Cloudera Data Services on premises 1.5.5

# Cloudera Data Services on premises Installation on the OpenShift Container Platform

**Date published: 2023-12-16**
**Date modified: 2025-11-08**

CLOUDERA

**https://docs.cloudera.com/**

# Legal Notice

# Contents

# Requirements for installing Cloudera Data Services on premises on OpenShift

Learn about the requirements of Cloudera Data Services on premises with the OpenShift Container Platform.

## Software Support Matrix for OpenShift

This support matrix lists the supported software for the Cloudera Base on premises cluster and the Cloudera Data Services on premises containerized cluster when installing using the OpenShift Container Platform (OCP).

| | | |
|---|---|---|
| Base Cluster | Version | • Cloudera Manager 7.13.1.501<br>• 7.3.1.500, 7.1.9 SP1 CHF11, 7.1.7 SP3 CHF10 |
| | Base OS | • See Cloudera Base on premises OS requirements |
| | TLS | • AutoTLS (Custom CMCA)<br>• AutoTLS (Self-signed)<br>• Manual TLS |
| | Kerberos | • AD<br>• FreeIPA |
| | JDK | • See Java Requirements |
| | Custom service principals | • Not supported |
| | Data Lake Storage | • HDFS<br>• Ozone<br>• Iceberg v2 (with HDFS and Ozone) |
| | Base DB (HMS access from CDW Data Services*) | • Oracle 19c<br>• Oracle 19.9<br>• MySQL 8<br>• MySQL 5.7<br>• MariaDB 10.2<br>• MariaDB 10.3<br>• MariaDB 10.4<br>• MariaDB 10.5<br>• MariaDB 10.6<br>• Postgres 12<br>• Postgres 13<br>• Postgres 14<br>• Postgres 15<br>• Postgres 16<br>* Cloudera Data Warehouse uses a TLS enabled connection |

| | | |
|---|---|---|
| Containerized Cluster | Kubernetes | • OCP 4.18 (K8s 1.31) [Fresh install]<br>• OCP 4.18 [Upgrade from 1.5.3 or 1.5.4 to 1.5.5 SP1, no fresh install]<br>**Important:**<br>• OCP upgrades from earlier versions are incremental. See OCP upgrade steps for Cloudera Data Services on premises 1.5.5 for more information. |
| | Control Plane Metadata DB | • Embedded |
| | Vault | • External v1.9 (OCP only)<br>• Embedded |
| | Docker registry type | • Secure registry with self signed CA certs (pwd protected + self signed certs), trusted CA certs |
| | Storage | • OCS (rebranded ODF) (SSD support only)<br>• Pure Portworx |
| | NFS | • Embedded<br>• External |
| | IdP | • FreeIPA<br>• ActiveDirectory (LDAP)<br>• OpenLDAPs |
| | Network Access | • Airgap<br>• HTTP proxy (CML) |
| | TLS | • Manual - CA signed |

## Red Hat OpenShift Container Platform hardware requirements

Cloudera on premises requires hardware for a dedicated OpenShift Container Platform (OCP) cluster. An OpenShift cluster consists of several master nodes for managing OpenShift and many worker nodes for running your application on Cloudera.

The sizing of the OpenShift cluster depends on:

- The OpenShift cluster setup on the master nodes
- Application workloads deployed on the worker nodes

The Cloudera Data Services on premises is installed on the OpenShift worker nodes with SSD disks.

**Important:** Cloudera Data Services on premises is only certified and supported for standalone OCP cluster deployments. You must not have any other third-party applications that use the same OCP cluster, as Cloudera will only support a single-tenant use on OCP.

Using the Cloudera certified cluster deployments has the following benefits:

- **Full supportability and accountability** - The single-tenant cluster deployment establishes a vertically integrated stack that provides full supportability and a single point of accountability for the entire software and Kubernetes environment. A single point of accountability accelerates issue resolution, eliminates dependency issues, and ensures Cloudera certification validation.
- **Predictable high performance** - A dedicated environment guarantees full resource isolation, ensuring that all CPU, memory, I/O, and network capacity are reserved solely for Cloudera workloads. This isolation is critical for achieving predictable high performance required for data processing Service Level Agreements (SLAs).
- **Strong security boundary** - Deploying Data Services on a separate cluster creates a strong security boundary. This strong isolation reduces the attack surface and simplifies the achievement of certifications and audits necessary for data governance and compliance mandates.
- **Optimization** - The dedicated cluster allows for the required deployment and full utilization of Apache YuniKorn. YuniKorn is necessary to efficiently schedule workloads with advanced capabilities such as queue-based quotas and priority scheduling.

**Note:** A dedicated cluster deployment is mandatory for mission-critical data workloads to ensure maximum reliability, security, and performance isolation.

The following table lists the hardware requirements for each node type. You require at least 3 minimum OpenShift Master Nodes + 1 Cluster System Admin Host (CSAH) Node + 1 Bootstrap Node. You need worker nodes based on your application workload requirements.

| Role | CPU cores | Memory | Storage (SSD support only) |
|---|---|---|---|
| Master | 4 | 16 GB | 120 GB |
| CSAH | 4 | 64 GB | 200 GB |
| Bootstrap | 4 | 16 GB | 120 GB |
| Worker | Depends on your workloads | Depends on your workloads | Depends on your workloads |

Additionally, if you plan to run Cloudera Data Warehouse or Cloudera AI data services workloads, you need to ensure that you meet the minimum requirements for each of those Data Services.

You can install Cloudera Data Services on premises in a low resource mode for Cloudera Data Warehouse workloads. For more information about OpenShift low resource mode requirements for Cloudera Data Warehouse, see *Get started with OpenShift low resource mode requirements* using the link in the related information section.

**Important:** Lowering the minimum hardware requirement reduces the up-front investment to deploy Cloudera Data Warehouse on OpenShift or ECS pods, but it does impact performance. Cloudera recommends that you use the Low Resource Mode option for proof of concept (POC) purposes only. This feature is not recommended for production deployment.

Complex queries and multiple queries on HS2 may fail due to limited memory configurations for HMS and HS2 in the low resource mode.

## Cloudera Data Warehouse hardware requirements

Review the requirements needed to get started with the Cloudera Data Warehouse service on Red Hat OpenShift.

You can also use the Cloudera Data Services on premises Spreadsheet to model the number and specification of hosts required for a deployment. See How to use the Cloudera Data Services on premises sizing spreadsheet.

- Cloudera Manager must be installed and running.
- Cloudera Data Services on premises must be installed and running. See Installing on OpenShift and Installing on ECS for more details.

- An environment must have been registered with Cloudera Management Console on the on premises. See Cloudera on premises Environments for more details.
- In addition to the general requirements, Cloudera Data Warehouse also has the following minimum memory, storage, and hardware requirements for each worker node using the standard resource mode:

Depending on the number of executors you want to run on each physical node, the per-node requirements change proportionally. For example, if you are running 3 executor pods per physical node, you require 384 GB of memory and approximately 1.8 TB of locally attached SSD/NVMe storage.

**Important:**

When you add memory and storage, it is very important that you add it in the increments stated:

- increments of 128 GB of memory
- increments of 600 GB of locally attached SSD/NVMe storage

If you add memory or storage that is not in the above increments, the memory and storage that exceeds these increments is not used for executor pods. Instead, the extra memory and storage can be used by other pods that require fewer resources.

