For instructions on how to access your Data Lake, refer to [Accessing Data Lake services](#). Once you've accessed Ranger, [create Ranger policies](#) to determine which users have access to which databases and tables.

Register an Azure environment from CDP CLI

Once you've met the Azure cloud provider requirements, register your Azure environment.

Before you begin

This assumes that you have already fulfilled the environment prerequisites described in [Azure requirements](#).

Required role: EnvironmentCreator

Steps

Unlike in the CDP web interface, in CDP CLI environment creation is a three-step process with environment creation, setting IDBroker mappings and Data Lake creation being three separate steps. The easiest way to obtain the correct commands is to provide all parameters in CDP web interface and then generate the CDP CLI commands on the last page of the wizard. For detailed steps, refer to [Obtain CLI commands for registering an environment](#).

To learn more about how to create Compute Cluster enabled environments with CLI, see [Enabling default Compute Cluster for new environments](#).

After you finish

After your environment is running, perform the following steps:

- You must assign roles to specific users and groups for the environment so that selected users or user groups can access the environment. Next, you need to perform user sync. For steps, refer to [Enabling admin and user access to environments](#).
- You must onboard your users and/or groups for cloud storage. For steps, refer to [Onboarding CDP users and groups for cloud storage](#).
- You must create Ranger policies for your users. For instructions on how to access your Data Lake, refer to [Accessing Data Lake services](#). Once you've accessed Ranger, [create Ranger policies](#) to determine which users have access to which databases and tables.

Register a GCP environment from CDP UI

Once you've met the Google Cloud cloud provider requirements, register your GCP environment.

Before you begin

This assumes that you have already fulfilled the environment prerequisites described in [GCP requirements](#).

Required role: EnvironmentCreator

Steps

1. Navigate to the Management Console > Environments > Register environment.
2. On the On the Register Environment page, provide the following information:

Parameter	Description
General Information	
Environment Name (Required)	Enter a name for your environment. The name: <ul style="list-style-type: none"> • Must be between 5 and 28 characters long. • Can only include lowercase letters, numbers, and hyphens. • Must start with a lowercase letter.
Description	Enter a description for your environment.
Select Cloud Provider (Required)	Select Google Cloud.
Google Cloud Platform Credential (Required)	
Select Credential	Select an existing credential or select Create new credential. For instructions on how to create a credential for Google Cloud, refer to Create a provisioning credential for GCP .

3. Click Next.
4. On the Data Access and Data Lake Scaling page, provide the following information:

Parameter	Description
Data Lake Settings	
Data Lake Name (Required)	Enter a name for the Data Lake cluster that will be created for this environment. The name: <ul style="list-style-type: none"> • Must be between 5 and 100 characters long • Must contain lowercase letters • Cannot contain uppercase letters • Must start with a letter • Can only include the following accepted characters are: a-z, 0-9, -, .
Data Lake Version (Required)	Select Cloudera Runtime version that should be deployed for your Data Lake. The latest stable version is used by default. All Data Hub clusters provisioned within this Data Lake will be using the same Runtime version. Note: Google Cloud environments can only be provisioned in CDP with Runtime version 7.2.8 or newer.
Data Access and Audit	
Assumer Service Account (Required)	Select the IDBroker service account created in Minimum setup for cloud storage .
Storage Location Base (Required)	Select the Google Storage location created for data in Minimum setup for cloud storage .
Data Access Service Account (Required)	Select the Data Lake Admin service account created in Minimum setup for cloud storage .
Ranger Audit Service Account (Required)	Select the Ranger Audit service account created in Minimum setup for cloud storage .
IDBroker Mappings	We recommend that you leave this out and set it up after registering your environment as part of Onboarding CDP users and groups for cloud storage .
Scale (Required)	Select Data Lake scale. By default, "Light Duty" is used. For more information on Data Lake scale, refer to Data Lake scale .

5. Click on Advanced Options to make additional configurations for your Data Lake. The following options are available:

Parameter	Description
Hardware and Storage	For each host group you can specify an instance type. For more information on instance types, see Machine type families .
Cluster Extensions	
Recipes	You can optionally select and attach previously registered recipes to run on a specific Data Lake host group. For more information, see Recipes .

6. Click Next.

7. On the Region, Networking and Security page, provide the following information:

Parameter	Description
Region	
Select Region (Required)	Select the region where your VPC network is located.
Select Zone (Required)	Select a zone within the selected region.
Network	
Use shared VPC	This option is disabled by default. Enable this if you would like to use your existing shared VPC. Next enter: <ul style="list-style-type: none"> Host project ID Network name Subnet name(s). If providing multiple, provide a comma separated list.
Select Network (Required)	Select the existing VPC network that you created as a prerequisite in the VPC network and subnets step. All CDP resources will be provisioned into this network.
Select Subnets (Required)	Select at least one existing subnet.
Create Public IPs	This option is disabled by default when CCM is enabled and enabled by default when CCM is disabled.
Proxies	Select a proxy configuration if previously registered. For more information refer to Setting up a proxy server .
Security Access Settings	
Select Security Access Type (Required)	You have two options: <ul style="list-style-type: none"> Do not create firewall rule: If you are using a shared VPC you can set the firewall rules directly on the VPC. If you did so, you can select this option. Provide existing firewall rules: If not all of your firewall rules are set directly on the VPC, provide the previously created firewall rules for SSH an UI access. You should select two existing firewall rules, one for Knox gateway-installed nodes and another for all other nodes. You may select the same firewall rule in both places if needed. For information on required ports, see Firewall rules .
SSH Settings	
New SSH public key (Required)	Upload a public key directly from your computer. Note: CDP does not use this SSH key. The matching private key can be used by your CDP administrator for root-level access to the instances provisioned for the Data Lake and Data Hub.
Add tags	You can optionally add tags to be created for your resources on GCP. Refer to Defining custom tags .

