Cloudera Management Console

Container Service

Date published: 2019-08-22
Date modified: 2025-08-18

https://docs.cloudera.com/

Legal Notice

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

Using Compute Clusters in environments on AWS....................4
Setting up Compute Cluster IAM permissions....................4
Creating IAM roles and instance profile for EKS....................4
Creating Compute Restricted IAM policy....................10
Enabling default Compute Cluster for new environments....................18
Enabling default Compute Cluster for existing environments....................20
Adding more Compute Clusters....................22
Managing Compute Clusters....................23
Suspending and resuming Compute Clusters....................24
Deleting Compute Clusters....................25
Upgrading Compute Clusters....................25

Using Compute Clusters in environments on Azure....................28
Enabling default Compute Cluster for new environments....................28
Enabling default Compute Cluster for existing environments....................31
Adding more Compute clusters....................32
Managing Compute Clusters....................33
Suspending and resuming Compute Clusters....................34
Deleting Compute Clusters....................35
Upgrading Compute Clusters....................35

Using Compute Clusters in environments on AWS

Cloudera Container Service enables you to deploy a containerized platform on Kubernetes for Data Services and shared services using Compute Clusters.

Cloudera Container Service offers simplified management, enhanced efficiency, and centralized control that leads to faster deployments, reduced configuration errors and improved system reliability. As multiple Data Services can optionally share the same Compute Cluster within the Container Service, it also lowers the cost of ownership.

When registering a Compute Cluster enabled environment, a default Compute Cluster is created. This default Compute Cluster is tied to the environment, and is used for running critical applications. The default Compute Cluster is labeled as **Default Cluster** in your environment to distinguish it from the other Compute Clusters.

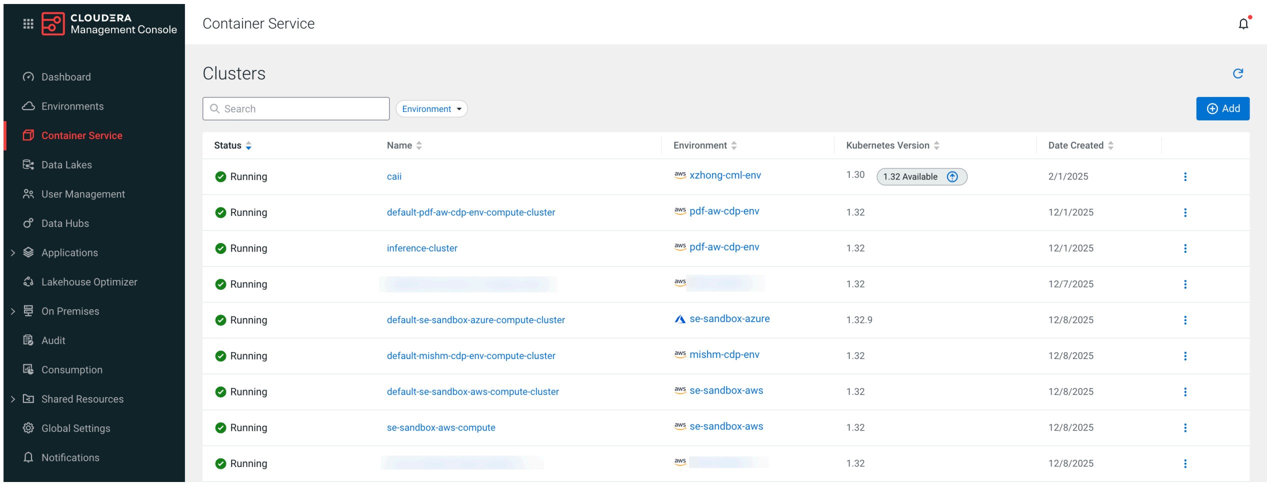

Container Service

Clusters

Status	Name	Environment	Kubernetes Version	Date Created
Running	caii	xzhong-cml-env	1.30 1.32 Available	2/1/2025
Running	default-pdf-aw-cdp-env-compute-cluster	pdf-aw-cdp-env	1.32	12/1/2025
Running	inference-cluster	pdf-aw-cdp-env	1.32	12/1/2025
Running	[redacted]	[redacted]	1.32	12/7/2025
Running	default-se-sandbox-azure-compute-cluster	se-sandbox-azure	1.32.9	12/8/2025
Running	default-mishm-cdp-env-compute-cluster	mishm-cdp-env	1.32	12/8/2025
Running	default-se-sandbox-aws-compute-cluster	se-sandbox-aws	1.32	12/8/2025
Running	se-sandbox-aws-compute	se-sandbox-aws	1.32	12/8/2025
Running	[redacted]	[redacted]	1.32	12/8/2025

You can create additional Compute Clusters that inherit the configuration of the default Compute Cluster. You can manage the lifecycle of the additional Compute Clusters as they are independent of the environment.

Setting up Compute Cluster IAM permissions

Before enabling Compute Clusters for your environment, you need to ensure that the required IAM roles and policies are set up.

Setting up the Compute Cluster IAM permissions are only needed when using **Reduced access policies** detailed on Cross-account access IAM role page, and completing these steps are not required when using default policies.

Creating IAM roles and instance profile for EKS

Complete the steps to create the required IAM roles and profile for EKS.

Procedure

1. Apply the following CloudFormation template to create the following:

   - IAM role called cdp-eks-master-role
   - IAM role and instance profile pair called cdp-liftie-instance-profile

   ```
   AWSTemplateFormatVersion: 2010-09-09
   Description: Creates Liftie IAM resources
   Parameters:
     TelemetryLoggingEnabled:
       Description: Telemetry logging is enabled
       Type: String
   ```

```
  TelemetryLoggingBucket:
    Description: Telemetry logging bucket where Liftie logs will be store
d.
    Type: String
  TelemetryKmsKeyARN:
    Description: KMS Key ARN For Telemetry logging bucket.
    Type: String
    Default: ""
  TelemetryLoggingRootDir:
    Description: Telemetry logging root directory inside telemetry logg
ing bucket used for storing logs.
    Default: "cluster-logs"
    Type: String
Conditions:
  TelemetryLoggingEnabled:
    Fn::Equals:
      - {Ref: TelemetryLoggingEnabled}
      - true
  KMSKeyARNForTelemetryLoggingBucketIsEmpty: !Not [!Equals [!Ref Telemetry
KmsKeyARN, ""]]
Resources:
  AWSServiceRoleForAmazonEKS:
    Type: AWS::IAM::Role
    Properties:
      AssumeRolePolicyDocument:
        Version: 2012-10-17
        Statement:
          - Effect: Allow
            Principal:
              Service:
                - eks.amazonaws.com
            Action:
              - sts:AssumeRole
      ManagedPolicyArns:
        - arn:aws:iam::aws:policy/AmazonEKSServicePolicy
        - arn:aws:iam::aws:policy/AmazonEKSClusterPolicy
      RoleName: cdp-eks-master-role
  NodeInstanceRole:
    Type: AWS::IAM::Role
    Properties:
      AssumeRolePolicyDocument:
        Version: 2012-10-17
        Statement:
          - Effect: Allow
            Principal:
              Service:
                - ec2.amazonaws.com
            Action:
              - sts:AssumeRole
      Path: "/"
      ManagedPolicyArns:
        - arn:aws:iam::aws:policy/AmazonEKSWorkerNodePolicy
        - arn:aws:iam::aws:policy/AmazonEKS_CNI_Policy
        - arn:aws:iam::aws:policy/AmazonEC2ContainerRegistryReadOnly
      RoleName: cdp-liftie-instance-profile
      Policies:
        - PolicyName: ssm-required
          PolicyDocument:
            Version: 2012-10-17
            Statement:
              - Effect: Allow
                Action:
                  - ssm:GetParameters
                Resource:
```

```
         - "*"
    - PolicyName: cluster-autoscaler
     PolicyDocument:
      Version: 2012-10-17
      Statement:
       - Effect: Allow
        Action:
         - autoscaling:DescribeAutoScalingGroups
         - autoscaling:DescribeAutoScalingInstances
         - autoscaling:DescribeLaunchConfigurations
         - autoscaling:DescribeScalingActivities
         - autoscaling:DescribeTags
         - ec2:DescribeImages
         - ec2:DescribeInstanceTypes
         - ec2:DescribeLaunchTemplateVersions
         - ec2:GetInstanceTypesFromInstanceRequirements
         - eks:DescribeNodegroup
        Resource:
         - "*"
       - Effect: Allow
        Action:
         - autoscaling:SetDesiredCapacity
         - autoscaling:TerminateInstanceInAutoScalingGroup
        Resource:
         - "*"
        Condition:
         StringEquals:
          "aws:ResourceTag/k8s.io/cluster-autoscaler/enabled":
"true"
    - PolicyName: ebs-csi
     PolicyDocument:
      Version: 2012-10-17
      Statement:
       - Effect: Allow
        Action:
         - ec2:CreateSnapshot
         - ec2:AttachVolume
         - ec2:DetachVolume
         - ec2:ModifyVolume
         - ec2:DescribeAvailabilityZones
         - ec2:DescribeInstances
         - ec2:DescribeSnapshots
         - ec2:DescribeTags
         - ec2:DescribeVolumes
         - ec2:DescribeVolumesModifications
        Resource: "*"
       - Effect: Allow
        Action:
         - ec2:CreateTags
        Resource:
         - "arn:aws:ec2:*:*:volume/*"
         - "arn:aws:ec2:*:*:snapshot/*"
        Condition:
         StringEquals:
          "ec2:CreateAction":
           - CreateVolume
           - CreateSnapshot
       - Effect: Allow
        Action:
         - ec2:DeleteTags
        Resource:
         - "arn:aws:ec2:*:*:volume/*"
         - "arn:aws:ec2:*:*:snapshot/*"
       - Effect: Allow
```

```
                Action:
                  - ec2:CreateVolume
                Resource: "*"
                Condition:
                  StringLike:
                    "aws:RequestTag/ebs.csi.aws.com/cluster": "true"
              - Effect: Allow
                Action:
                  - ec2:CreateVolume
                Resource: "*"
                Condition:
                  StringLike:
                    "aws:RequestTag/CSIVolumeName": "*"
              - Effect: Allow
                Action:
                  - ec2:CreateVolume
                Resource: "*"
                Condition:
                  StringLike:
                    "aws:RequestTag/kubernetes.io/cluster/*": "owned"
              - Effect: Allow
                Action:
                  - ec2:DeleteVolume
                Resource: "*"
                Condition:
                  StringLike:
                    "ec2:ResourceTag/ebs.csi.aws.com/cluster": "true"
              - Effect: Allow
                Action:
                  - ec2:DeleteVolume
                Resource: "*"
                Condition:
                  StringLike:
                    "ec2:ResourceTag/CSIVolumeName": "*"
              - Effect: Allow
                Action:
                  - ec2:DeleteVolume
                Resource: "*"
                Condition:
                  StringLike:
                    "ec2:ResourceTag/kubernetes.io/created-for/pvc/name":
 "*"
              - Effect: Allow
                Action:
                  - ec2:DeleteSnapshot
                Resource: "*"
                Condition:
                  StringLike:
                    "ec2:ResourceTag/CSIVolumeSnapshotName": "*"
              - Effect: Allow
                Action:
                  - ec2:DeleteSnapshot
                Resource: "*"
                Condition:
                  StringLike:
                    "ec2:ResourceTag/ebs.csi.aws.com/cluster": "true"
        - PolicyName: efs-csi
          PolicyDocument:
            Version: 2012-10-17
            Statement:
              - Effect: Allow
                Action:
                  - elasticfilesystem:DescribeAccessPoints
                  - elasticfilesystem:DescribeFileSystems
                  - elasticfilesystem:DescribeMountTargets
                Resource: "*"
              - Effect: Allow
                Action:
                  - elasticfilesystem:CreateAccessPoint
                Resource: "*"
                Condition:
                  StringLike:
                    "aws:RequestTag/efs.csi.aws.com/cluster": "true"
              - Effect: Allow
                Action:
                  - elasticfilesystem:DeleteAccessPoint
                Resource: "*"
                Condition:
                  StringEquals:
                    "aws:ResourceTag/efs.csi.aws.com/cluster": "true"
        - !If
          - TelemetryLoggingEnabled
          - PolicyName: telemetry-s3-list-bucket
            PolicyDocument:
              Version: 2012-10-17
              Statement:
                - Effect: Allow
                  Action:
                    - s3:ListBucket
                  Resource:
                    - !Sub 'arn:aws:s3:::${TelemetryLoggingBucket}'
                    - !Sub 'arn:aws:s3:::${TelemetryLoggingBucket}/${Tele
metryLoggingRootDir}/*'
          - !Ref 'AWS::NoValue'
        - !If
          - TelemetryLoggingEnabled
          - PolicyName: telemetry-s3-read-write
            PolicyDocument:
              Version: 2012-10-17
              Statement:
                - Effect: Allow
                  Action:
                    - s3:*Object
                    - s3:AbortMultipartUpload
                    - s3:GetBucketAcl
                  Resource:
                    - !Sub 'arn:aws:s3:::${TelemetryLoggingBucket}'
                    - !Sub 'arn:aws:s3:::${TelemetryLoggingBucket}/${Tel
emetryLoggingRootDir}/*'
          - !Ref 'AWS::NoValue'
        - !If
          - KMSKeyARNForTelemetryLoggingBucketIsEmpty
          - PolicyName: s3-kms-read-write-policy
            PolicyDocument:
              Version: 2012-10-17
              Statement:
                - Effect: Allow
                  Action:
                    - kms:CreateGrant
                    - kms:Decrypt
                    - kms:GenerateDataKey
                    - kms:GenerateDataKeyWithoutPlainText
                  Resource:
                    - !Sub ${TelemetryKmsKeyARN}
          - !Ref 'AWS::NoValue'
        - PolicyName: calico-cni
          PolicyDocument:
            Version: 2012-10-17
            Statement:
              - Effect: Allow
                Action:
                  - ec2:ModifyInstanceAttribute
                Resource:
                  - "*"
                Condition:
                  StringEquals:
                    "ec2:Attribute": "SourceDestCheck"
  NodeInstanceProfile:
    Type: AWS::IAM::InstanceProfile
    Properties:
      Path: /
      InstanceProfileName: cdp-liftie-instance-profile
      Roles:
        - !Ref NodeInstanceRole
```

2. In the AWS console Cloudformation wizard, provide values for the following properties:

- Stack Name: Provide an appropriate name. Example: compute-precreated-roles-and-instanceprofile)
- TelemetryLoggingBucket: Name of the log bucket. Example: compute-logging-bucket
- TelemetryLoggingEnabled: Set it to true.
- TelemetryLoggingRootDir: Verify that it is set to the default value cluster-logs.
- TelemetryKMSKeyARN: If the telemetry bucket is encrypted, specify the KMS Key ARN. The default value is null.