For example, if you add 200 GB of memory, only 128 GB is used by the executor pods. If you add 2 TB of locally attached storage, only 1.8 TB is used by the executor pods.

## Security requirements

The Cloudera Data Warehouse service requires the "cluster-admin" role on the OpenShift and Cloudera Embedded Container Service cluster in order to install correctly. The "cluster-admin" role enables namespace creation and the use of the OpenShift Local Storage Operator for local storage.

## Low resource mode requirements

Review the memory, storage, and hardware requirements for getting started with the Cloudera Data Warehouse service in low resource mode on Red Hat OpenShift and Embedded Container Service (ECS). This mode reduces the minimum amount of hardware needed.

To get started with the Cloudera Data Warehouse service on Red Hat OpenShift or ECS low resource mode, make sure you have fulfilled the following requirements:

**Important:** Lowering the minimum hardware requirement reduces the up-front investment to deploy Cloudera Data Warehouse on OpenShift or ECS pods, but it does impact performance. Cloudera recommends that you use the Low Resource Mode option for proof of concept (POC) purposes only. This feature is not recommended for production deployment.

Complex queries and multiple queries on HS2 may fail due to limited memory configurations for HMS and HS2 in the low resource mode.

- Cloudera Manager must be installed and running.
- Cloudera Data Services on premises must be installed and running. See Installing on OpenShift and Installing on ECS for more details.
- An environment must have been registered with Cloudera Management Console on the on premises. See Cloudera on premises Environments for more details.
- In addition to the general requirements, Cloudera Data Warehouse also has the following minimum memory, storage, and hardware requirements for each worker node using the standard resource mode:

| Component | Low resource mode deployment |
|---|---|
| Nodes | 4 |
| CPU | 4 |
| Memory | 48 GB |
| Storage | 3 x 100 GB (SATA) or 2 x 200 GB (SATA) |
| Network Bandwidth | 1 GB/s guaranteed bandwidth to every Cloudera Base on premisesnode |

**Important:** When you add memory and storage for low resource mode, it is very important that you add it in the increments stated in the above table:

- increments of 48 GB of memory
- increments of at least 100 GB or 200 GB of SATA storage

If you add memory or storage that is not in the above increments, the memory and storage that exceeds these increments is not used for executor pods. Instead, the extra memory and storage can be used by other pods that require fewer resources.

### Virtual Warehouse low resource mode resource requirements

The following requirements are in addition to the low resource mode requirements listed in the previous section.

#### Table 1: Impala Virtual Warehouse low resource mode requirements

| Component | vCPU | Memory | Local Storage | Number of pods in XSMALL Virtual Warehouse |
|---|---|---|---|---|
| Coordinator (2) | 2 x 0.4 | 2 x 24 GB | 2 x 100 GB | 2 |
| Executor (2) | 2 x 3 | 2 x 24 GB | 2 x 100 GB | 2 |
| Statestore | 0.1 | 512 MB | -- | 1 |
| Catalogd | 0.4 | 16 GB | -- | 1 |
| Auto-scaler | 0.1 | 1 GB | -- | 1 |
| Hue (backend) | 1 | 8 GB | -- | 1 |
| Hue (frontend) | 0.5 | 8 GB | -- | 1 |
| Total for XSMALL Virtual Warehouse | 8 (7.9) | 121.5 GB | 400 GB - 3 volumes | -- |

Impala Admission Control Configuration

- Maximum concurrent queries per executor: 4
- Maximum query memory limit: 8 GB

#### Table 2: Hive Virtual Warehouse low resource mode requirements

| Component | vCPU | Memory | Local Storage | Number of pods in XSMALL Virtual Warehouse |
|---|---|---|---|---|
| Coordinator (2) | 2 x 1 | 2 x 4 GB | 2 x 100 GB | 2 |
| Executor (2) | 2 x 4 | 2 x 48 GB (16 GB heap; 32 GB off-heap) | 2 x 100 GB | 2 |
| HiveServer2 | 1 | 16 GB | -- | 1 |
| Hue (backend) | 1 | 8 GB | -- | 1 |
| Hue (frontend) | 0.5 | 8GB | -- | 1 |
| Standalone compute operator | 0.1 | 100 MB (.1 GB) | -- | -- |
| Standalone query executor (separate) | Same as executor | Same as executor | Same as executor | -- |
| Total for XSMALL Virtual Warehouse | 21 (20.6) | 237 GB (236.1) | 400 GB - 4 volumes | -- |

### Database Catalog low resource mode requirements

The HiveMetaStore (HMS) requires 2 CPUs and 8 GB of memory. Because HMS pods are in High Availability mode, they need a total of 4 CPUs and 16 GB of memory.

### Data Visualization low resource requirements

#### Table 3: Data Visualization low resource mode requirements

| vCPU | Memory | Local Storage | Number of pods in XSMALL Virtual Warehouse |
|---|---|---|---|
| 0.5 | 8 GB | -- | 1 |

## Requirements for Cloudera AI on Openshift Container Platform

To launch the Cloudera AI service, the OpenShift Container Platform (OCP) host must meet several requirements. Review the following Cloudera AI-specific software, NFS server, and storage requirements.

### Requirements

**Note:**

Only the usage of SSD disks is supported with Cloudera on premises on OpenShift Container Platform.

If needed, reach out to your Administrator to ensure the following requirements are met.

Compatibility requirements

- If you use OpenShift, check that the version of the installed OpenShift Container Platform is exactly as listed in Software Support Matrix for OpenShift.

Required accesses

- If Cloudera AI needs access to a database on the Cloudera Base on premises cluster, then the user must be authenticated using Kerberos and must have Ranger policies set up to allow read/write operations to the default (or other specified) database.
- Ensure that Kerberos is enabled for all services in the cluster. Custom Kerberos principals are not supported currently. For more information, see Authenticating Hue users with Kerberos.
- Cloudera AI assumes it has cluster-admin privileges on the cluster.
- If external NFS is used, the NFS directory and assumed permissions must be those of the cdsw user. For details see Using an External NFS Server.
- If you intend to access a workbench over https, see Deploying a Cloudera AI Workbench with support for TLS.

Requirements for functioning

- On OpenShift Container Platform, CephFS is used as the underlying storage provisioner for any new internal workbench on Cloudera on premises 1.5.x. A storage class named ocs-storagecluster-cephfs with csi driver set to openshift-storage.cephfs.csi.ceph.com must exist in the cluster for new internal workbenches to get provisioned.
- A block storage class must be marked as default in the cluster. This may be rook-ceph-block, Portworx, or another storage system. Confirm the storage class by listing the storage classes (run `oc get sc`) in the cluster, and check that one of them is marked default.

DNS-related requirements

- Forward and reverse DNS must be working.
- DNS lookups to sub-domains and the Cloudera AI Workbench itself shall work properly.
- In DNS, wildcard subdomains (such as *.cml.yourcompany.com) must be set to resolve to the master domain (such as cml.yourcompany.com). The TLS certificate (if TLS is used) must also include the wildcard subdomains. When a session or job is started, an engine is created for it, and the engine is assigned to a random, unique subdomain.