8. Click on Advanced Options to make additional configurations for FreeIPA. The following options are available:

Parameter	Description
Hardware and Storage	For each host group you can specify an instance type. For more information on instance types, see Machine type families .
Cluster Extensions	
Recipes	You can optionally select and attach previously registered recipes to run on FreeIPA nodes. For more information, see Recipes .

9. Click Next.

10. On the Storage page, provide the following information:

Parameter	Description
Logs	
Logger Service Account (Required)	Select the Logger service account created in Minimum setup for cloud storage .
Logs Location Base (Required)	Select the Google Storage location created for logs in Minimum setup for cloud storage .
Backup Location Base	Select the Google Storage location created for FreeIPA backups in Minimum setup for cloud storage . If not provided, the default Backup Location Base uses the Logs Location Base.
Telemetry	
Enable Workload Analytics	Enables Cloudera Observability support for workload clusters created within this environment. When this setting is enabled, diagnostic information about job and query execution is sent to the Cloudera Observability.
Enable Deployment Cluster Logs Collection	When this option is enabled, the logs generated during deployments will be automatically sent to Cloudera.

11. Click Register Environment to trigger environment registration.

12. The environment creation takes about 60 minutes. The creation of the FreeIPA server and Data Lake cluster is triggered. You can monitor the progress from the web UI. Once the environment creation has been completed, its status will change to “Running”.

After you finish

After your environment is running, perform the following steps:

- You must assign roles to specific users and groups for the environment so that selected users or user groups can access the environment. Next, you need to perform user sync. For steps, refer to [Enabling admin and user access to environments](#).
- You must onboard your users and/or groups for cloud storage. For steps, refer to [Onboarding CDP users and groups for cloud storage](#).
- You must create Ranger policies for your users. For instructions on how to access your Data Lake, refer to [Accessing Data Lake services](#). Once you've accessed Ranger, [create Ranger policies](#) to determine which users have access to which databases and tables.

Register a GCP environment from CDP CLI

Once you've met the Google Cloud cloud provider requirements, register your GCP environment.

Before you begin

This assumes that you have already fulfilled the environment prerequisites described in [GCP requirements](#).

Required role: EnvironmentCreator

Steps

Unlike in the CDP web interface, in CDP CLI environment creation is a two-step process with environment creation and data lake creation being two separate steps. The following commands can be used to create an environment in CDP.

1. Once you've met the prerequisites, register your GCP environment in CDP using the `cdp environments create-gcp-environment` command and providing the CLI input parameters. For example:

```
cdp environments create-gcp-environment --cli-input-json '{
  "environmentName": "test-env",
  "description": "Test GCP environment",
  "credentialName": "test-gcp-crd",
  "region": "us-west2",
  "publicKey": "ssh-rsa AAAAB3NzaZ1yc2EAAAADAQABAAQDwCI/wmQzbNn9YcA8v
dU+Ot4lIIUWJfOfiDrUuNcULOQL6ke5qcEKuboXzbLxV0YmQcPFvswbM5S4FlHjy2VrJ5spy
GhQajfEm9+PgrsybgzHkssziX0zRq7U4BVD68kSn6CuAHj9L4wx8WBwefMzkw7u01CkfiFI
p8UE6ZcKKKwe2fLR6ErDa9jQxIWhTPEiFjIhItPHrnOcfGKY/p6OlDDUOuMRiFZh7qMzfg
vWI+UdN/qjnTlc/M53JftK6GJqK6osN+j7fCwKENpWC/gmy8El7ZMH1IENxDut6X0qj9Okc/
JMmG0ebkSZAebhgNOBNLZYdP0oeQGCXjqdv",
  "enableTunnel": true,
  "usePublicIp": true,
  "existingNetworkParams": {
    "networkName": "eng-private",
    "subnetNames": [
      "private-us-west2"
    ],
    "sharedProjectId": "dev-project"
  },
  "logStorage": {
    "storageLocationBase": "gs://logs",
    "serviceAccountEmail": "logger@dev-project.iam.gserviceaccount.com"
  }
}'
```

Parameter	Description
environmentName	Provide a name for your environment.
credentialName	Provide the name of the credential created earlier.
region	Specify the region where your existing VPC network is located. For example "us-west2" is a valid region.
publicKey	Paste your SSH public key.
existingNetworkParams	<p>Provide a JSON specifying the following:</p> <pre>{ "networkName": "string", "subnetNames": ["string", ...], "sharedProjectId": "string" }</pre> <p>Replace the values with the actual VPC network name, one or more subnet names and shared project ID.</p> <p>The sharedProjectId value needs to be set in the following way:</p> <ul style="list-style-type: none"> • For a shared VPC, set it to the GCP host project ID • For a non-shared VPC, set it to the GCP project ID of the project where CDP is being deployed.