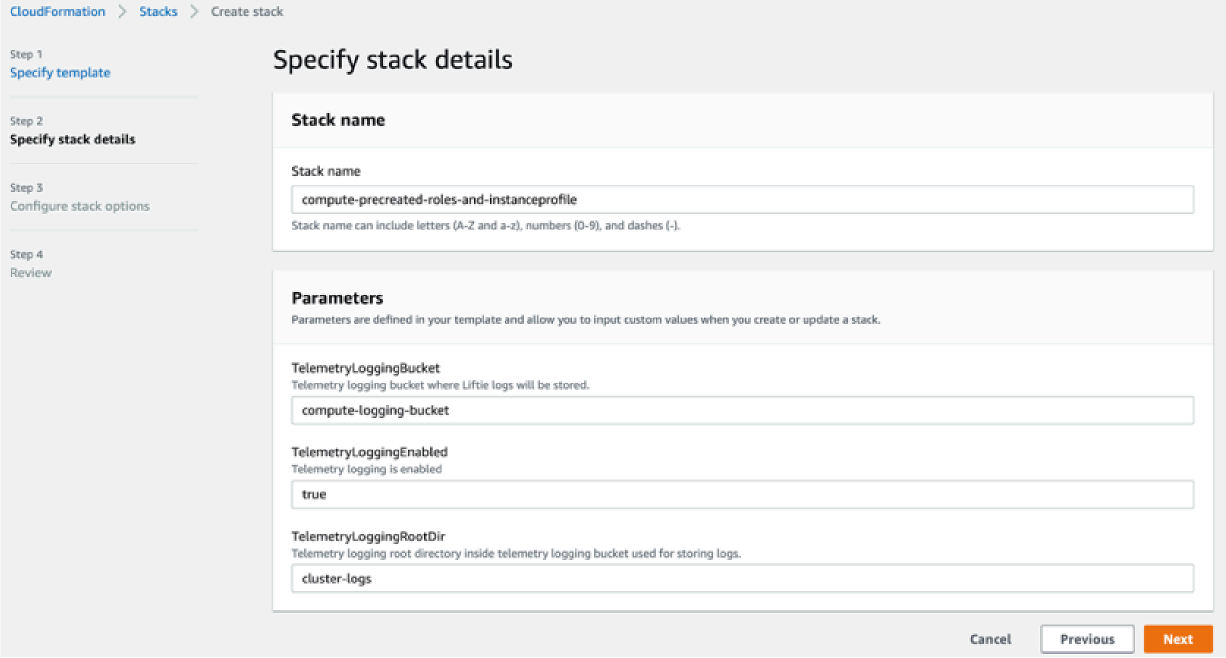

Specify stack details

Stack name

Stack name: compute-precreated-roles-and-instanceprofile

Stack name can include letters (A-Z and a-z), numbers (0-9), and dashes (-).

Parameters

Parameters are defined in your template and allow you to input custom values when you create or update a stack.

TelemetryLoggingBucket
Telemetry logging bucket where Liftie logs will be stored.
compute-logging-bucket

TelemetryLoggingEnabled
Telemetry logging is enabled
true

TelemetryLoggingRootDir
Telemetry logging root directory inside telemetry logging bucket used for storing logs.
cluster-logs

3. On the last page in the wizard process, click the I acknowledge... checkbox to allow creation of IAM resources with special names.

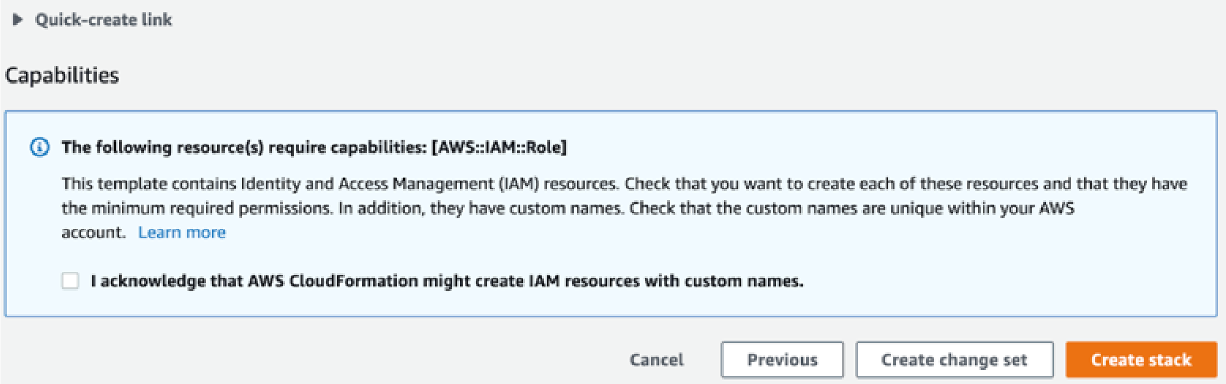

The following resource(s) require capabilities: [AWS::IAM::Role]

This template contains Identity and Access Management (IAM) resources. Check that you want to create each of these resources and that they have the minimum required permissions. In addition, they have custom names. Check that the custom names are unique within your AWS account. Learn more

□ **I acknowledge that AWS CloudFormation might create IAM resources with custom names.**

4. Click Create stack.

Results

On the Cloudformation **Resources** tab, you find the precreated role and instance profile.

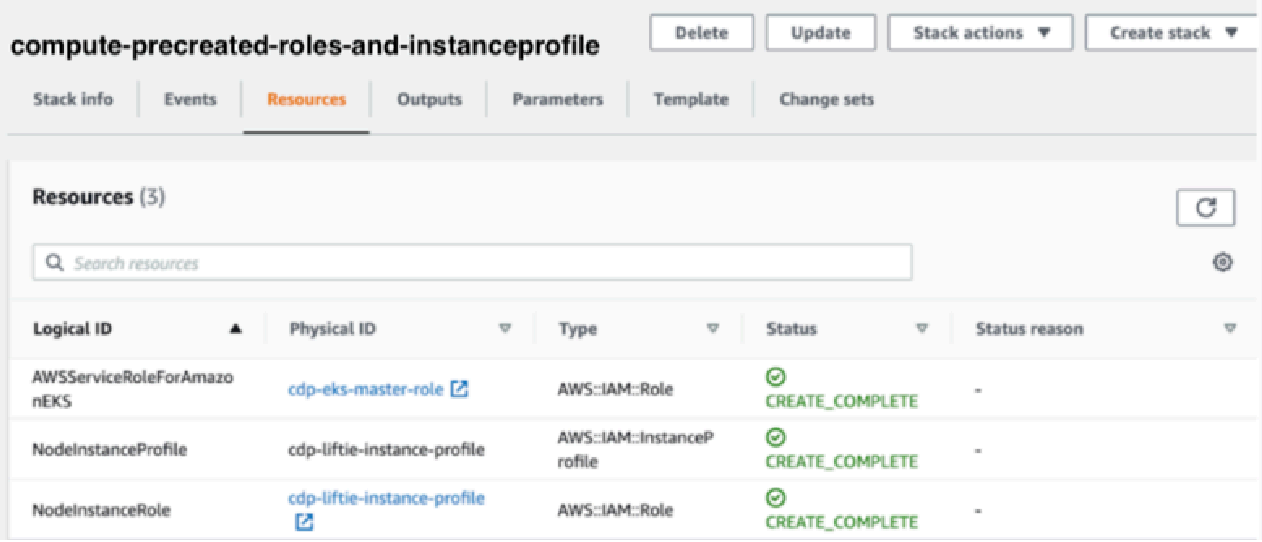

Logical ID ▲	Physical ID	Type	Status	Status reason
AWSServiceRoleForAmazonEKS	cdp-eks-master-role	AWS::IAM::Role	CREATE_COMPLETE	-
NodeInstanceProfile	cdp-liftie-instance-profile	AWS::IAM::InstanceProfile	CREATE_COMPLETE	-
NodeInstanceRole	cdp-liftie-instance-profile	AWS::IAM::Role	CREATE_COMPLETE	-

What to do next

Update the environment role to use the restricted role and policy.

Creating Compute Restricted IAM policy

Complete the steps to attach the Compute Restricted IAM policy with the cross-account role associated with your environment.

Procedure

1. Go to the **Environments** page.

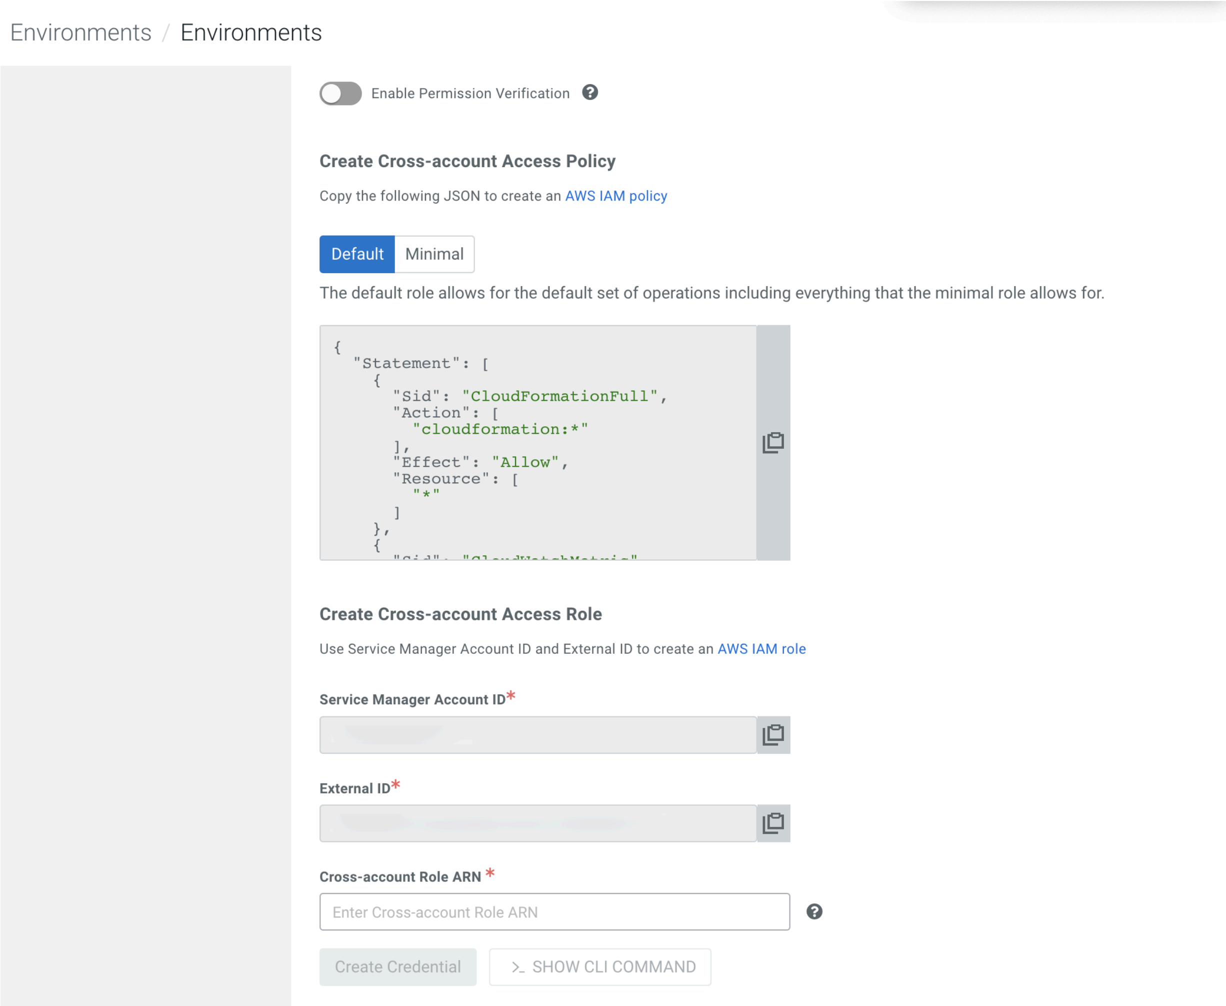

Environments / Environments

Enable Permission Verification

Create Cross-account Access Policy

Copy the following JSON to create an AWS IAM policy

Default | Minimal

The default role allows for the default set of operations including everything that the minimal role allows for.

```
{
  "Statement": [
    {
      "Sid": "CloudFormationFull",
      "Action": [
        "cloudformation:*"
      ],
      "Effect": "Allow",
      "Resource": [
        "*"
      ]
    },
    {
```

Create Cross-account Access Role

Use Service Manager Account ID and External ID to create an AWS IAM role

Service Manager Account ID*: [redacted]

External ID*: [redacted]

Cross-account Role ARN *: Enter Cross-account Role ARN

2. In the Create Cross-account Access Policy field, attach the Compute Restricted IAM policy:

   Replace the following placeholders in the JSON file:

   - *[YOUR-ACCOUNT-ID]* with your account ID in use.
   - *[YOUR-IAM-ROLE-NAME]* with the IAM restricted role associated with this policy.
   - *[YOUR-SUBNET-ARN-*]* supplied during the Cloudera Environment(s) creation.