Configuration requirements

- The external load balancer server timeout needs to be set to 5 min. Without this, creating a project in a Cloudera AI Workbench with `git clone` or with the API may result in API timeout errors. For workarounds, see *Known Issue DSE-11837*.
- For non-TLS Cloudera AI Workbench, websockets need to be allowed for port 80 on the external load balancer.
- Only a TLS-enabled custom Docker Registry is supported. Ensure that you use a TLS certificate to secure the custom Docker Registry. The TLS certificate can be self-signed, or signed by a on premises or on cloud trusted Certificate Authority (CA).
- On OpenShift, due to a Red Hat issue with OpenShift Container Platform 4.17, the image registry cluster operator configuration must be set to Managed.
- Check if storage is set up in the cluster image registry operator. See *Known Issues DSE-12778* for further information.

For more information on requirements, see Install Cloudera Base on premises.

## Hardware requirements

Storage

The cluster must have persistent storage classes defined for both block and filesystem volume Modes of storage. Ensure that a block storage class is set up. The exact amount of storage classified as block or filesystem storage depends on the specific workload used:

### Table 4: Storage requirements for Cloudera AI on OCP

| | Local Storage (for example, ext4) | Block PV (for example, Ceph or Portworx) | NFS (for Cloudera AI user project files) |
|---|---|---|---|
| Cloudera Control Plane | N/A | 250 GB | N/A |
| Cloudera AI | N/A | The total (not per node) storage needed only for Cloudera AI in OCP without disaster recovery (drs) is 1800 Gi per workbench with the external NFS.<br>If the Cloudera AI Workbench is using the internal NFS, the minimum storage needed per workbench is 3800 Gi, considering the replication factor of 2. If you have a different replication configured, this will change accordingly.<br>Considering the DRS and a single backup of the workbench, the total storage needed is 1800 Gi * 2 = 3600 Gi for the workbench with external NFS. If the workbench uses internal NFS, the total storage needed is 7600Gi. | 1 TB per workbench (dependent on the size of the Cloudera AI user files) |

**Note:**

NFS storage must be routable from all pods running in the cluster.

**Note:**

For monitoring, the recommended volume size is 60 GB.

**Note:**

The storage calculation is based on the replication factor of two. If a different replication factor is used, the calculation will adjust accordingly.

External NFS considerations

Cloudera AI requires NFS 4.1 for storing project files and folders. NFS storage is to be used only for storing project files and folders, and not for any other Cloudera AI data, such as PostgreSQL database and LiveLog.

OpenShift requirements for NFS storage

An internal user-space NFS server can be deployed into the cluster which serves a block storage device (persistent volume) managed by the cluster's software defined storage (SDS) system, such as Ceph or Portworx. This is the recommended option for Cloudera AI on OpenShift. Alternatively, the NFS server can be external to the cluster, such as a NetApp filer that is accessible from the on premises cluster nodes. NFS storage is to be used only for storing project files and folders, and not for any other Cloudera AI data, such as PostgreSQL database and LiveLog.

Cloudera AI does not support shared volumes, such as Portworx shared volumes, for storing project files. A read-write-once (RWO) persistent volume must be allocated to the internal NFS server (for example, NFS server provisioner) as the persistence layer. The NFS server uses the volume to dynamically provision read-write-many (RWX) NFS volumes for the Cloudera AI clients.

## Cloudera Data Engineering hardware requirements

Review the requirements needed to get started with the Cloudera Data Engineering service on Red Hat OpenShift.

### Requirements

- Cloudera Data Engineering assumes it has cluster-admin privileges on the OpenShift cluster.
- Openshift cluster should be configured with route admission policy set to namespaceOwnership: InterNamespaceAllowed. This allows Openshift cluster to run applications in multiple namespaces with the same domain name. The

```
oc -n openshift-ingress-operator patch ingresscontroller/default --patch
 '{"spec":{"routeAdmission":
{"namespaceOwnership":"InterNamespaceAllowed"}}}' --type=merge
```

- Cloudera Data Engineering Service requirements: Overall for a Cloudera Data Engineering service, it requires 110 GB Block PV or NFS PV, 9 CPU cores, and 18 GB memory. The following are the Cloudera Data Engineering Service requirements:

#### Table 5: Cloudera Data Engineering Service requirements

| Component | vCPU | Memory | Block PV or NFS PV | Number of replicas |
|---|---|---|---|---|
| Embedded DB | 4 | 8 GB | 100 GB | 1 |
| Admission Controller | 250 m | 512 MB | -- | 1 |
| Config Manager | 500 m | 1 GB | -- | 2 |
| Dex Downloads | 250 m | 512 MB | -- | 1 |
| Knox | 250 m | 1 GB | -- | 1 |
| Management API | 1 | 2 GB | -- | 1 |
| NGINX Ingress Controller | 100 m | 90 MB | -- | 1 |
| Tgt Generator | 350 m | 630 MB | -- | 1 |
| FluentD Forwarder | 250 m | 512 MB | -- | 1 to 5 |
| Grafana | 250 m | 512 MB | 10 GB | 1 |
| Keytab Management | 250 m | 512 MB | -- | 1 |
| Data Connector | 250 m | 512 MB | -- | 1 |
| Total | 8600 m | 17.71 GB | 110 GB | |

- Cloudera Data Engineering Virtual Cluster requirements:
  - For Spark 3: Overall storage of 400 GB Block PV or Shared Storage PV, 5.35 CPU cores, and 15.6 GB per virtual cluster.
  - For Spark 2: If you are using Spark 2, you need additional 500 m CPU, 4.5 GB memory and 100 GB storage, that is, the overall storage of 500 GB Block PV or Shared Storage PV, 5.85 CPU cores, and 20.1 GB per virtual cluster.

  The following are the Cloudera Data Engineering Virtual Cluster requirements for Spark 3:

  **Table 6: Cloudera Data Engineering Virtual Cluster requirements for Spark 3**

| Component | vCPU | Memory | Block PV or NFS PV | Number of replicas |
|---|---|---|---|---|
| Airflow API | 350 m | 612 MB | 100 GB | 1 |
| Airflow Scheduler | 1 | 1 GB | 100 GB | 1 |
| Airflow Web | 250 m | 512 MB | -- | 1 |
| Runtime API | 250 m | 512 MB | 100 GB | 1 |
| Livy | 3 | 12 GB | 100 GB | 1 |
| SHS | 250 m | 1 GB | | 1 |
| Pipelines | 250 m | 512 MB | -- | 1 |
| Total | 5350 m | 16.1 GB | 400 GB | |

- **Important:** The above requirements does not include workloads. See the below workload information on the additional resources based on workload.
- Workloads: Depending upon the workload, you must configure resources.
  - The Spark Driver container uses resources based on the configured driver cores and driver memory and additional 40% memory overhead.
  - In addition to this, Spark Driver uses 110 m CPU and 232 MB for the sidecar container.
  - The Spark Executor container uses resources based on the configured executor cores and executor memory and additional 40 % memory overhead.
  - In addition to this, Spark Executor uses 10 m CPU and 32 MB for the sidecar container.
  - Minimal Airflow jobs need 100 m CPU and 200 MB memory per Airflow worker.

## Red Hat OpenShift Container Platform software requirements

You must understand the various OpenShift Container Platform (OCP) requirements before you install Cloudera Data Services on premises. Cloudera Data Services on premises requires at least one OpenShift cluster for the control plane and the environments. The Cloudera Data Warehouse, Cloudera AI, and Cloudera Data Engineering Data Services run on these environments.

**Note:** Do not configure policies that impose resource quotas on individual projects. Defining CPU, memory, limit, request, or object#count quotas at the project level is not supported.