Parameter	Description
enableTunnel	By default CCM is enabled (set to "true"). If you would like to disable it, set it to "false". If you disable it, then you must also add the following to your JSON definition to specify two security groups as follows: <pre> "securityAccess": { "securityGroupIdForKnox": "string", "defaultSecurityGroupId": "string" } </pre>
usePublicIp	Set this to "true" or "false", depending on whether or not you want to create public IPs.
logStorage	Provide a JSON specifying your configuration for cluster and audit logs: <pre> { "storageLocationBase": "string", "serviceAccountEmail": "string" } </pre> <p>The storageLocationBase should be in the following format: gs://my-bucket-name.</p>



Note: CDP CLI includes the `cdp environments create-gcp-environment --generate-cli-skeleton` command option, which allows you to generate a CLI JSON template. You can also use CLI help to get some information about specific CLI command options.

2. To verify that your environment is running, use:

```
cdp environments list-environments
```

You can also log in to the CDP web interface to check the deployment status.

3. Once your environment and Data Lake are running, you should set IDBroker Mappings. To create the mappings, run the `cdp environments set-id-broker-mappings` command. For example:

```
cdp environments set-id-broker-mappings \
--environment-name test-env \
--data-access-role dl-admin@dev-project.iam.gserviceaccount.com \
--ranger-audit-role ranger-audit@dev-project.iam.gserviceaccount.com \
--mappings '[{"accessorCrn": "crn:altus:iam:us-west-1:45ca3068-42a6-4227-8394-13a4493e2ac0:user:430c534d-8a19-4d9e-963d-8af377d16963", "role": "data-science@dev-project.iam.gserviceaccount.com"}, {"accessorCrn": "crn:altus:iam:us-west-1:45ca3068-42a6-4227-8394-13a4493e2ac0:machineUser:mfox-gcp-idbmms-test-mu/2cbca867-647b-44b9-8e41-47a01dea6c19", "role": "data-eng@dev-project.iam.gserviceaccount.com"}]'
```

Parameter	Description
environment-name	Specify a name of the environment created earlier.
data-access-role	Specify an email address of the Data Lake admin service account created earlier.
ranger-audit-role	Specify an email address of the Ranger audit service account created earlier.

Parameter	Description
mappings	<p>Map CDP users or groups to GCP service accounts created earlier. Use the following syntax:</p> <pre>[{ "accessorCrn": "string", "role": "string" } ...]</pre> <p>You can obtain user or group CRN from the Management Console > User Management by navigating to details of a specific user or group.</p> <p>The role should be specified as service account email.</p>

4. Next, sync IDBroker mappings:

```
cdp environments sync-id-broker-mappings --environment-name demo3
```

5. Finally, check the sync status:

```
cdp environments get-id-broker-mappings-sync-status --environment-name demo3
```

6. Once your environment is running, you can create a Data Lake using the `cdp datalake create-gcp-datalake` command and providing the CLI input parameters:

```
cdp datalake create-gcp-datalake --cli-input-json '{
  "datalakeName": "my-dl",
  "environmentName": "test-env",
  "scale": "LIGHT_DUTY",
  "cloudProviderConfiguration": {
    "serviceAccountEmail": "idbroker@dev-project.iam.gserviceaccount.com",
    "storageLocation": "gs://data-storage"
  }
}'
```

Parameter	Description
datalakeName	Provide a name for your Data Lake.
environmentName	Provide a name of the environment created earlier.
scale	<p>Provide Data Lake scale. It must be one of:</p> <ul style="list-style-type: none"> LIGHT_DUTY or MEDIUM_DUTY_HA.
cloudProviderConfiguration	Provide the name of the data storage bucket and the email of the IDBroker service account.



Note: CDP CLI includes the `cdp datalake create-gcp-datalake --generate-cli-skeleton` command option, which allows you to generate a CLI JSON template. You can also use CLI help to get some information about specific CLI command options.

7. To verify that your Data lake is running, use:

```
cdp datalake list-datalakes
```

You can also log in to the CDP web interface to check the deployment status.

After you finish

After your environment is running, perform the following steps:

- You must assign roles to specific users and groups for the environment so that selected users or user groups can access the environment. Next, you need to perform user sync. For steps, refer to [Enabling admin and user access to environments](#).
- You must onboard your users and/or groups for cloud storage. For steps, refer to [Onboarding CDP users and groups for cloud storage](#).
- You must create Ranger policies for your users. For instructions on how to access your Data Lake, refer to [Accessing Data Lake services](#). Once you've accessed Ranger, [create Ranger policies](#) to determine which users have access to which databases and tables.