     Note: Provide all the subnets present in all the Cloudera Environment(s) that you intend to use it for the experience. If at any point a new Cloudera Environment is created or an existing one is updated for subnets, provide it here.
   - *[YOUR-IDBROKER-ROLE-NAME]* with the ID Broker Role name in use.
   - *[YOUR-LOG-ROLE-NAME]* with the Log Role name in use.
   - *[YOUR-KMS-CUSTOMER-MANAGED-KEY-ARN]* with KMS key ARN.

```
{
  "Version": "2012-10-17",
  "Id": "ComputePolicy_v12",
  "Statement": [
    {
      "Sid": "SimulatePrincipalPolicy",
      "Effect": "Allow",
      "Action": [
        "iam:SimulatePrincipalPolicy"
      ],
      "Resource": [
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/[YOUR-IAM-ROLE-NAME]"
      ]
    },
    {
      "Sid": "RestrictedPermissionsViaClouderaRequestTag",
      "Effect": "Allow",
      "Action": [
        "cloudformation:CreateStack",
        "cloudformation:CreateChangeSet",
        "ec2:createTags",
        "eks:TagResource"
      ],
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "aws:RequestTag/Cloudera-Resource-Name": [
            "crn:cdp:*"
          ]
        }
      }
    },
    {
      "Sid": "RestrictedPermissionsViaClouderaResourceTag",
      "Effect": "Allow",
      "Action": [
        "autoscaling:DeleteTags",
        "autoscaling:DetachInstances",
        "autoscaling:ResumeProcesses",
        "autoscaling:SetDesiredCapacity",
        "autoscaling:SuspendProcesses",
        "autoscaling:TerminateInstanceInAutoScalingGroup",
        "autoscaling:UpdateAutoScalingGroup",
        "cloudformation:DeleteChangeSet",
        "cloudformation:DeleteStack",
        "cloudformation:DescribeChangeSet",
        "cloudformation:DescribeStacks",
        "cloudformation:CancelUpdateStack",
        "cloudformation:ContinueUpdateRollback",
        "cloudformation:DescribeStackEvents",
        "cloudformation:DescribeStackResource",
        "cloudformation:DescribeStackResources",
        "cloudformation:ExecuteChangeSet",
        "cloudformation:ListStacks",
        "cloudwatch:deleteAlarms",
        "cloudwatch:putMetricAlarm",
        "ec2:AttachVolume",
        "ec2:CreateNetworkInterface",
        "ec2:CreateVolume",
        "ec2:DeleteVolume",
        "ec2:RunInstances",
        "eks:DescribeUpdate",
        "eks:ListUpdates",
        "eks:UpdateClusterConfig",
        "eks:UpdateClusterVersion",
        "iam:GetRolePolicy",
        "iam:ListInstanceProfiles",
        "iam:ListRoleTags",
        "iam:RemoveRoleFromInstanceProfile",
        "iam:TagRole",
        "iam:UntagRole",
        "logs:DescribeLogStreams",
        "logs:FilterLogEvents"
      ],
      "Resource": "*",
      "Condition": {
        "StringLike": {
          "aws:ResourceTag/Cloudera-Resource-Name": [
            "crn:cdp:*"
          ]
        }
      }
    },
    {
      "Sid": "RestrictedPermissionsViaCloudFormation",
      "Effect": "Allow",
      "Action": [
        "autoscaling:CreateAutoScalingGroup",
        "autoscaling:CreateLaunchConfiguration",
        "autoscaling:CreateOrUpdateTags",
        "autoscaling:DeleteAutoScalingGroup",
        "autoscaling:DeleteLaunchConfiguration",
        "autoscaling:DescribeAutoScalingInstances",
        "autoscaling:DescribeLaunchConfigurations",
        "autoscaling:DescribeScalingActivities",
        "autoscaling:DescribeScheduledActions",
        "autoscaling:DescribeTags",
        "ec2:AuthorizeSecurityGroupEgress",
        "ec2:CreateLaunchTemplate",
        "ec2:CreateSecurityGroup",
        "ec2:DeleteLaunchTemplate",
        "ec2:DeletePlacementGroup",
        "ec2:DeleteSecurityGroup",
        "ec2:DescribeAccountAttributes",
        "ec2:DescribeImages",
        "ec2:DescribeInstanceStatus",
        "ec2:DescribeInstances",
        "ec2:DescribeKeyPairs",
        "ec2:DescribeLaunchTemplateVersions",
        "ec2:DescribeLaunchTemplates",
        "ec2:DescribePlacementGroups",
        "ec2:DescribeRegions",
        "ec2:DescribeRouteTables",
        "ec2:DescribeSecurityGroups",
        "ec2:DescribeVolumes",
        "ec2:RevokeSecurityGroupEgress",
        "ec2:RevokeSecurityGroupIngress",
        "eks:CreateCluster",
        "eks:DeleteCluster"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:CalledVia": [
            "cloudformation.amazonaws.com"
          ]
        }
      }
    },
    {
      "Sid": "RestrictedEC2PermissionsViaClouderaResourceTag",
      "Effect": "Allow",
      "Action": [
        "ec2:RebootInstances",
        "ec2:StartInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "ForAnyValue:StringLike": {
          "ec2:ResourceTag/Cloudera-Resource-Name": [
            "crn:cdp:*"
          ]
        }
      }
    },
    {
      "Sid": "RestrictedIamPermissionsToClouderaResources",
      "Effect": "Allow",
      "Action": [
        "iam:PassRole"
      ],
      "Resource": [
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/[YOUR-IDBROKER-ROLE-NAME]",
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/[YOUR-LOG-ROLE-NAME]",
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/liftie-*-eks-service-role",
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/liftie-*-eks-worker-nodes",
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/cdp-eks-master-role",
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/cdp-liftie-instance-profile"
      ]
    },
    {
      "Sid": "RestrictedKMSPermissionsUsingCustomerProvidedKey",
      "Effect": "Allow",
      "Action": [
        "kms:CreateGrant",
        "kms:DescribeKey",
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*"
      ],
      "Resource": [
        "[YOUR-KMS-CUSTOMER-MANAGED-KEY-ARN]"
      ]
    },
    {
      "Sid": "AllowCreateDeleteTagsForSubnets",
      "Effect": "Allow",
      "Action": [
        "ec2:CreateTags",
        "ec2:DeleteTags"
      ],
      "Resource": [
        "arn:aws:ec2:[YOUR-SUBNET-REGION]:[YOUR-ACCOUNT-ID]:subnet/*"
      ]
    },
    {
      "Sid": "ModifyInstanceAttribute",
      "Effect": "Allow",
      "Action": [
        "ec2:ModifyInstanceAttribute"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "StringEquals": {
          "ec2:Attribute": "SourceDestCheck"
        }
      }
    },
    {
      "Sid": "OtherPermissions",
      "Effect": "Allow",
      "Action": [
        "autoscaling:DescribeAutoScalingGroups",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateLaunchTemplateVersion",
        "ec2:CreatePlacementGroup",
        "ec2:DeleteKeyPair",
        "ec2:DeleteNetworkInterface",
        "ec2:DescribeAvailabilityZones",
        "ec2:DescribeInstanceTypes",
        "ec2:DescribeNetworkInterfaces",
        "ec2:DescribeSubnets",
        "ec2:DescribeVpcAttribute",
        "ec2:DescribeVpcs",
        "ec2:ImportKeyPair",
        "ec2:UpdateSecurityGroupRuleDescriptionsIngress",
        "ec2:GetInstanceTypesFromInstanceRequirements",
        "eks:DescribeCluster",
        "elasticloadbalancing:DescribeLoadBalancers",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:GetInstanceProfile"
      ],
      "Resource": [
        "*"
      ]
    },
    {
      "Sid": "AllowSsmParams",
      "Effect": "Allow",
      "Action": [
        "ssm:DescribeParameters",
        "ssm:GetParameter",
        "ssm:GetParameters",
        "ssm:GetParameterHistory",
        "ssm:GetParametersByPath"
      ],
      "Resource": [
        "arn:aws:ssm:*:*:parameter/aws/service/eks/optimized-ami/*"
      ]
    },
    {
      "Sid": "CfDeny",
      "Effect": "Deny",
      "Action": [
        "cloudformation:*"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "ForAnyValue:StringLike": {
          "cloudformation:ImportResourceTypes": [
            "*"
          ]
        }
      }
    },
    {
      "Sid": "ForAutoscalingLinkedRole",
      "Effect": "Allow",
      "Action": [
        "iam:CreateServiceLinkedRole"
      ],
      "Resource": [
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/aws-service-role/autoscaling-
plans.amazonaws.com/AWSServiceRoleForAutoScalingPlans_EC2AutoScaling"
      ],
      "Condition": {
        "StringLike": {
          "iam:AWSServiceName": "autoscaling-plans.amazonaws.com"
        }
      }
    },
    {
      "Sid": "ForEksLinkedRole",
      "Effect": "Allow",
      "Action": [
        "iam:CreateServiceLinkedRole"
      ],
      "Resource": [
        "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/aws-service-role/eks.amazo
naws.com/AWSServiceRoleForEKS"
      ],
      "Condition": {
        "StringLike": {
          "iam:AWSServiceName": "eks.amazonaws.com"
        }
      }
    }
  ]
}
```

3. Provide and verify your Customer Managed Key (CMK) to be used for EBS encryption.

   Along with providing the KMS Customer Managed Key (CMK) for volume encryption in the policy section with Sid: RestrictedKMSPermissionsUsingCustomerProvidedKey, you need to verify that the policy for the Customer Managed Key (CMK) at KMS (this is not an IAM policy) has the following three permission blocks defined for AWSServiceRoleForAutoScaling.

```
{
  "Statement": [
    {
      "Sid": "AllowAutoscalingServiceLinkedRoleForAttachmentOfPersisten
tResources",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/aws-service-role/auto
scaling.amazonaws.com/AWSServiceRoleForAutoScaling"
      },
      "Action": "kms:CreateGrant",
      "Resource": "*",
      "Condition": {
        "Bool": {
          "kms:GrantIsForAWSResource": "true"
        }
      }
    },
    {
      "Sid": "AllowAutoscalingServiceLinkedRoleUseOfTheCMK",
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::[YOUR-ACCOUNT-ID]:role/aws-service-role/autos
caling.amazonaws.com/AWSServiceRoleForAutoScaling"
      },
      "Action": [
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey"
      ],
      "Resource": "*"
    },
    {
      "Sid": "Allow EKS access to EBS.",
      "Effect": "Allow",
      "Principal": {
        "AWS": "*"
      },
      "Action": [
        "kms:CreateGrant",
        "kms:Encrypt",
        "kms:Decrypt",
        "kms:ReEncrypt*",
        "kms:GenerateDataKey*",
        "kms:DescribeKey"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "kms:CallerAccount": "[YOUR-ACCOUNT-ID]",
          "kms:viaService": "ec2.[YOUR-ACCOUNT-REGION].amazonaws.com"
        }
      }
    }
  ]
```

```
}
```

After the policy is attached, the KMS service page will show the CMS as having the policy attached as shown in the following example:

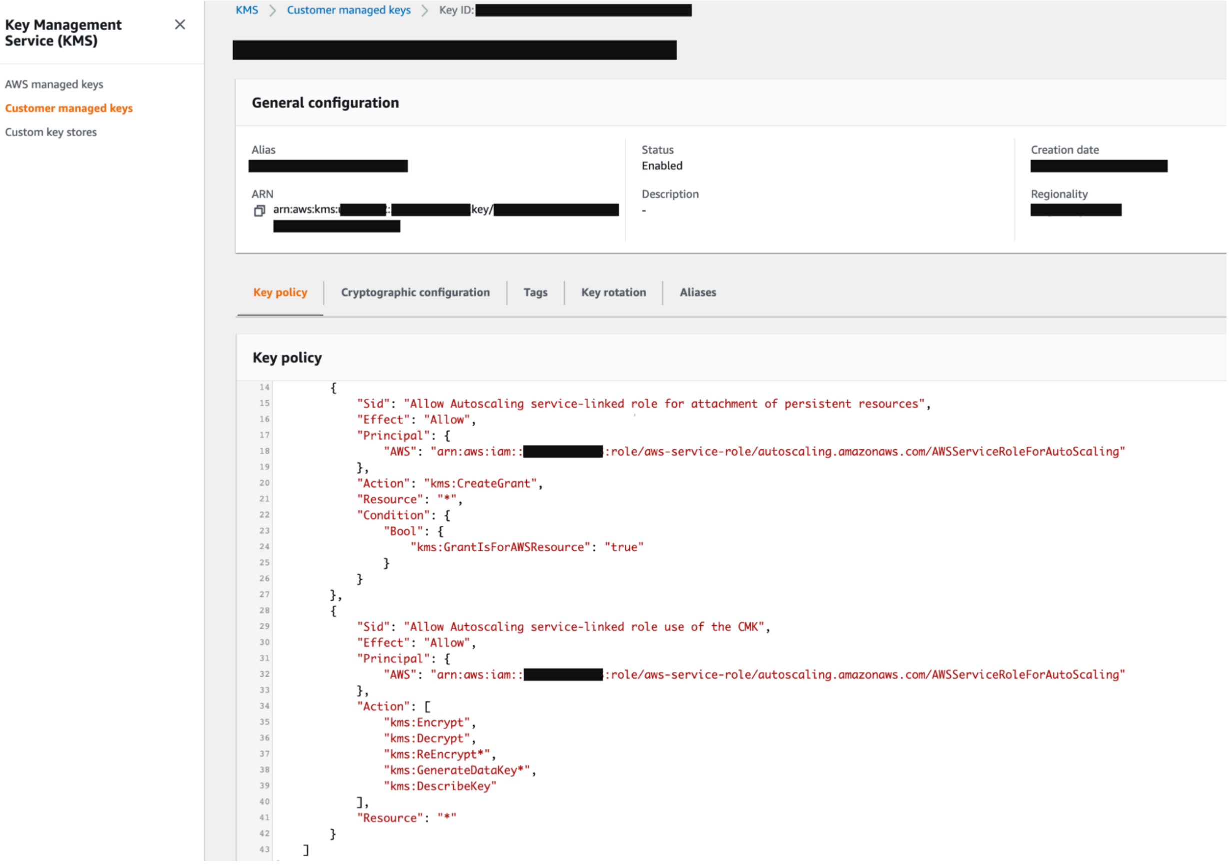
Key Management Service (KMS)
AWS managed keys
Customer managed keys
Custom key stores
KMS > Customer managed keys > Key ID: [redacted]
Key policy
Cryptographic configuration
Tags
Key rotation
Aliases

[redacted]

General configuration

Alias
[redacted]

ARN
arn:aws:kms:[redacted]:[redacted]key/[redacted]
[redacted]

Status
Enabled

Description
-

Creation date
[redacted]

Regionality
[redacted]

Key policy

```
        {
            "Sid": "Allow Autoscaling service-linked role for attachment of persistent resources",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::[redacted]:role/aws-service-role/autoscaling.amazonaws.com/AWSServiceRoleForAutoScaling"
            },
            "Action": "kms:CreateGrant",
            "Resource": "*",
            "Condition": {
                "Bool": {
                    "kms:GrantIsForAWSResource": "true"
                }
            }
        },
        {
            "Sid": "Allow Autoscaling service-linked role use of the CMK",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::[redacted]:role/aws-service-role/autoscaling.amazonaws.com/AWSServiceRoleForAutoScaling"
            },
            "Action": [
                "kms:Encrypt",
                "kms:Decrypt",
                "kms:ReEncrypt*",
                "kms:GenerateDataKey*",
                "kms:DescribeKey"
            ],
            "Resource": "*"
        }
    ]
```

Enabling default Compute Cluster for new environments

When creating your environment, you can enable the default Compute Cluster using the Cloudera Management Console or CDP CLI to be able to run your data and shared services on the containerized platform.