Running oc describe AppliedClusterResourceQuota in a given project will return no results.

For more details, refer to the RedHat OpenShift Documentation on quotas and multi#project ClusterResou rceQuota configurations.

Review the Software Support Matrix for OpenShift on page 4.

Read the following topics to understand the various OpenShift integration requirements:

- Credentials
- Security context credentials
- Load balancing and ingress

- Certificate management and DNS
- Storage classes
- Docker registry access

## Credentials

You must have a kubeconfig file that has the cluster access information and authentication information for a single user, who has the “cluster-admin” pre-provisioned ClusterRole assigned.

Cloudera recommends that you use a kubeconfig file that does not expire, to avoid access issues to the installed software.

## Security context credentials

The Cloudera software must have privileged access at runtime. Cloudera recommends that you configure security context in your OpenShift cluster to ensure access to Cloudera Data Services on premises.

You must install additional scc definitions into OpenShift that Cloudera provides as part of the installation software. For more information about security context credentials in OpenShift, see Introduction to Security Contexts and SCCs.

Check with your Openshift Admin to add the registry entry of the external Docker repository server to the allowedR egistries in the Image Controller configuration (image.config.openshift.io/cluster), to avoid policy denial during deployment.

For example:

For more information please see this link from RedHat

```
[..]
spec:
registrySources:
  allowedRegistries:
  - ##your_external docker repository server name###
  - quay.io
  - registry.redhat.io
  - image-registry.openshift-image-registry.svc:5000
  - registry.example.com:5000
```

## Load balancing and ingress

OpenShift Route must be the default ingress controller setup on the cluster.

A non-terminating external load balancer must be configured to route ingress traffic on HTTP/HTTPS to the OpenShift cluster.

When a load balancer is used in front of the OCP external API, it must allow “Websocket traffic”, in addition to https.

## Certificate management and DNS

You must be aware of the reasons why an external DNS is required for Cloudera Data Services on premises installation along with the required setup in the cluster.

An external DNS must be available to route inbound traffic to the cluster through the load balancer. The external DNS should contain forward and reverse zones for both the OpenShift and the Cloudera Base on premises cluster nodes.

Ensure that the canonical load balancers required for OpenShift is routable from within the OpenShift cluster and from any other location that you want to access resources in the Cloudera Management Console; this is a standard requirement for on-premises load balancers communicating Kubernetes clusters.

There must also be a set of certificates set up for use by the OpenShift Route ingress controller as defined in the *OpenShift bare metal install guide* that the Cloudera services use.

## Storage classes

You need to have persistent storage classes defined in your OpenShift cluster. Storage classes can be defined by OpenShift cluster administrators.

The exact amount of storage classified as block or filesystem storage depends on the specific workloads (Cloudera AI or Cloudera Data Warehouse) and how they are used.

See the *Red Hat OpenShift documentation* for more information about OpenShift storage classes and persistent volumes.

To use Portworx as a storage platform, you must first create a Portworx storage class on your OCP cluster and then specify it in the Storage Class field while installing Cloudera Private Cloud Data Services on the OCP cluster. For information on how to create the storage class, see Step 4: StorageClass Setup in the Portworx documentation. See Installing in internet environment for information on how to install Cloudera Private Cloud Data Services on OCP.

After you specify the storage class while installing Cloudera Private Cloud Data Services, all other data services can use it.

## Volume snapshot support

Volume snapshot support for the storage class must be installed in your OpenShift cluster.

Run the following commands to determine whether or not volume support for the storage class is installed in your OpenShift cluster:

```
kubectl get sc
NAME                                        PROVISIONER
     RECLAIMPOLICY   VOLUMEBINDINGMODE      ALLOWVOLUMEEXPANSION   AGE
cdw-scratch                                 kubernetes.io/no-provisioner
    Delete          WaitForFirstConsumer   false                  82d
localblock                                  kubernetes.io/no-provisioner
     Delete          WaitForFirstConsumer   false                   82d
nfs                                         nfs-server-provisioner
    Delete          Immediate              true                   82d
ocs-storagecluster-ceph-rbd (default)   openshift-storage.rbd.csi.ceph.com
     Delete          Immediate              true                   82d
ocs-storagecluster-ceph-rgw             openshift-storage.ceph.rook.io/buck
et   Delete          Immediate              false                  82d
ocs-storagecluster-cephfs               openshift-storage.cephfs.csi.ceph
.com   Delete          Immediate              true                   82d
openshift-storage.noobaa.io             openshift-storage.noobaa.io/obc
    Delete          Immediate              false                  82d
```

```
kubectl get volumesnapshotclasses
NAME                                          DRIVER
       DELETIONPOLICY   AGE
ocs-storagecluster-cephfsplugin-snapclass   openshift-storage.cephfs.csi.
ceph.com   Delete           82d
ocs-storagecluster-rbdplugin-snapclass      openshift-storage.rbd.csi.ceph.c
om      Delete           82d
```

A storage class has a volume snapshot installed if there is an entry with the value in the DRIVER column returned by the second command that matches one of the values in the PROVISIONER column returned by the first command. In the example above, the following storage classes have volume snapshot support:

- ocs-storagecluster-ceph-rbd (default)
- ocs-storagecluster-cephfs

### Related Information

CSI volume snapshots

## Cloudera Base on premises requirements

Your Cloudera Base on premises cluster must have the operating system, JDK, database, Cloudera components, and Cloudera Runtime version required to install Cloudera Data Services on premises.

Operating system, JDK, and database:

- See Cloudera Private Cloud Base Requirements and Supported Versions

**Note:**

- Ozone is not a mandatory requirement for installing Cloudera AI without the model registry.
- Configure Atlas with Hive for Cloudera Data Engineering to work properly.

The PostgreSQL database instance must be configured to accept inbound TLS requests to the Hive Metastore database. A TLS connection is required when initiated from Cloudera Data Warehouse in OpenShift.

Cloudera Runtime components (services):

- Hive Metastore (HMS)
- Ranger
- Atlas
- HDFS
- Ozone
- YARN
- Kafka
- Solr

Additionally, do the following:

- Set up Kerberos on these clusters using an Active Directory.
- Enable TLS on the Cloudera Manager cluster for communication with components and services.
- Ensure that the Cloudera Private Cloud Base cluster is on the same network as the OpenShift cluster.
- Configure PostgreSQL database as an external database for the Cloudera Private Cloud Base cluster components.
- Configure the Cloudera Private Cloud Base cluster hostnames to be forward and reverse resolvable in DNS from the OpenShift cluster.
- Allow websocket traffic and https traffic when you use a load balancer with the OpenShift external API.
- Ensure `hive` user is able to create and list an Ozone bucket. For information about creating and listing ozone bucket, see *Managing buckets*.

You can use the Cloudera Management Console to create one or more environments. These environments can be associated with any of the Data Lake from the Cloudera Private Cloud Base clusters. The Cloudera Private Cloud Base Cloudera Manager deploys the Cloudera Management Console.

Cloudera currently does not support associating an environment with many Cloudera Private Cloud Base cluster installations.

### Related Information

Managing buckets

## Preparing Cloudera Base on premises

Use Cloudera Manager to configure your Cloudera Base on premises in preparation for the Cloudera Data Services on premises installation.

### Procedure

1. Configure the Cloudera Base on premises cluster to use TLS.

   For configuration steps, see Configuring TLS Encryption for Cloudera Manager Using Auto-TLS.

2. Configure Cloudera Manager with a JKS-format (not PKCS12) TLS truststore.

   For configuration steps, see Database requirements.
3. Configure Cloudera Manager to include a root certificate that trusts the certificate for all Cloudera Manager server hosts expected to be used with on premises.

   a. Import the necessary certificates into the truststore configured in Configure Administration Settings Security TLS/SSL Client Trust Store File .

   **Note:** This requires a Cloudera Manager restart.
4. Configure Ranger and LDAP for user authentication. Ensure that you have configured Ranger user synchronization.

   For configuration steps, see Configure Ranger authentication for LDAP and Ranger usersync.

   **Note:** Upgrading to Oracle JDK 1.8.351 causes a Kerberos issue when deprecated 3DES and RC4 permitted encryption types are used.

   Workaround: Remove the deprecated 3DES and RC4 encryption types in the krb5.conf and kdc.conf files.
5. Enable Kerberos for all the services in the cluster.

   For configuration steps, see Enabling Kerberos for authentication.

   **Note:** To avoid connection issues from pods to FreeIPA leading to data services install issues, Cloudera recommends to update the krb5.conf file with the following value:

   ```
   udp_preference_limit = 1
   ```
6. Configure LDAP using Cloudera Manager. Only Microsoft Active Directory (AD) and OpenLDAP are currently supported.

   For configuration steps, see Configure authentication using an LDAP-compliant identity service.
7. Check if all the running services in the cluster are healthy. To check this using Cloudera Manager, go to Cloudera Manager Clusters [***CLUSTER NAME***] Health Issues . If there are no health issues, the No Health Issues message is displayed.
8. Verify if you have the necessary Cloudera entitlements from Cloudera to access the on premises installation. To check this using Cloudera Manager, go to Private Cloud Select Repository [***REPOSITORY URL***] . If you have the required entitlements, the You are about to install Cloudera on premises version [*VERSION*] message with a list of prerequisites is displayed. An error message is displayed if you do not have the necessary entitlements.

   Contact your Cloudera account team to get the necessary entitlements.
9. If you want to reuse data from your legacy CDH or HDP deployment in your on premises, ensure that you have migrated that data into your Cloudera Base on premises. You must be using Cloudera Runtime 7.1.9 for migrating your data from your CDH or HDP cluster.

   For more information about data migration, see the Data Migration Guide.
10. For installing Cloudera Base on premises, see Install Cloudera Base on premises

## Cloudera Data Services on premises Hardware Requirements

You must learn about the minimum and recommended hardware and network infrastructure requirements before deploying Cloudera Data Services on premises.

Architects and infrastructure administrators must understand these requirements to install Cloudera Data Services on premises in your data center.

You must know the minimum hardware requirements prior to:

- Installing a dedicated Red Hat OpenShift Container Platform cluster required for Cloudera on premises
- Installing and configuring Cloudera Data Services on premises

- Deploying and running the Cloudera Data Warehouse and Cloudera AI Data services

### Related Information

Cloudera Data Warehouse hardware requirements

## Cloudera Data Services on premises deployment considerations

You must understand the deployment requirements to sufficiently provision node counts, CPU, memory, and other hardware resources required to install Cloudera on premises.

The Cloudera Data Services on premises are installed on the OpenShift Cluster and run on the provisioned worker nodes. Cloudera Data Services on premises deployment consists of a Cloudera Management Console on premises and one or more environments that are created for deploying the Data Services. The Cloudera Management Console is a service used by Cloudera administrators to manage environments, users, and services.

The worker node hardware requirements are described below. The number of worker nodes needed depends on factors such as the number of virtual warehouses or machine learning workspaces required for your workloads. The recommendation here is a guideline for a basic Cloudera Data Services on premises installation. For hardware sizing in production environments, contact Cloudera Support or your Cloudera Account Team.

| Component | Minimum | Recommended |
|---|---|---|
| Node Count | 10 | 20 |
| CPU | 16 | 32 + |
| Memory | 128 GB | 384 GB |
| Storage | 2 TB (SATA) | 4 TB (SSD/NVMe) |
| Network Bandwidth | 1 Gbps guaranteed bandwidth (minimum) dedicated to every Cloudera Base on premises node | 10 Gbps guaranteed bandwidth (minimum) dedicated to every Cloudera Base on premises node |

**Important:**

- You need a Cluster Admin role for OpenShift system administration.
- You need the bootstrap node for the initial installation. It can be converted into an OpenShift worker after initial deployment.

To know about architecture, design choices, and deployment guidelines to use Cloudera Data Services on premises with Dell EMC and Intel Infrastructure, and Cisco Intelligent Data Platform, see Dell EMC and Intel Infrastructure Guide for Cloudera Data Services on premises and Cisco Data Intelligence Platform on Cisco UCS C240 M5 with Cloudera Data Services on premises Plus Design Guide.

## Storage requirements

Learn about the storage requirements necessary for deploying Cloudera Data Engineering, Cloudera Data Warehouse, Cloudera Control Plane, Cloudera Data Catalog, and Cloudera AI.

### Storage Requirements

| Data Services | Storage type | Storage required | Purpose |
|---|---|---|---|
| Cloudera Data Engineering | Block | 500 GB per Virtual Cluster | Stores all information related to virtual clusters |
| Cloudera Data Warehouse | Local | 100 GB per executor in REDUCED profile mode and 600 GB per executor in FULL mode | Used for caching |

| Data Services | Storage type | Storage required | Purpose |
|---|---|---|---|
| Cloudera Control Plane | Block | 118 GB total if using an External Database, 318 GB total if using the Embedded Database (SSD support only) | Storage for Cloudera infrastructure including Fluentd logging, Prometheus monitoring, and Vault. Backing storage for an embedded DB for control plane configuration purpose, if applicable |
| Cloudera AI | Block | 600 GB per node (minimum), 4.5 TB (recommended) | Stores all Cloudera AI Workbench information |
| | External NFS or Block | 1 TB per workbench (dependent on the size of the Cloudera AI user files) | Stores all user project files. |
| MonitoringApp | Block | 30 GB + (Env cnt x 100 GB) | Stores metrics collected by Prometheus. |
| Cloudera Data Catalog | Requires Control Plane database and not a dedicated storage space | 100 GB extra in Cloudera Control Plane database | Stores profiling metadata. |

## Cloudera Data Services on premises network infrastructure considerations

Learn about the networking infrastructure consideration necessary to install Cloudera on premises. The networking considerations for Cloudera Data Services on premises are similar to the networking requirements for Cloudera Manager Virtual Private Clusters (CM VPC).

In Cloudera Data Services on premises, the network bandwidth requirements are less stringent than those of the Cloudera Manager Virtual Private Cluster (VPC) because of data caching technology introduced at the compute layer, which is not available in VPCs.

While the initial load of data from the remote storage would require significant bandwidth between the compute and storage clusters, subject to the quantity of data ingested; subsequently, the network bandwidth requirements are lower.

The following list of network considerations will help you plan your network infrastructure before you install Cloudera Data Services on premises:

- Use 1 Gbps guaranteed bandwidth between each OpenShift worker node and each Cloudera Base on premises DataNode. Cloudera recommends 10 Gbps guaranteed bandwidth.
- Stress test the network infrastructure with all the OpenShift nodes trying to read or write from the Cloudera Data Services on premises nodes at the same time.
- Use the Spine-Leaf network architecture with no more than a 4:1 oversubscription between the spine and leaf switches.
- Check the applicable ports used by Cloudera Runtime components.