Required role: EnvironmentAdmin

Before you begin

- Ensure that your AWS account has all the resources required by Cloudera.

  For more information, see AWS account requirements.
- Ensure that the IAM permissions are correctly set up for your environment.

  For more information, see Setting up Compute Cluster IAM permissions.

Using Cloudera Management Console

When creating your environment in **Cloudera Management Console**, ensure that you use the **Enable Compute Cluster** setting to create the Compute Cluster enabled environment.

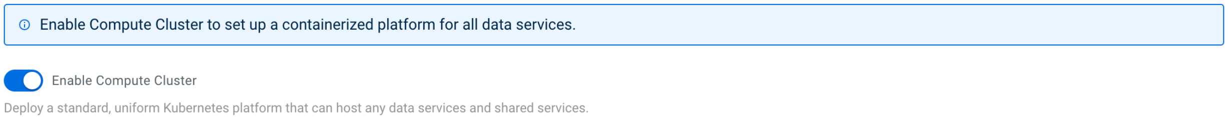

ⓘ Enable Compute Cluster to set up a containerized platform for all data services.

Enable Compute Cluster

Deploy a standard, uniform Kubernetes platform that can host any data services and shared services.

After completing the step for **Data Access and Data Lake Scaling**, configure the networking settings for Kubernetes with either selecting the **Private Kubernetes Cluster** or providing **Authorized IP Ranges** on the **Region, Networking and Security** page. **Worker Node Subnets** are automatically pre-filled with the same set of subnets provided in **Network** section, but you have the option to not use all of the available subnets.

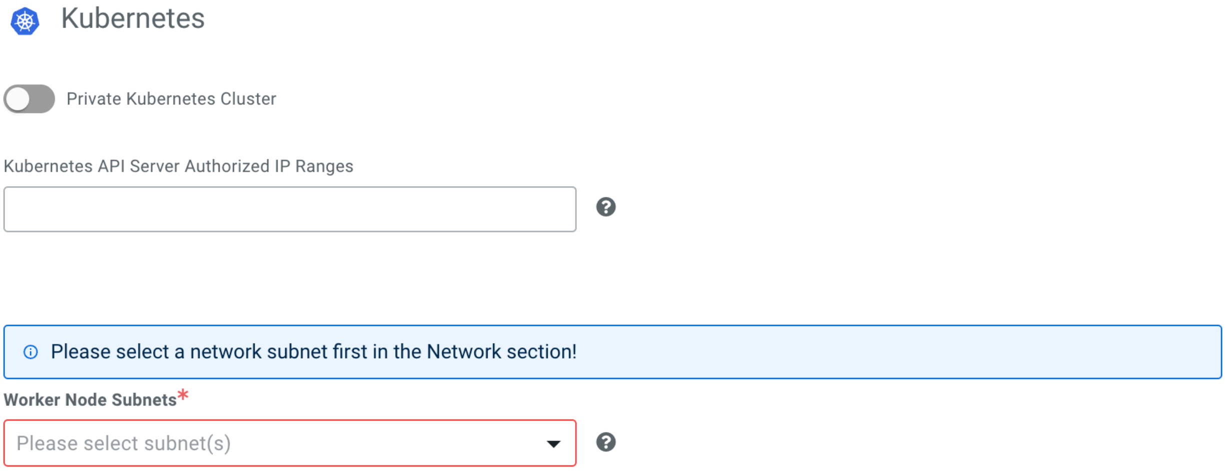

Kubernetes

Private Kubernetes Cluster

Kubernetes API Server Authorized IP Ranges

ⓘ Please select a network subnet first in the Network section!

Worker Node Subnets*

Please select subnet(s)

For more information about creating your environment, see the Register environment (UI) documentation.

Using CDP CLI

Run the following command to create the Compute Cluster enabled environment:

For Without private cluster

```
cdp environments create-aws-environment
--environment-name [***ENVIRONMENT NAME***] \
--credential-name [***CREDENTIAL NAME***] \
--region [***REGION***] \
--security-access [***SECURITY CONTROL CONFIGURATIONS***] \
--authentication [***PUBLIC SSH KEY***] \
--log-storage [***STORAGE CONFIGURATION***] \
--enable-compute-cluster \
--compute-cluster-configuration \
privateCluster=FALSE, \
kubeApiAuthorizedIpRanges=[***CIDR1***],[***CIDR2***]
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

For With private cluster

```
cdp environments create-aws-environment
--environment-name [***ENVIRONMENT NAME***] \
--credential-name [***CREDENTIAL NAME***] \
--region [***REGION***] \
--security-access [***SECURITY CONTROL CONFIGURATIONS***] \
--authentication [***SSH KEY***] \
--log-storage [***STORAGE CONFIGURATION***] \
--enable-compute-cluster \
--compute-cluster-configuration \
privateCluster=TRUE
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

After the command runs, you can verify if the environment was successfully created with the default Compute Cluster with using the following commands:

- Describing the environment:

```
cdp environments describe-environment --env-name-or-crn [***ENVIRONMENT
 NAME OR CRN***]

...
        "awsComputeClusterConfiguration": {
            "privateCluster": false,
            "kubeApiAuthorizedIpRanges": [
                "0.0.0.0/0"
            ]
        },
        "enableComputeCluster": "true"
...
```

- Listing Compute Clusters:

```
cdp compute list-clusters --env-name-or-crn [***ENVIRONMENT NAME OR
 CRN***]
```

Note: FreeIPA must be created and running before the default Compute Cluster is created.

You can use the following command to retry the environment creation with the default Compute Cluster:

For Without private cluster

```
cdp environments initialize-aws-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=FALSE, \
kubeApiAuthorizedIpRanges=[***CIDR1***],[***CIDR2***]
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

For With private cluster

```
cdp environments initialize-aws-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=TRUE
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

Enabling default Compute Cluster for existing environments

In case you already have an environment, but would like to have your services to run on the containerized platform enabled by Compute Clusters, you can add the default Compute Cluster to your existing environment.

Required resource role: EnvironmentAdmin

Before you begin

- Ensure that the environment has no default Compute Cluster provisioned.
- Ensure that the environment is started and available.
- Ensure that IAM permissions are correctly set up during the environment creation.

  For more information, see Setting up Compute Cluster IAM permissions.

Using Cloudera Management Console

1. Navigate to your environment.
2. Click Compute Clusters.

   The following message will indicate that Compute Cluster are not enabled for your environment:

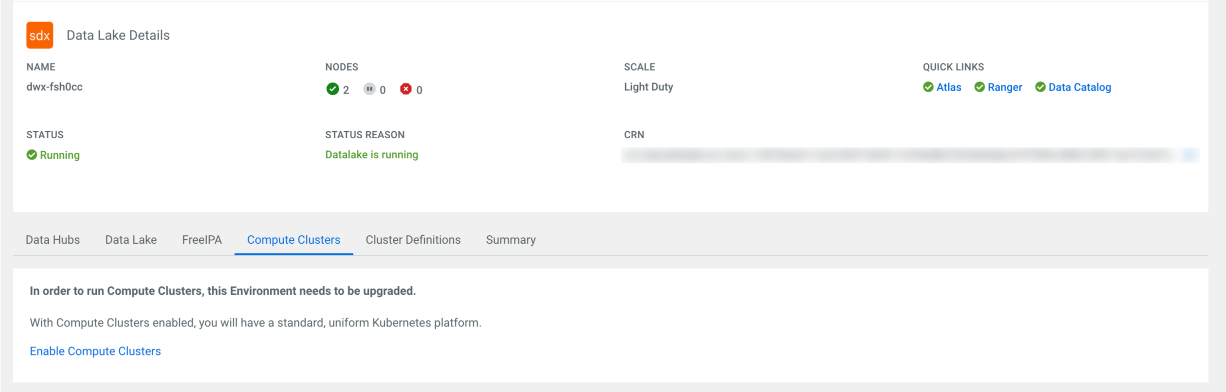

sdx Data Lake Details

NAME: dwx-fsh0cc

NODES: ✓ 2 ⏸ 0 ✗ 0

SCALE: Light Duty

QUICK LINKS: ✓ Atlas ✓ Ranger ✓ Data Catalog

STATUS: ✓ Running

STATUS REASON: Datalake is running

CRN: [redacted]

In order to run Compute Clusters, this Environment needs to be upgraded.

With Compute Clusters enabled, you will have a standard, uniform Kubernetes platform.

Enable Compute Clusters

3. Click Enable Compute Clusters.
4. Provide the necessary networking information for the **Kubernetes** cluster.

   a. If you need to create a Private Cluster, enable Private Kubernetes Cluster to create a private cluster that blocks all access to the API Server endpoint.

   b. If you do not need to create a Private Cluster, 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.

   c. **Worker Node Subnets** are automatically pre-filled with the same set of subnets provided during environment registration, but you have the option to not use all of the available subnets.
5. Click Submit.

   You will be redirected to the **Compute Clusters** tab, where you can track the creation process of the default Compute Cluster.

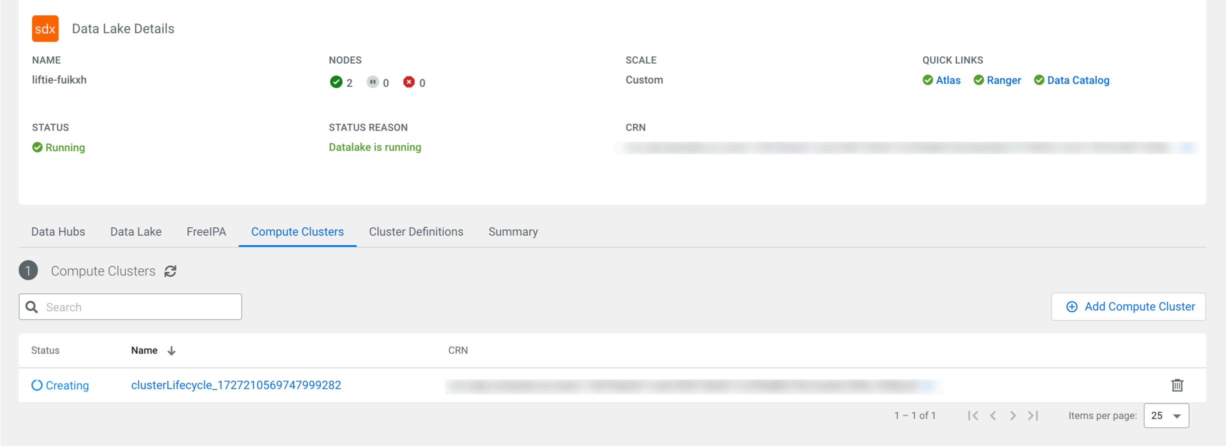

sdx Data Lake Details

NAME: liftie-fuikxh

NODES: ✓ 2 ⏸ 0 ✗ 0

SCALE: Custom

QUICK LINKS: ✓ Atlas ✓ Ranger ✓ Data Catalog

STATUS: ✓ Running

STATUS REASON: Datalake is running

CRN: [redacted]

Status	Name ↓	CRN
Creating	clusterLifecycle_1727210569747999282	[redacted]

Using CDP CLI

Run the following command to add the default Compute Cluster to the environment:

For Without private cluster

```
cdp environments initialize-aws-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=FALSE, \
kubeApiAuthorizedIpRanges=[***CIDR1***],[***CIDR2***]
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

For With private cluster

```
cdp environments initialize-aws-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=TRUE
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

The environment will have COMPUTE_CLUSTER_CREATION_IN_PROGRESS status. You can use the following command to check the status of the environment creation, the statusReason field will contain the information about the process:

```
cdp environments describe-environment --env-name-or-crn [***ENVIRONMENT NAME
 OR CRN***]
```

For more detailed status information about the cluster creation, you can use the following command:

```
cdp compute list-clusters --env-name-or-crn [***ENVIRONMENT NAME OR CRN***]
```

Adding more Compute Clusters

You can add as many additional Compute Clusters as required beside the default Compute Cluster using Cloudera Management Console or CDP CLI.

Required role: EnvironmentAdmin

Using Cloudera Management Console

You can create additional Compute Clusters beside the default Compute Cluster using Cloudera Management Console.

1. Navigate to your environment.
2. Select Compute Clusters tab.

3. Click Add Compute Cluster.

   The **Add Compute Cluster** wizard appears.

   Add Compute Cluster

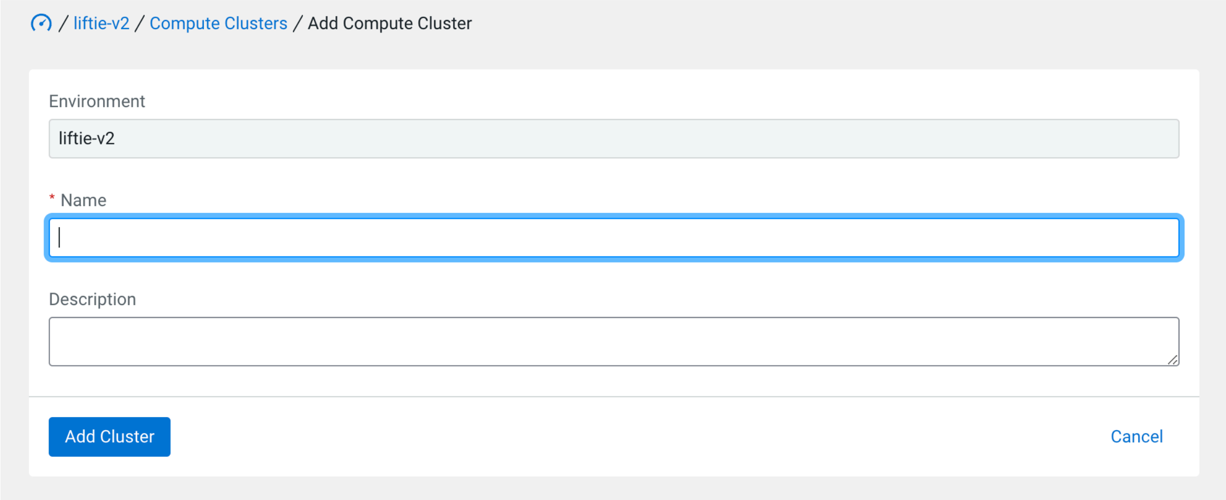

/ liftie-v2 / Compute Clusters / Add Compute Cluster

Environment: liftie-v2

4. Provide a Name to the cluster, and optionally a Description.
5. Click Add Cluster.

   You will be redirected to the **Compute Clusters** tab, where you can track the creation process of the additional Compute Cluster.