For more information about minimum network performance requirements, network sizing, and designing a network topology, see Networking Considerations for Virtual Private Clusters.

## Cloudera Data Services on premises Software Requirements

You must learn about the software and configuration requirements before deploying Cloudera on premises. Administrators and operators must understand these requirements to install Cloudera Data Services on premises in your data center.

You must understand the following software requirements before you install Cloudera on premises:

- OpenShift integration requirements
- Cloudera Base on premises requirements
- External database requirements

- External vault requirements

## External vault requirements

You can learn about how to configure an external vault, build on HashiCorp, for Cloudera Data Services on premises. Hashicorp Vault securely stores your passwords, tokens, certificates, and encryption keys.

Cloudera supports all the external vaults that uses Hashcorp.

**Note:** Vault namespaces are not supported.

### Vault Token Policy

Cloudera Data Services on premises can be installed using an internal or external Vault. If you are installing Cloudera Data Services on premises with an external Vault, a Vault token with the following permissions is required.

- Create/Update/List/Read a secret engine of type kv-2 at the applicable path.
- Create/Update/List/Read auth of type kubernetes at the applicable path.
- Create/Update/List/Read policies.
- Access to List and Read the Vault token details.

Example Vault policy:

```
# Manage auth methods broadly across Vault
path "auth/*"
{
  capabilities = ["create", "read", "update", "list"]
}
# Create, update auth methods
path "sys/auth/*"
{
  capabilities = ["create", "update", "sudo"]
}

# List auth methods
path "sys/auth"
{
  capabilities = ["read"]
}

# List existing policies
path "sys/policies/acl"
{
  capabilities = ["list"]
}

# Create and manage ACL policies via API & UI
path "sys/policies/acl/*"
{
  capabilities = ["create", "read", "update", "list"]
}


# Manage secrets engines
path "sys/mounts/*"
{
  capabilities = ["create", "read", "update", "list"]
}

# List existing secrets engines.
path "sys/mounts"
{
```

```
  capabilities = ["read"]
}
```

For more information, see HashiCorp Vault Policy Requirements.

### Vault Token Use

The Vault token should be created using the preceding policy. It is recommended that the Vault administrator delete this token after the installation is complete.

### External Vault Installation Parameters

- Vault Address – The external Vault FQDN (Fully Qualified Domain Name) with the port number.
- Token – The Vault token described above
- CA Certificate – A valid certificate for the Vault server in PEM format.

### Vault Secrets Engine, Auth, and Policies

During installation, Cloudera data platform enables a kv-v2 secrets engine and kubernetes authentication at unique paths in the following format:

```
cloudera-[***CONTROL PLANE NAMESPACE***]-[***SERVER-URL***]
```

It is recommended that you do not have any kv-v2 secrets and kubernetes auth enabled at the same path in your Vault server.

Cloudera data platform also creates Vault policies that provide access to control plane services to write their protected data. These two policies have the following format:

```
[***NAMESPACE***]-[***SERVER URL***]
```

```
admin-[***NAMESPACE***]-[***SERVER URL***]
```

## Docker repository access

You must ensure that the cluster has access to the Docker Container Repository in order to retrieve the container images for deployment.

There are several types of Docker Repositories you can use:

**Cloudera Repository**

Using the Cloudera Repository requires that the cluster have internet connectivity to the Cloudera public repository. Using the Cloudera Repository is the fastest option.

The Cloudera-hosted Docker Repository option may increase the time required to deploy or start the services in the cluster. Cloudera generates Docker Repository credentials that are identical to your paywall credentials. Refer to your welcome letter for the credentials or use the credential generator on Cloudera.com to generate credentials from your license key.

This option is best suited for proof-of-concept, non-production deployments or deployments that do not have security requirements that disallow internet access.

**Custom Repository**

A Custom Repository is a repository that you manage in your environment and can be Enterprise grade and highly available.

During installation and upgrade, a custom script is generated that you use to copy the images. Copying images can take 4 - 5 hours.

Only TLS-enabled custom Docker Registry is supported. Ensure that you use a TLS certificate to secure the custom Docker Registry. The TLS certificate can be self-signed, or signed by a on premises or on cloud trusted Certificate Authority (CA).

**Important:** When using an Cloudera Embedded Container Service cluster, passwords must not contain the $ character.

### Related Information

Installing Cloudera Data Services on premises with OpenShift Container Platform

Installation using the Cloudera Embedded Container Service

## Cloudera AI on premises software requirements

To launch the Cloudera AI service, the on premises host must meet several software requirements. Review the following Cloudera AI-specific software requirements.

### Requirements

**Note:**

Only the usage of SSD disks is supported with Cloudera on premises on OpenShift Container Platform.

If needed, reach out to your Administrator to ensure the following requirements are met.

Compatibility requirements

- If you use OpenShift, check that the version of the installed OpenShift Container Platform is exactly as listed in Software Support Matrix for OpenShift.

Required accesses

- If Cloudera AI needs access to a database on the Cloudera Base on premises cluster, then the user must be authenticated using Kerberos and must have Ranger policies set up to allow read/write operations to the default (or other specified) database.
- Ensure that Kerberos is enabled for all services in the cluster. Custom Kerberos principals are not supported currently. For more information, see Authenticating Hue users with Kerberos.
- Cloudera AI assumes it has cluster-admin privileges on the cluster.
- If external NFS is used, the NFS directory and assumed permissions must be those of the cdsw user. For details see Using an External NFS Server.
- If you intend to access a workbench over https, see Deploying a Cloudera AI Workbench with support for TLS.

Requirements for functioning

- On OpenShift Container Platform, CephFS is used as the underlying storage provisioner for any new internal workbench on Cloudera on premises 1.5.x. A storage class named ocs-storagecluster-cephfs with csi driver set to openshift-storage.cephfs.csi.ceph.com must exist in the cluster for new internal workbenches to get provisioned.
- A block storage class must be marked as default in the cluster. This may be rook-ceph-block, Portworx, or another storage system. Confirm the storage class by listing the storage classes (run `oc get sc`) in the cluster, and check that one of them is marked default.

DNS-related requirements

- Forward and reverse DNS must be working.
- DNS lookups to sub-domains and the Cloudera AI Workbench itself shall work properly.
- In DNS, wildcard subdomains (such as *.cml.yourcompany.com) must be set to resolve to the master domain (such as cml.yourcompany.com). The TLS certificate (if TLS is used) must also include the wildcard subdomains. When a session or job is started, an engine is created for it, and the engine is assigned to a random, unique subdomain.

Configuration requirements

- The external load balancer server timeout needs to be set to 5 min. Without this, creating a project in a Cloudera AI Workbench with `git clone` or with the API may result in API timeout errors. For workarounds, see *Known Issue DSE-11837*.
- For non-TLS Cloudera AI Workbench, websockets need to be allowed for port 80 on the external load balancer.
- Only a TLS-enabled custom Docker Registry is supported. Ensure that you use a TLS certificate to secure the custom Docker Registry. The TLS certificate can be self-signed, or signed by a on premises or on cloud trusted Certificate Authority (CA).
- On OpenShift, due to a Red Hat issue with OpenShift Container Platform 4.17, the image registry cluster operator configuration must be set to Managed.
- Check if storage is set up in the cluster image registry operator. See *Known Issues DSE-12778* for further information.

For more information on requirements, see Install Cloudera Base on premises.

# Installing Cloudera Data Services on premises with OpenShift Container Platform

Learn about installing Cloudera Data Services on premises with the OpenShift Container Platform (OCP).

## Cloudera Data Services on premises pre-installation checklist

Before starting the installation, you must ensure that you have configured all the required hardware and software. There are several pre-installation tasks that you must complete using Cloudera Manager and OpenShift Container Platform.

**Note:**

For Cloudera Data Services on premises running on OCP, UTC is the recommended timezone setup for nodes and pods under OpenShift environment. This is based on RedHat recommendations: Setting timezone for pods in OpenShift Container Platform.

Use the following checklists to ensure that you have completed all the pre-installation tasks:

- Cloudera Base on premises
- OpenShift Container Platform
- Cloudera Data Warehouse
- Cloudera AI
- Cloudera Data Engineering

### Cloudera Base on premises checklist

Use this checklist to ensure that your Cloudera Base on premises is configured and ready for installing Cloudera Data Services on premises.

**Note:** The Cloudera Manager mentioned in this checklist is the Cloudera Base on premises using the Cloudera Manager which you want to install Cloudera Data Services on premises on.

**Table 7: Cloudera Base on premises checklist to install Cloudera Data Services on premises**

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| Runtime components | Ensure that you have Ranger, Atlas, Hive, HDFS, and Ozone installed in your Cloudera Base on premises. | • Software Support Matrix for OpenShift on page 4<br>• Cloudera Base on premises requirements | If you do not install these components, you see an error when creating an environment in Cloudera Data Services on premises. |

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| Network requirement | Ensure that all the network routing hops in production. Cloudera recommends not to use more than 4:1 oversubscription between the spine-leaf switches. | | |
| Cloudera Manager database requirement | Refer to the the Cloudera Base on premises database requirements. | • Database Requirements<br>• Cloudera Support Matrix | N/A |
| Cloudera Manager TLS configuration | Ensure that Cloudera Manager in the Cloudera Base on premises cluster is configured to use TLS. | Configuring TLS Encryption for Cloudera Manager Using Auto-TLS | You can also manually configure TLS to complete this task. See<br>Manually Configuring TLS Encryption for Cloudera Manager |
| Cloudera Manager JKS-format TLS truststore | Ensure that the Cloudera Manager is configured with a JKS-format (not PKCS12) TLS truststore. | Obtain and Deploy Keys and Certificates for TLS/SSL | N/A |
| Cloudera Manager truststore and root certificate | Ensure that the Cloudera Manager truststore contains a root certificate that trusts the certificate for all Cloudera Manager server hosts used with Cloudera Data Services on premises. | How to Add Root and Intermediate CAs to Truststore for TLS/SSL | Import the necessary certificates into the truststore configured in Configure Administration > Settings > Security > Cloudera ManagerTLS/SSL Client Trust Store File . |
| LDAP configuration | Ensure that you configure LDAP using Cloudera Manager. | N/A | Only Microsoft Active Directory (AD) and OpenLDAP are currently supported. |
| Apache Ranger configuration for LDAP | Ensure that the Cloudera Base on premises cluster is configured with Apache Ranger and LDAP for user authentication. | Configure Ranger authentication for LDAP | N/A |
| Apache Ranger usersync configuration | Ensure that you have configured Apache Ranger and Apache Ranger usersync. | Ranger usersync | Apache Ranger user synchronization is used to get users and groups from the corporate ActiveDirectory to use in policy definitions. |
| Kerberos configuration | Ensure that Kerberos is enabled for all services in the cluster. | Enabling Kerberos for authentication | Custom Kerberos principals are not currently supported. |
| Internet access or air gap installation | Ensure that Cloudera Base on premises and the Cloudera Embedded Container Service hosts have access to the Internet. If you do not have access to the Internet, you must do an air gap installation. | Install Cloudera Data Services on premises in air gap environment | You need access to the Docker registries and the Cloudera repositories during the installation process. |
| Services health check | Ensure that all services running in the cluster are healthy. | Cloudera Manager Health Tests | N/A |
| Cloudera on premises entitlement | Ensure that you have the necessary Cloudera entitlement to access the on premises installation. | N/A | |
| Reuse data from CDH or HDP (Optional) | To reuse data from your legacy CDH or HDP deployment in your on premises, ensure that you have migrated that data into your Cloudera Base on premises. You must be using Cloudera Runtime 7.1.7 for migrating data from your CDH or HDP cluster. | Data Migration Guide | N/A |

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| (Recommended) Configure HDFS properties to optimize logging | Cloudera uses “out_webhdfs” Fluentd output plugin to write records into HDFS, in the form of log files, which are then used by different data services to generate diagnostic bundles. To optimize the size of logs that are captured and stored on HDFS, you must update a few HDFS configurations in the hdfs-site.xml file using Cloudera Manager. | Configuring HDFS properties to optimize logging | N/A |

## OpenShift Container Platform (OCP) Checklist

Use this checklist to ensure that your OpenShift Container Platform (OCP) is configured and ready for installing Cloudera Data Services on premises.

### Table 8: OpenShift Container Platform (OCP) checklist to install Cloudera Data Services on premises

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| Network requirements | | Validate if there were any stale/ predefined IP tables that interfere with Kubernetes cluster/service IP ranges. | Cloudera Data Services on premises requires a single ethernet interface. Multihoming is currently not supported. |
| OpenShift Platform version | Check the the installed OpenShift Container Platform version. | • OpenShift software requirements<br>• Software Support Matrix for OpenShift on page 4 | N/A |
| DNS configuration | Ensure that you have set up the DNS and Reverse DNS between OpenShift Container Platform (OCP) hosts and Cloudera Base on premises. This is required for obtaining Kerberos ticket-granting tickets. | Certificate management and DNS | |
| Check if you can access the OpenShift hostnames outside the cluster | Ensure that OpenShift Container Platform (OCP) application hostnames can be accessed from outside the cluster. | A minimal Ingress resource example | Perform a DNS query on the route generated, to check if you can access the hostnames outside the cluster. |
| Storage classes configuration | Ensure that you have configured separate storage classes for the control plane and the compute clusters. Both the storage classes must be provisioned from Persistent Volumes. | Storage classes | N/A |
| OpenShift Container Platform (OCP) Kubeconfig file | Ensure that you have access to the OpenShift Container Platform (OCP) Kubeconfig file, cluster administrator privileges, and sufficient expiry time for you to complete your installation. | Download Kubernetes Configuration | The kubeconfig should have valid certificates in it for the cluster. If the kubeconfig does not have certificates, then the you must upload custom certifications during Cloudera Data Services on premises installation. |

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| Allow WebSocket traffic in addition to HTTPS | When a load balancer is used for your OpenShift Container Platform external API, you must allow WebSocket traffic in addition to HTTPS. The load balancer must allow WebSockets on port 80. Also, ensure that you set the load balancer server timeout to 5 minutes. | N/A | N/A |
| Clock time from NTP source | Ensure that the NTP clock in Cloudera Base on premises is in sync with the time configured in the OpenShift Container Platform (OCP) cluster. This is an important step if your setup does not have access to the Internet. | Enable an NTP Service | Install Cloudera Data Services on premises in air gap environment |
| Route admission policy | Ensure OpenShift Container Platform (OCP) cluster is configured to run applications in multiple namespaces with the same domain name. | Configuring the route admission policy | N/A |

## Cloudera Data Warehouse checklist

Use this checklist to ensure that you have all the requirements for Cloudera Data Warehouse in Cloudera Data Services on premises.