Using CDP CLI

You can use the following CDP CLI command to create additional Compute Clusters after the default Compute Cluster is created:

```
cdp compute create-cluster
--environment --env-name-or-crn [***ENVIRONMENT NAME OR CRN***] \
--name [***CLUSTER NAME***]
```

After the command runs, you can verify if the additional Compute Cluster creation was successful using the following command:

```
cdp compute list-clusters --env-name-or-crn [***ENVIRONMENT NAME OR CRN***]
```

The additional Compute Cluster status should be in RUNNING state.

Managing Compute Clusters

After creation, you can view the Compute Cluster details, manage the access of the clusters, download the Kubeconfig file, and delete the created additional Compute Clusters.

Required resource role: EnvironmentAdmin or Owner

Required account role: IamViewer

Using Cloudera Management Console

1. Click on the Name of the Compute Cluster.

   You will be redirected to the **Cluster Details** page.

   On the **Cluster Details** page, you can view the **Status**, **Creation Date**, **Created By**, **Description** and **CRN** of the Compute Cluster.
2. Click Actions to open the drop-down menu.

   Note: Ensure that you have one of the following roles to access the **Actions** menu

   - EnvironmentOwner or EnvironmentAdmin
   - ClusterCreator or ClusterOwner

   a. Click Manage Access to update the list of users who have access to manage the Compute Cluster and use it for installing Data Services.

      1. Click Update roles to update the **Resource Role** of a user or group.

         You can assign the Owner resource role to a user or group to provide permission to manage the Compute Cluster and install services on it.

   b. Click Kubernetes Access to grant access for users to the Kubernetes API server.

      1. Enter the User ARN.
      2. Click Grant Access.

      You also have the option to Download the Kubeconfig from this page.

   c. Click Delete Compute Cluster, if you no longer need the additional Compute Cluster.

      Warning: Ensure that the data services running on the Compute Cluster are deleted before deleting the Compute Cluster itself.

      1. Confirm the deletion of the additional Compute Cluster by clicking Remove.
3. Click Networking tab to view the subnet information of the Compute Cluster.

   a. Click ✎ to update the **Kubernetes API Server Authorized IP Ranges**.

      Note: Ensure that you have one of the following roles to access the **Actions** menu

      - EnvironmentOwner or EnvironmentAdmin
      - ClusterCreator or ClusterOwner

      1. Add or remove the CIDRs, and click Save.
4. Click **Encryption** tab to view the encryption key of the Compute Cluster.
5. Click **Node Groups** tab to have an illustrated overview of the resource utilization of the different services running on the Compute Cluster.
6. Click **Compute Cluster Version** to view the Kubernetes and Insfrastructure Service versions of the cluster.
7. Click **Logs** to check the different events of the Compute Cluster.

Suspending and resuming Compute Clusters

You can suspend and resume your default and additional Compute Clusters in an environment to manage your cloud costs using CDP CLI.

Before you begin

- Ensure that the environment is started and available.
- Ensure that there are no Data Service workload instances running in the Compute Cluster.

Required resource role: EnvironmentAdmin or Owner

Procedure

1. Run the following command to suspend the Compute Cluster:

   ```
   cdp compute suspend-cluster --cluster-crn [***CLUSTER CRN***]
   ```

   After running the command, the cluster status changes to SUSPENDING, and will progress to SUSPENDED once the operation completes. The status of the cluster can be polled using describe-cluster command:

   ```
   cdp compute describe-cluster --cluster-crn [***CLUSTER CRN***]
   ```

2. Run the following command to resume the Compute Cluster with SUSPENDED status:

   ```
   cdp compute resume-cluster --cluster-crn [***CLUSTER CRN***]
   ```

   After running the command, the cluster status changes to RESUMING, and will progress to RUNNING status once the operation completes.

   Note: Only the following operations are supported on a cluster with SUSPENDED status:

   - DescribeCluster
   - ListClusters
   - DeleteCluster

Deleting Compute Clusters

You can delete additional Compute Clusters in an environment to manage your cloud costs using CDP CLI.

Before you begin:

- Ensure that the environment is started and available.
- Ensure that there are no Data Service workload instances running in the Compute Cluster.
- You can only delete additional Compute Clusters.
- You can only delete the default Cluster by deleting the environment.

Run the following command to delete the Compute Cluster:

```
cdp compute delete-cluster --cluster-crn [***CLUSTER CRN***]
```

Upgrading Compute Clusters

You can upgrade default and additional Compute Clusters in an environment when there is a new Kubernetes version available from Cloudera Management Console and from CDP CLI.

Required role: EnvironmentAdmin or EnvironmentOwner

Before you begin

- Ensure that the Cloudera environment is running and available..

Using Cloudera Management Console

1. Navigate to Container Service in Cloudera Management Console.

2. Find the compute cluster that you want to upgrade.

   You can scroll through the list of compute clusters or use the search box to filter down the list of clusters.

   In case there is an upgrade available for the compute cluster, the new version is indicated next to the Kubernetes Version.

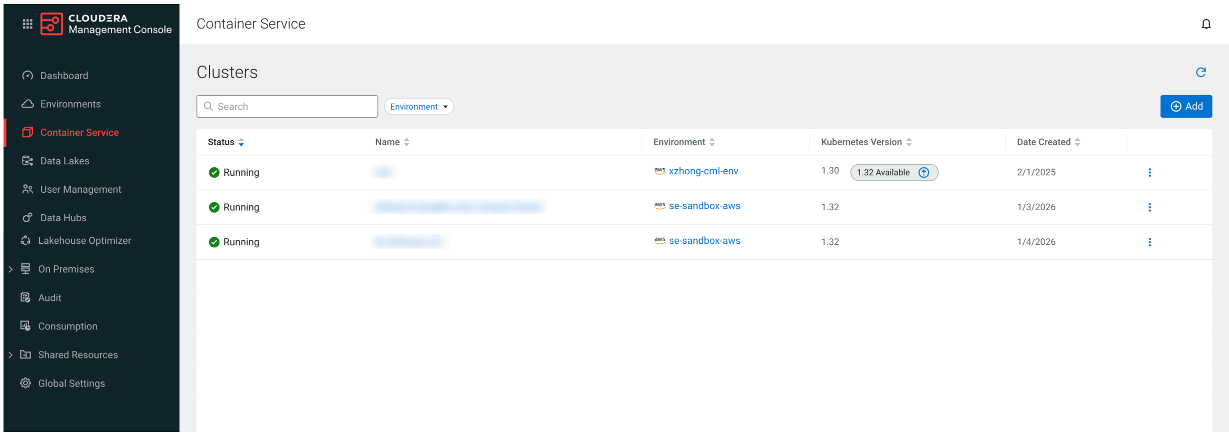

Clusters

Status	Name	Environment	Kubernetes Version	Date Created
Running	[redacted]	xzhong-cml-env	1.30 1.32 Available	2/1/2025
Running	[redacted]	se-sandbox-aws	1.32	1/3/2026
Running	[redacted]	se-sandbox-aws	1.32	1/4/2026

3. Start the upgrade process by using one of the following steps:

   a. Click on the cluster that you want to upgrade. You will be redirected to the **Compute Cluster Details** page, and select Actions Upgrade .

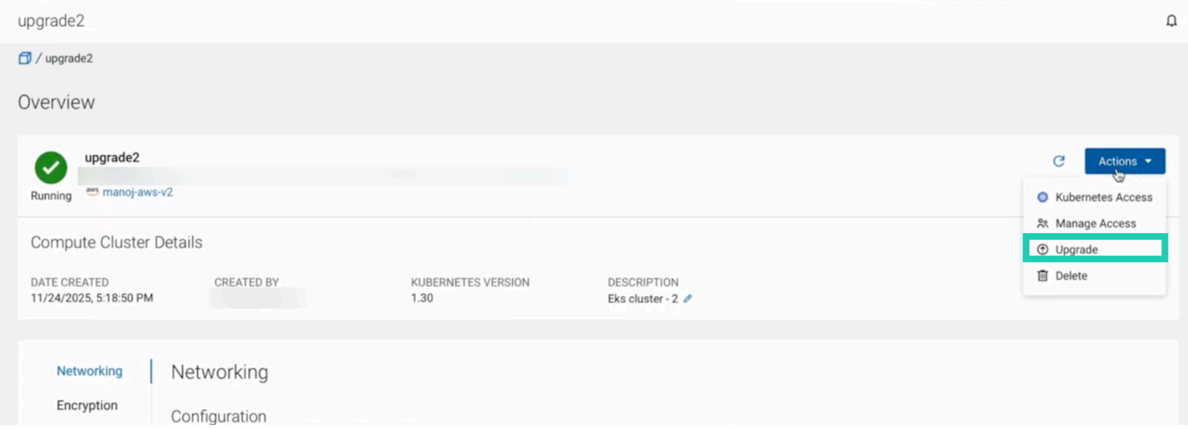

upgrade2 [redacted]
Running manoj-aws-v2

Compute Cluster Details

DATE CREATED	CREATED BY	KUBERNETES VERSION	DESCRIPTION
11/24/2025, 5:18:50 PM	[redacted]	1.30	Eks cluster - 2

   b. Click on the ⊕ next to the new version.

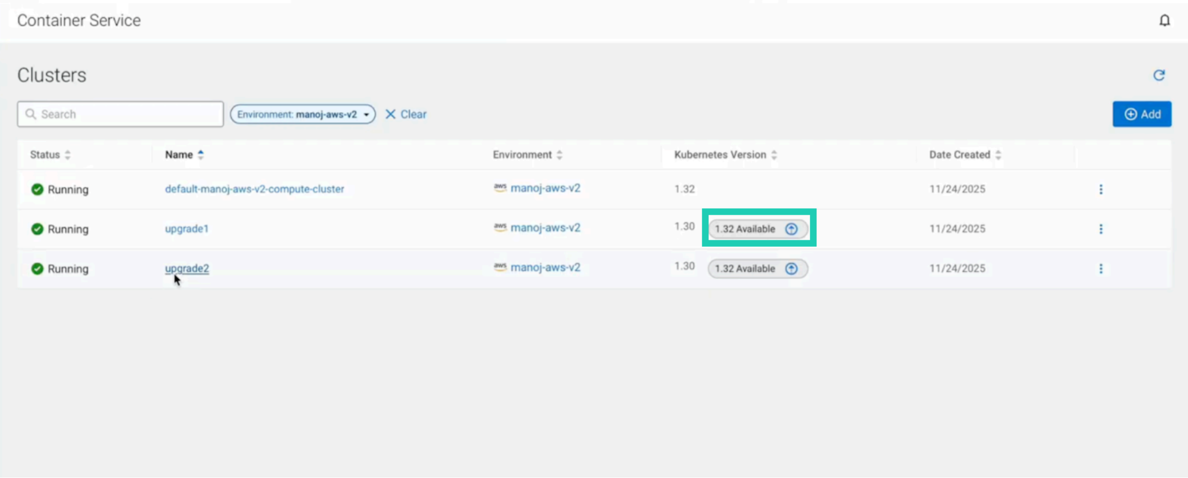

Clusters

Status	Name	Environment	Kubernetes Version	Date Created
Running	default-manoj-aws-v2-compute-cluster	manoj-aws-v2	1.32	11/24/2025
Running	upgrade1	manoj-aws-v2	1.30 1.32 Available	11/24/2025
Running	upgrade2	manoj-aws-v2	1.30 1.32 Available	11/24/2025

4. Click Ok to confirm the upgrade.

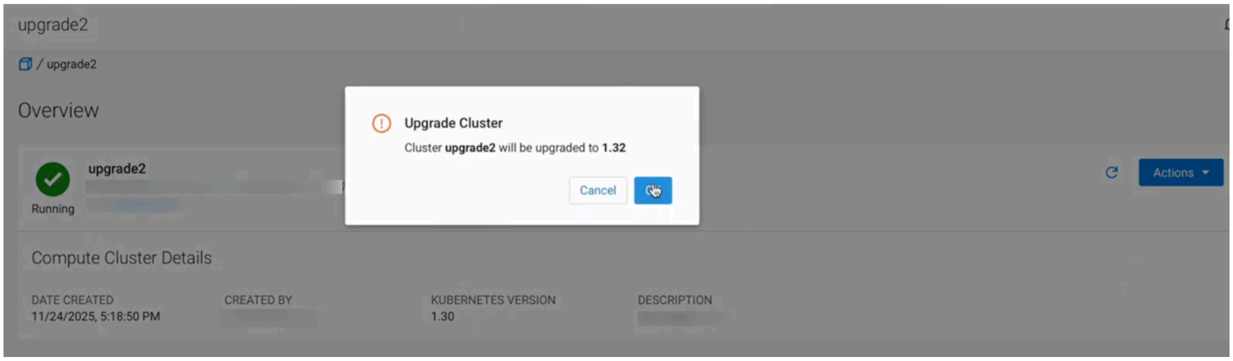

Upgrade Cluster

Cluster **upgrade2** will be upgraded to **1.32**

Compute Cluster Details

DATE CREATED: 11/24/2025, 5:18:50 PM
CREATED BY: [redacted]
KUBERNETES VERSION: 1.30
DESCRIPTION: [redacted]

The upgrade process will be triggered. You can track the phases of the upgrade process under the **Cluster Upgrade Operation**.

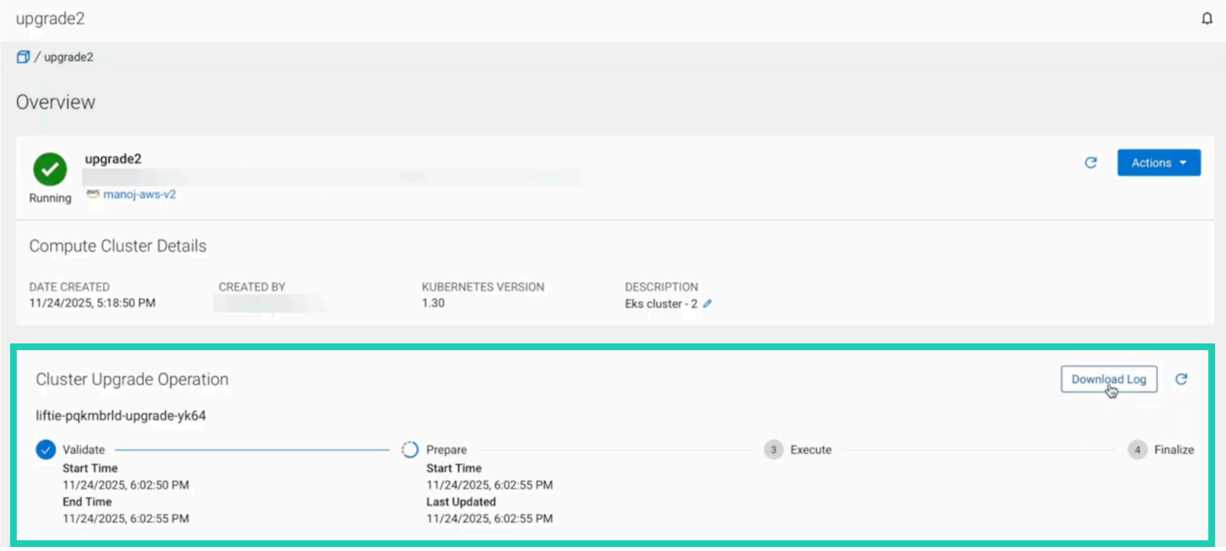

Compute Cluster Details

DATE CREATED: 11/24/2025, 5:18:50 PM
CREATED BY: [redacted]
KUBERNETES VERSION: 1.30
DESCRIPTION: Eks cluster - 2

Cluster Upgrade Operation

liftie-pqkmbrld-upgrade-yk64

✓ Validate — Start Time 11/24/2025, 6:02:50 PM; End Time 11/24/2025, 6:02:55 PM
○ Prepare — Start Time 11/24/2025, 6:02:55 PM; Last Updated 11/24/2025, 6:02:55 PM
3 Execute
4 Finalize

Note: In case there is any error during the upgrade process, you can view the logs by using the Download Log button. After fixing the error that caused the upgrade process failure, you can retry the upgrade process again using ↺ next to the Download Log button.

A check mark is shown at each phase, Validate, Prepare, Execute and Finalize, when the given phase is executed successfully as shown in the following example:

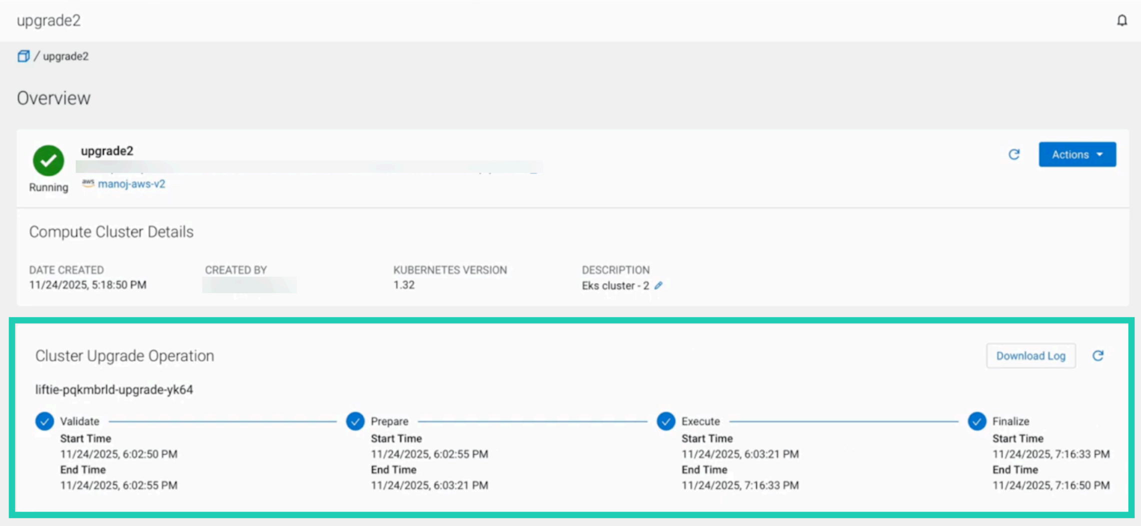

Compute Cluster Details

DATE CREATED: 11/24/2025, 5:18:50 PM
CREATED BY: [redacted]
KUBERNETES VERSION: 1.32
DESCRIPTION: Eks cluster - 2

Cluster Upgrade Operation

liftie-pqkmbrld-upgrade-yk64

✓ Validate — Start Time 11/24/2025, 6:02:50 PM; End Time 11/24/2025, 6:02:55 PM
✓ Prepare — Start Time 11/24/2025, 6:02:55 PM; End Time 11/24/2025, 6:03:21 PM
✓ Execute — Start Time 11/24/2025, 6:03:21 PM; End Time 11/24/2025, 7:16:33 PM
✓ Finalize — Start Time 11/24/2025, 7:16:33 PM; End Time 11/24/2025, 7:16:50 PM

Using CDP CLI

Run the following command to upgrade the Compute Cluster:

```
cdp compute upgrade-cluster --cluster-crn [***CLUSTER CRN ***]
```

The above command returns an operationId that can be used to further track the upgrade progress:

```
cdp compute get-operation-details --operation-id [***OPERATION ID ***]
cdp compute get-operation-status --operation-id [***OPERATION ID ***]
```

In case the upgrade process has failed, you can use the retry-operation command with the operationId to trigger the upgrade again:

```
cdp compute retry-operation --operation-id [***OPERATION ID ***]
```

Using Compute Clusters in environments on Azure

Cloudera Container Service enables you to deploy a containerized platform on Kubernetes for Data Services and shared services.

Cloudera Container Service offers simplified management, enhanced efficiency, and centralized control that leads to faster deployments, reduced configuration errors and improved system reliability. As multiple Data Services can optionally share the same Compute Cluster within the Container Service, it also lowers the cost of ownership.

When registering a Compute Cluster enabled environment, a default Compute Cluster is created. This default Compute Cluster is tied to the environment, and is used for running critical applications. The default Compute Cluster is labeled as **Default Cluster** in your environment to distinguish it from the other Compute Clusters.

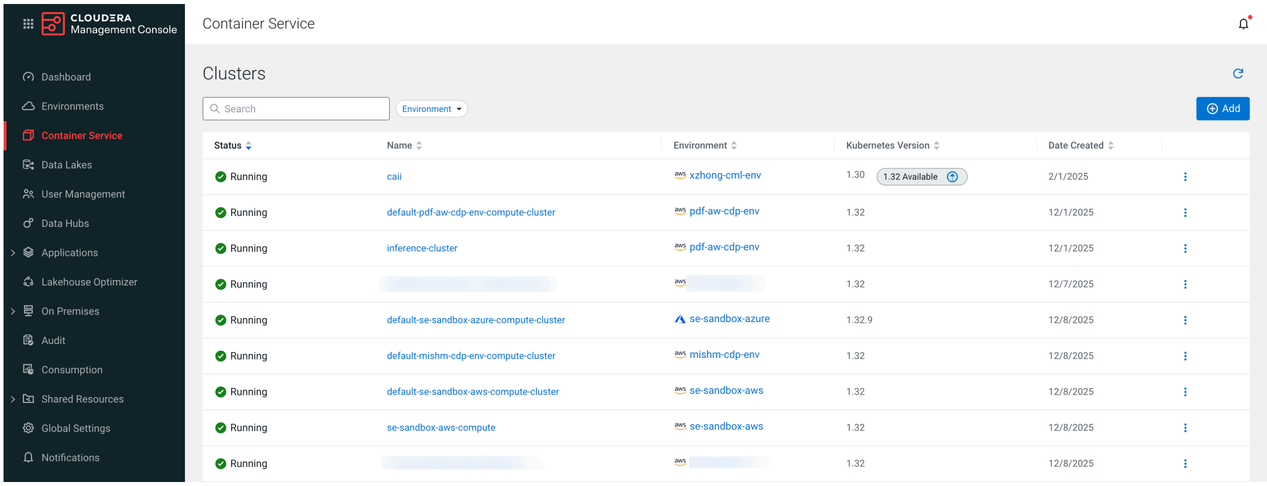

Status	Name	Environment	Kubernetes Version	Date Created
Running	caii	xzhong-cml-env	1.30 1.32 Available	2/1/2025
Running	default-pdf-aw-cdp-env-compute-cluster	pdf-aw-cdp-env	1.32	12/1/2025
Running	inference-cluster	pdf-aw-cdp-env	1.32	12/1/2025
Running	[redacted]	[redacted]	1.32	12/7/2025
Running	default-se-sandbox-azure-compute-cluster	se-sandbox-azure	1.32.9	12/8/2025
Running	default-mishm-cdp-env-compute-cluster	mishm-cdp-env	1.32	12/8/2025
Running	default-se-sandbox-aws-compute-cluster	se-sandbox-aws	1.32	12/8/2025
Running	se-sandbox-aws-compute	se-sandbox-aws	1.32	12/8/2025
Running	[redacted]	[redacted]	1.32	12/8/2025

You can create additional Compute Clusters that inherit the configuration of the default Compute Cluster. You can manage the lifecycle of the additional Compute Clusters as they are independent of the environment.

Enabling default Compute Cluster for new environments

When creating your environment, you can enable the default Compute Cluster using the Cloudera Management Console or CDP CLI to be able to run your data and shared services on the containerized platform.

Required role: EnvironmentAdmin

Before you begin

- Ensure that your Azure account has all the resources required by Cloudera.

  For more information, see Azure subscription requirements.

Using Cloudera Management Console

When creating your environment in **Cloudera Management Console**, ensure that you use the **Enable Compute Cluster** setting to create the Compute Cluster enabled environment.

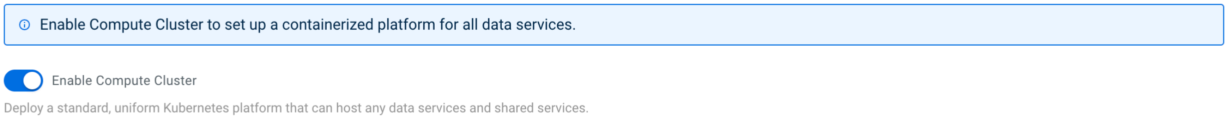

ⓘ Enable Compute Cluster to set up a containerized platform for all data services.

● Enable Compute Cluster

Deploy a standard, uniform Kubernetes platform that can host any data services and shared services.

After completing the step for **Data Access and Data Lake Scaling**, configure the networking settings for Kubernetes with enabling **User Defined Routing**, and selecting **Private Kubernetes Cluster** or providing **Authorized IP Ranges** on the **Region, Networking and Security** page. When selecting **Private Kubernetes Cluster**, you also need to select an existing private DNS zone or select creating a new private DNS zone by Cloudera on your Azure account for the database. **Worker Node Subnets** are automatically pre-filled with the same set of subnets provided in **Network** section, but you have the option to not use all of the available subnets.

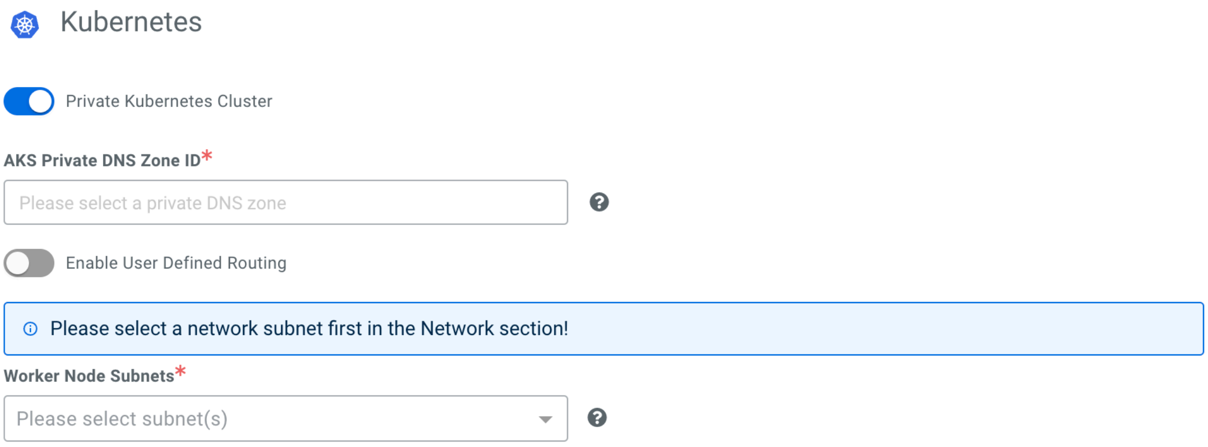

Kubernetes

● Private Kubernetes Cluster

AKS Private DNS Zone ID*

Please select a private DNS zone

○ Enable User Defined Routing

ⓘ Please select a network subnet first in the Network section!

Worker Node Subnets*

Please select subnet(s)

For more information about creating your environment, see the Registering environment (UI) documentation.

Using CDP CLI

Run the following command to create the Compute Cluster enabled environment:

For Without private cluster

```
cdp environments create-azure-environment \
--environment-name [***ENVIRONMENT NAME***] \
--credential-name [***CREDENTIAL NAME***] \
--region [***REGION***] \
--public-key [***PUBLIC SSH KEY***] \
--security-access [***SECURITY ACCESS CONFIGURATION***] \
--use-public-ip | --no-use-public-ip \
--log-storage [***STORAGE CONFIGURATION***] \
--enable-compute-cluster \
--compute-cluster-configuration \
privateCluster=FALSE, \
kubeApiAuthorizedIpRanges=[***CIDR1***],[***CIDR2***]
```

```
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

For With private cluster

```
cdp environments create-azure-environment \
--environment-name [***ENVIRONMENT NAME***] \
--credential-name [***CREDENTIAL NAME***] \
--region [***REGION***] \
--public-key [***PUBLIC SSH KEY***] \
--security-access [***SECURITY ACCESS CONFIGURATION***] \
--use-public-ip | --no-use-public-ip \
--log-storage [***STORAGE CONFIGURATION***] \
--enable-compute-cluster \
--compute-cluster-configuration \
privateCluster=TRUE
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

After the command runs, you can verify if the environment was successfully created with the default Compute Cluster with using the following commands:

- Describing the environment:

```
cdp environments describe-environment --env-name-or-crn [***ENVIRONMENT
 NAME OR CRN***]

...
        "azureComputeClusterConfiguration": {
            "privateCluster": false,
            "kubeApiAuthorizedIpRanges": [
                "0.0.0.0/0"
            ]
        },
        "enableComputeCluster": "true"
...
```

- Listing Compute Clusters:

```
cdp compute list-clusters --env-name-or-crn [***ENVIRONMENT NAME OR
 CRN***]
```

Note: FreeIPA must be created and running before the default Compute Cluster is created.

You can use the following command to retry the environment creation with the default Compute Cluster:

For Without private cluster

```
cdp environments initialize-azure-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=FALSE, \
kubeApiAuthorizedIpRanges=[***CIDR1***],[***CIDR2***]
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

For With private cluster

```
cdp environments initialize-azure-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=TRUE
```

```
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

Enabling default Compute Cluster for existing environments

In case you already have an environment, but would like to have your services to run on the containerized platform enabled by Compute Clusters, you can add the default Compute Cluster to your existing environment.

Required role: EnvironmentAdmin

Before you begin

- Ensure that the environment has no default Compute Cluster provisioned.
- Ensure that the environment is started and available.

Using Cloudera Management Console

1. Navigate to your environment.
2. Click Compute Clusters.

   The following message will indicate that Compute Cluster are not enabled for your environment:

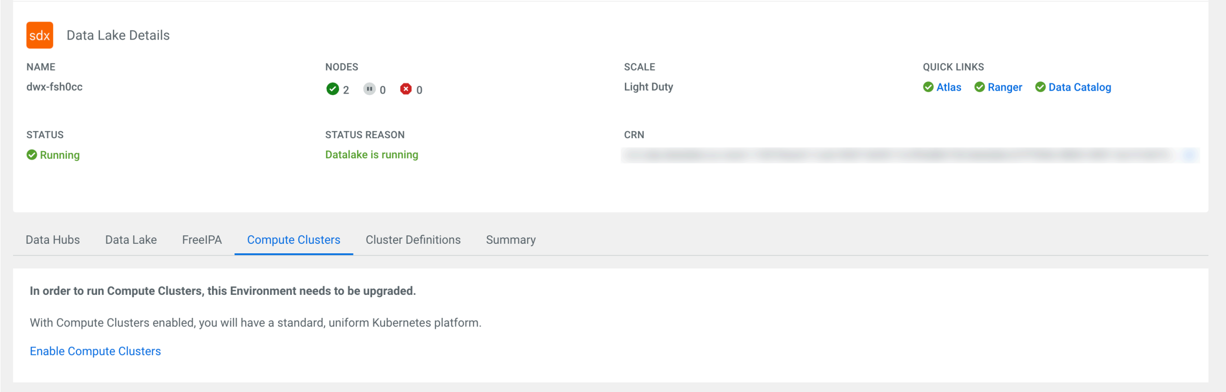

NAME dwx-fsh0cc

NODES 2 0 0

SCALE Light Duty

STATUS Running

STATUS REASON Datalake is running

CRN [redacted]

In order to run Compute Clusters, this Environment needs to be upgraded.

With Compute Clusters enabled, you will have a standard, uniform Kubernetes platform.

3. Click Enable Compute Clusters.
4. Provide the necessary networking information for the **Kubernetes** cluster.

   a. 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.
   b. If you need to create a Private Cluster, enable Private Kubernetes Cluster to create a private cluster that blocks all access to the API Server endpoint.
   c. If you do not need to create a Private Cluster, 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.

5. Click Submit.

You will be redirected to the **Compute Clusters** tab, where you can track the creation process of the default Compute Cluster.

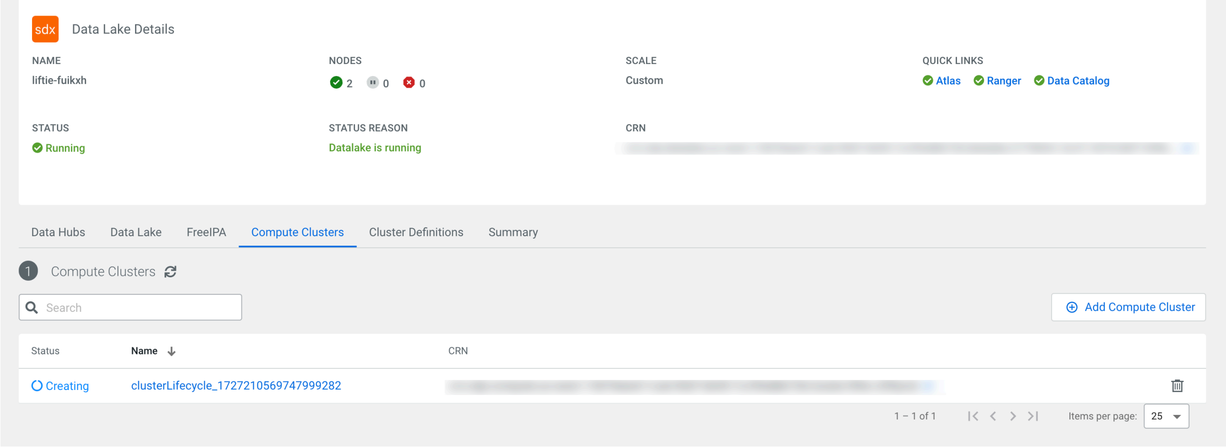

sdx Data Lake Details

NAME	NODES	SCALE	QUICK LINKS
liftie-fuikxh	✓ 2 ● 0 ✗ 0	Custom	✓ Atlas ✓ Ranger ✓ Data Catalog

STATUS	STATUS REASON	CRN
✓ Running	Datalake is running	[redacted]

Status	Name	CRN
Creating	clusterLifecycle_1727210569747999282	[redacted]

Using CDP CLI

Run the following command to add the default Compute Cluster to the environment:

For Without private cluster

```
cdp environments initialize-azure-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=FALSE, \
kubeApiAuthorizedIpRanges=[***CIDR1***],[***CIDR2***]
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

For With private cluster

```
cdp environments initialize-azure-compute-cluster
--environment-name [***ENVIRONMENT NAME***] \
--compute-cluster-configuration \
privateCluster=TRUE
workerNodeSubnets=[***SUBNET1***],[***SUBNET2***]
```

The environment will have COMPUTE_CLUSTER_CREATION_IN_PROGRESS status. You can use the following command to check the status of the environment creation, the statusReason field will contain the information about the process:

```
cdp environments describe-environment --env-name-or-crn [***ENVIRONMENT
 NAME***]
```

For more detailed status information about the cluster creation, you can use the following command:

```
cdp compute list-clusters --env-name-or-crn [***ENVIRONMENT NAME***]
```

Adding more Compute clusters

You can add as many additional Compute Clusters as required beside the default Compute Cluster using Cloudera Management Console or CDP CLI.

Required role: EnvironmentAdmin

Using Cloudera Management Console

You can create additional Compute Clusters beside the default Compute Cluster using Cloudera Management Console.

1. Navigate to your environment.
2. Select Compute Clusters tab.
3. Click Add Compute Cluster.

   The **Add Compute Cluster** wizard appears.

   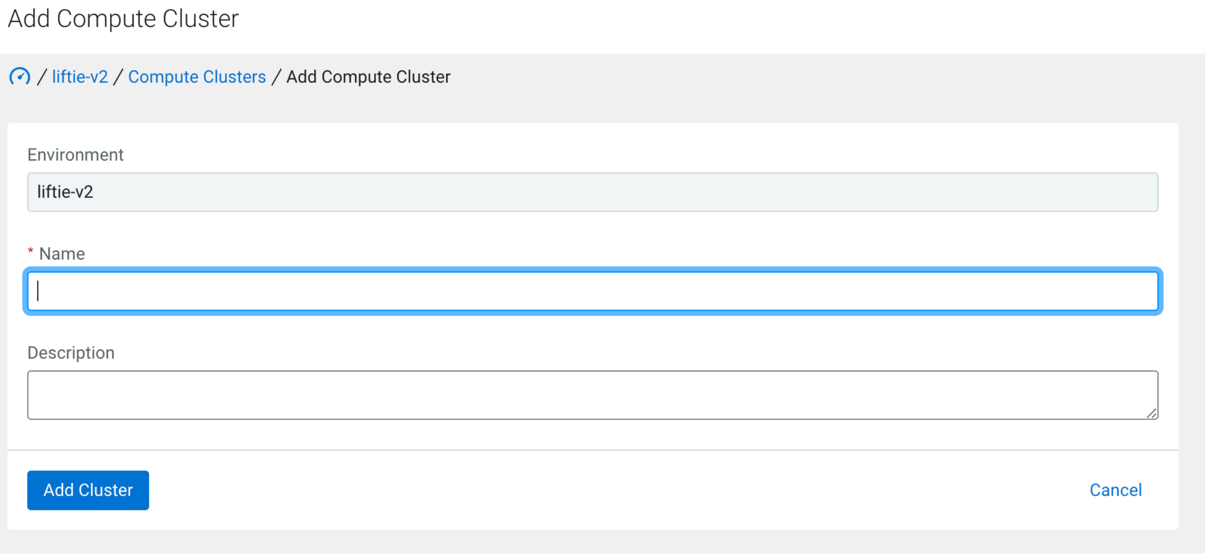

   Environment

   liftie-v2

4. Provide a Name to the cluster, and optionally a Description.
5. Click Add Cluster.

   You will be redirected to the **Compute Clusters** tab, where you can track the creation process of the additional Compute Cluster.

Using CDP CLI

You can use the following CDP CLI command to create additional Compute Clusters after the default Compute Cluster is created:

```
cdp compute create-cluster
--environment --env-name-or-crn [***ENVIRONMENT NAME OR CRN***] \
--name [***CLUSTER NAME***]
```

After the command runs, you can verify if the additional Compute Cluster creation was successful using the following command:

```
cdp compute list-clusters --env-name-or-crn [***ENVIRONMENT NAME OR CRN***]
```

The additional Compute Cluster status should be in RUNNING state.

Managing Compute Clusters

After creation, you can view the Compute Cluster details, manage the access of the clusters, download the Kubeconfig file, and delete the created additional Compute Clusters.

Required resource role: EnvironmentAdmin or Owner

Required account role: IamViewer

Using Management Console

1. Click on the Name of the Compute Cluster.

   You will be redirected to the **Cluster Details** page.

   On the **Cluster Details** page, you can view the **Status**, **Creation Date**, **Created By**, **Description** and **CRN** of the Compute Cluster.
2. Click Actions to open the drop-down menu.

   > **Note:** Ensure that you have one of the following roles to access the **Actions** menu
   >
   > - EnvironmentOwner or EnvironmentAdmin
   > - ClusterCreator or ClusterOwner

   a. Click Manage Access to update the list of users who have access to manage the Compute Cluster and use it for installing Data Services.

      1. Click Update roles to update the **Resource Role** of a user or group.

         You can assign the Owner resource role to a user or group to provide permission to manage the Compute Cluster and install services on it.

   b. Click Kubeconfig to download the Kubeconfig file.

      1. Click Download Kubeconfig.

   c. Click Delete Compute Cluster, if you no longer need the additional Compute Cluster.

      > **Warning:** Ensure that the data services running on the Compute Cluster are deleted before deleting the Compute Cluster itself.

      1. Confirm the deletion of the additional Compute Cluster by clicking Remove.
3. Click Networking tab to view the subnet information of the Compute Cluster.

   a. Click ✎ to update the **Kubernetes API Server Authorized IP Ranges**.

      > **Note:** Ensure that you have one of the following roles to access the **Actions** menu
      >
      > - EnvironmentOwner or EnvironmentAdmin
      > - ClusterCreator or ClusterOwner

      1. Add or remove the CIDRs, and click Save.
4. Click **Encryption** tab to view the encryption key of the Compute Cluster.
5. Click **Node Groups** tab to have an illustrated overview of the resource utilization of the different services running on the Compute Cluster.
6. Click **Compute Cluster Version** to view the Kubernetes and Insfrastructure Service versions of the cluster.
7. Click **Logs** to check the different events of the Compute Cluster.

Suspending and resuming Compute Clusters

You can suspend and resume your default and additional Compute Clusters in an environment to manage your cloud costs using CDP CLI.

Before you begin

- Ensure that the environment is started and available.
- Ensure that there are no Data Service workload instances running in the Compute Cluster.

Required resource role: EnvironmentAdmin or Owner

Procedure

1. Run the following command to suspend the Compute Cluster:

   ```
   cdp compute suspend-cluster --cluster-crn [***CLUSTER CRN***]
   ```

   After running the command, the cluster status changes to SUSPENDING, and will progress to SUSPENDED once the operation completes. The status of the cluster can be polled using describe-cluster command:

   ```
   cdp compute describe-cluster --cluster-crn [***CLUSTER CRN***]
   ```

2. Run the following command to resume the Compute Cluster with SUSPENDED status:

   ```
   cdp compute resume-cluster --cluster-crn [***CLUSTER CRN***]
   ```

   After running the command, the cluster status changes to RESUMING, and will progress to RUNNING status once the operation completes.

   Note: Only the following operations are supported on a cluster with SUSPENDED status:

   - DescribeCluster
   - ListClusters
   - DeleteCluster

Deleting Compute Clusters

You can delete additional Compute Clusters in an environment to manage your cloud costs using CDP CLI.

Before you begin:

- Ensure that the environment is started and available.
- Ensure that there are no Data Service workload instances running in the Compute Cluster.
- You can only delete additional Compute Clusters.
- You can only delete the default Cluster by deleting the environment.

Run the following command to delete the Compute Cluster:

```
cdp compute delete-cluster --cluster-crn [***CLUSTER CRN***]
```

Upgrading Compute Clusters

You can upgrade default and additional Compute Clusters in an environment when there is a new Kubernetes version available from Cloudera Management Console and from CDP CLI.

Required role: EnvironmentAdmin or EnvironmentOwner

Before you begin

- Ensure that the Cloudera environment is running and available..

Using Cloudera Management Console

1. Navigate to Container Service in Cloudera Management Console.

2. Find the compute cluster that you want to upgrade.

   You can scroll through the list of compute clusters or use the search box to filter down the list of clusters.

   In case there is an upgrade available for the compute cluster, the new version is indicated next to the Kubernetes Version.

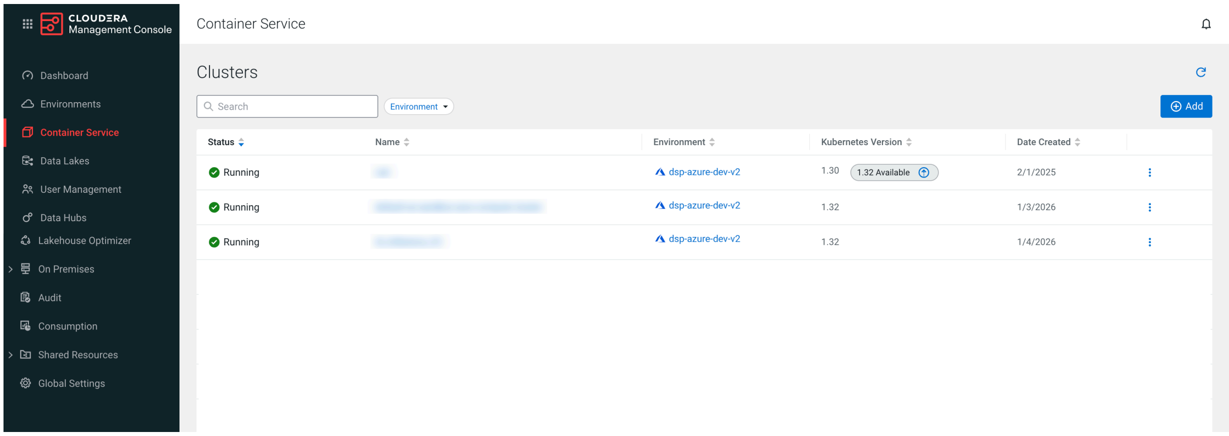

Status	Name	Environment	Kubernetes Version	Date Created
Running	[redacted]	dsp-azure-dev-v2	1.30 1.32 Available	2/1/2025
Running	[redacted]	dsp-azure-dev-v2	1.32	1/3/2026
Running	[redacted]	dsp-azure-dev-v2	1.32	1/4/2026

3. Start the upgrade process by using one of the following steps:

   a. Click on the cluster that you want to upgrade. You will be redirected to the **Compute Cluster Details** page, and select Actions Upgrade .

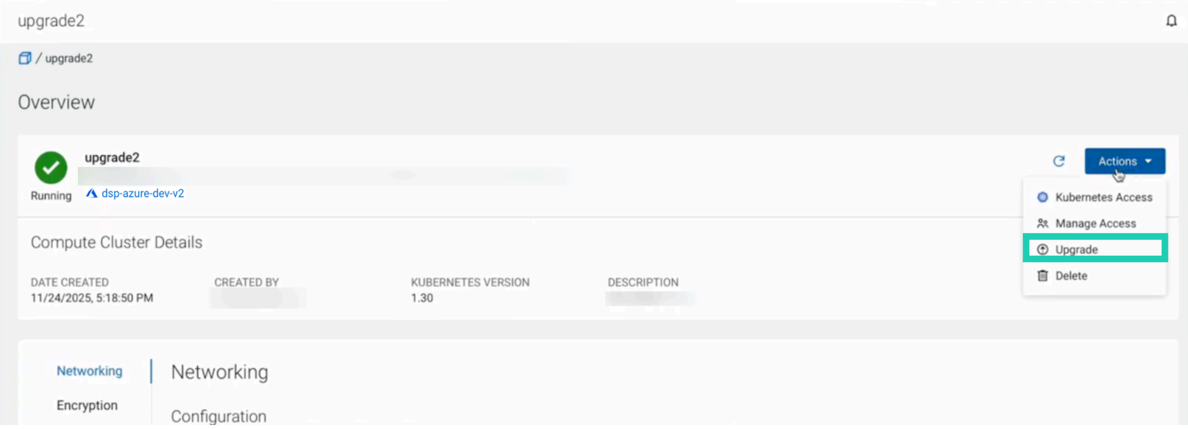

upgrade2 [redacted]

Running dsp-azure-dev-v2

Compute Cluster Details

DATE CREATED	CREATED BY	KUBERNETES VERSION	DESCRIPTION
11/24/2025, 5:18:50 PM	[redacted]	1.30	[redacted]

   b. Click on the ⬆ next to the new version.

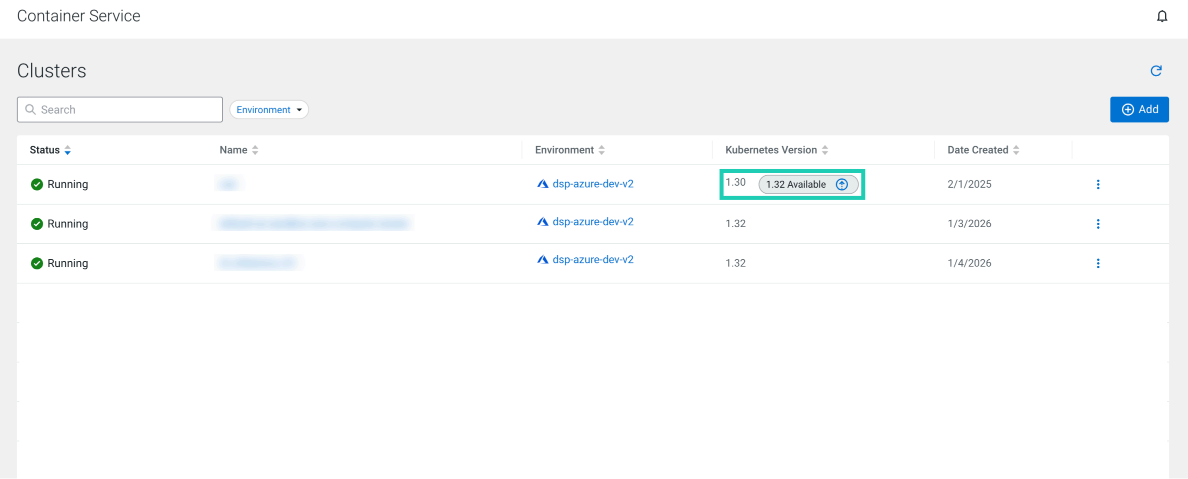

Status	Name	Environment	Kubernetes Version	Date Created
Running	[redacted]	dsp-azure-dev-v2	1.30 1.32 Available	2/1/2025
Running	[redacted]	dsp-azure-dev-v2	1.32	1/3/2026
Running	[redacted]	dsp-azure-dev-v2	1.32	1/4/2026

4. Click Ok to confirm the upgrade.

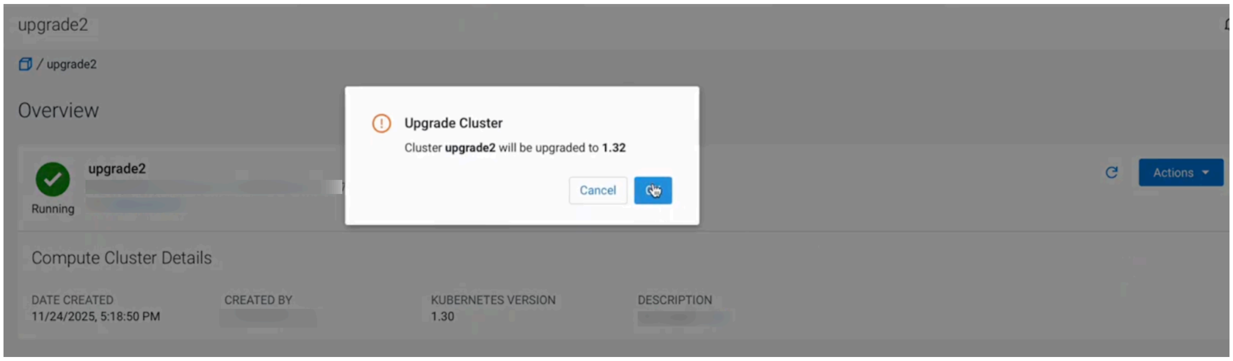

upgrade2

/ upgrade2

Overview

Upgrade Cluster

Cluster **upgrade2** will be upgraded to **1.32**

upgrade2 [redacted]

Running

Compute Cluster Details

DATE CREATED
11/24/2025, 5:18:50 PM

CREATED BY
[redacted]

KUBERNETES VERSION
1.30

DESCRIPTION
[redacted]

The upgrade process will be triggered. You can track the phases of the upgrade process under the **Cluster Upgrade Operation**.

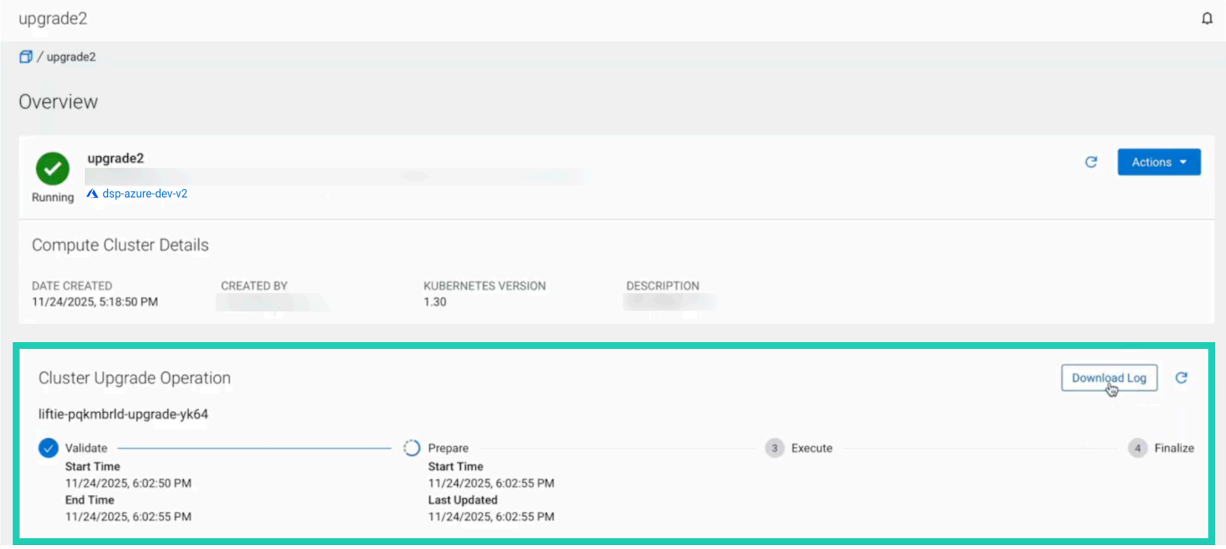

upgrade2

/ upgrade2

Overview

upgrade2 [redacted]

Running dsp-azure-dev-v2

Compute Cluster Details

DATE CREATED
11/24/2025, 5:18:50 PM

CREATED BY
[redacted]

KUBERNETES VERSION
1.30

DESCRIPTION
[redacted]

Cluster Upgrade Operation

liftie-pqkmbrld-upgrade-yk64

Validate
Start Time
11/24/2025, 6:02:50 PM
End Time
11/24/2025, 6:02:55 PM

Prepare
Start Time
11/24/2025, 6:02:55 PM
Last Updated
11/24/2025, 6:02:55 PM

3 Execute

4 Finalize

Note: In case there is any error during the upgrade process, you can view the logs by using the Download Log button. After fixing the error that caused the upgrade process failure, you can retry the upgrade process again using ↺ next to the Download Log button.

A check mark is shown at each phase, Validate, Prepare, Execute and Finalize, when the given phase is executed successfully as shown in the following example:

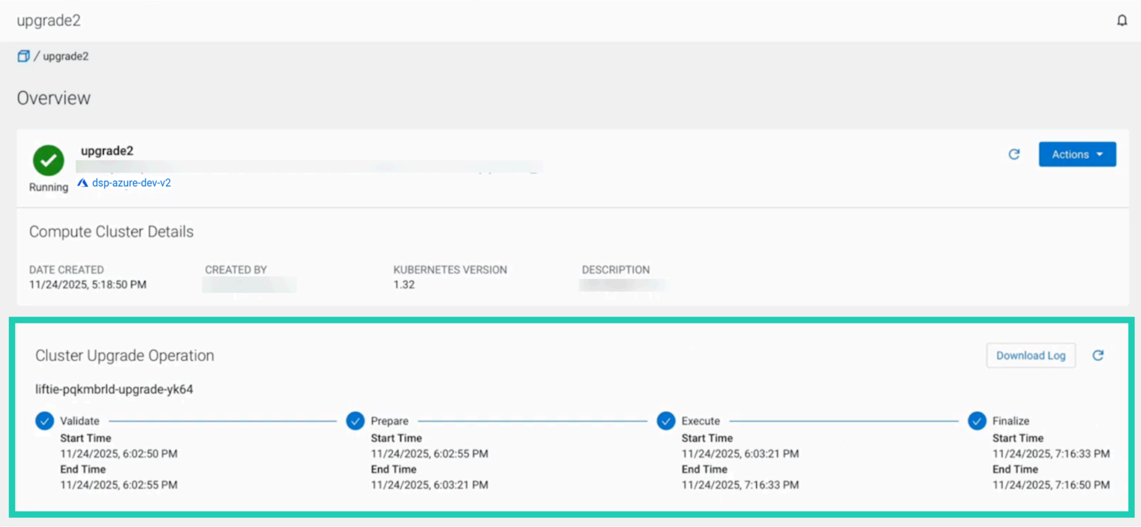

upgrade2

/ upgrade2

Overview

upgrade2 [redacted]

Running dsp-azure-dev-v2

Compute Cluster Details

DATE CREATED
11/24/2025, 5:18:50 PM

CREATED BY
[redacted]

KUBERNETES VERSION
1.32

DESCRIPTION
[redacted]

Cluster Upgrade Operation

liftie-pqkmbrld-upgrade-yk64

Validate
Start Time
11/24/2025, 6:02:50 PM
End Time
11/24/2025, 6:02:55 PM

Prepare
Start Time
11/24/2025, 6:02:55 PM
End Time
11/24/2025, 6:03:21 PM

Execute
Start Time
11/24/2025, 6:03:21 PM
End Time
11/24/2025, 7:16:33 PM

Finalize
Start Time
11/24/2025, 7:16:33 PM
End Time
11/24/2025, 7:16:50 PM

Using CDP CLI

Run the following command to upgrade the Compute Cluster:

```
cdp compute upgrade-cluster --cluster-crn [***CLUSTER CRN ***]
```

The above command returns an operationId that can be used to further track the upgrade progress:

```
cdp compute get-operation-details --operation-id [***OPERATION ID ***]
cdp compute get-operation-status --operation-id [***OPERATION ID ***]
```

In case the upgrade process has failed, you can use the retry-operation command with the operationId to trigger the upgrade again:

```
cdp compute retry-operation --operation-id [***OPERATION ID ***]
```