### Table 9: Cloudera Data Warehouse installation checklist for Cloudera Private Cloud Data Services

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| OpenShift requirements | Ensure that you have the required memory, storage, and hardware requirements for getting started with the Cloudera Data Warehouse service on Red Hat OpenShift. | OpenShift requirements | N/A |
| Security requirements | Ensure that you have all the security requirements needed to install and run the Cloudera Data Warehouse service on Red Hat OpenShift clusters. | Security requirements for Cloudera Data Warehouse Private Cloud | N/A |
| Database requirements | Ensure that you fulfill the requirements for the database that is used for the Hive Metastore on the base cluster (Cloudera Manager side) for Cloudera Data Warehouse on premises. | Database requirements | N/A |

## Cloudera AI checklist

Use this checklist to ensure that you have all the requirements for Cloudera AI in Cloudera Data Services on premises.

### Table 10: Cloudera AI installation checklist for Cloudera Data Services on premises

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| Network File System (NFS) support | Ensure that you have either configured an external or an embedded NFS. | Cloudera AI requirements | N/A |

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| NFS Provisioner | When using OCP 4.17, the supported NFS versions are 3.0, 4.0, 4.1, and 4.2. | | |
| Ranger policy configuration | Ensure that the user who is authenticated using Kerberos has Ranger policies that are configured to allow read/write to the default (or other specified) databases. | Cloudera AI requirements | N/A |

## Cloudera Data Engineering checklist

Use this checklist to ensure that you have all the requirements for Cloudera Data Engineering in Cloudera Data Services on premises.

### Table 11: Cloudera Data Engineering installation checklist for Cloudera Data Services on premises

| Item | Summary | Documentation | Notes |
|---|---|---|---|
| Ozone in Base cluster | For workloads to store logs, Ozone in Base cluster is a must. Ensure Ozone is installed on Cloudera Base on premises cluster. | Install Cloudera Base on premises | N/A |
| Ranger policy configuration | Ensure that the user who is authenticated using Kerberos needs to have Ranger policies that are configured to allow read/write to the default (or other specified) databases. | Kerberos authentication for Apache Ranger | N/A |

# Installing in internet environment

Follow the steps in this topic to install Cloudera on premises.

### Before you begin

- Ensure that your Kubernetes kubeconfig has permissions to create Kubernetes namespaces.
- You require persistent storage classes defined in your OpenShift cluster. Storage classes can be defined by OpenShift cluster administrators.
- Only TLS-enabled custom Docker Registry is supported. Ensure that you use a TLS certificate to secure the custom Docker Registry. The TLS certificate can be self-signed, or signed by a private or public trusted Certificate Authority (CA).
- Only TLS 1.2 is supported for authentication with Active Directory/LDAP. You require TLS 1.2 to authenticate the Cloudera Control Plane with your LDAP directory service like Active Directory.
- OCP network configurations that restrict pod communication are not supported. For example, multi-tenancy isolation with network policy is not supported.

## Procedure

1. In Cloudera Manager, on the top right corner, click Add > Add Cluster. The Select Cluster Type page appears.

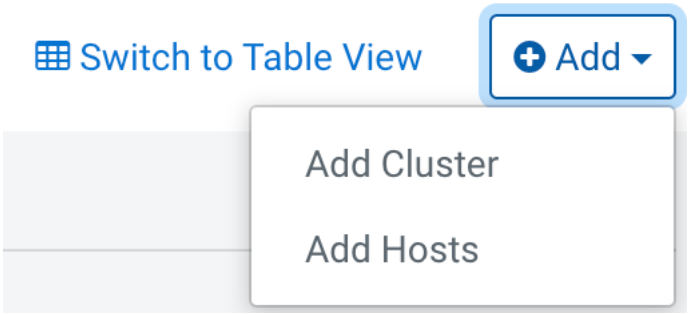


2. On the Select Cluster Type page, select the cluster type as **Cloudera on premises** Containerized Cluster. Under Other Options, click here to install Cloudera Data Services on premises, then click Continue.

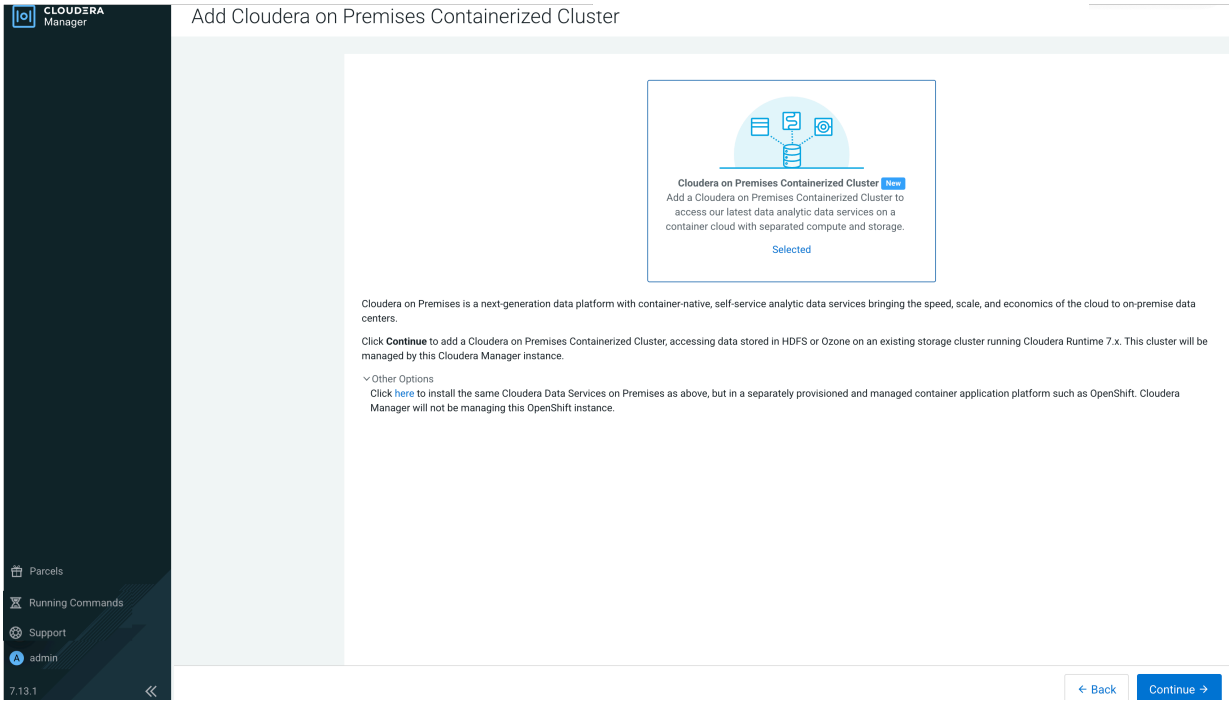

3. On the Getting Started page of the installation wizard, select Internet as the Install Method. To use a custom repository link provided to you by Cloudera, click Custom Repository. Click Next.

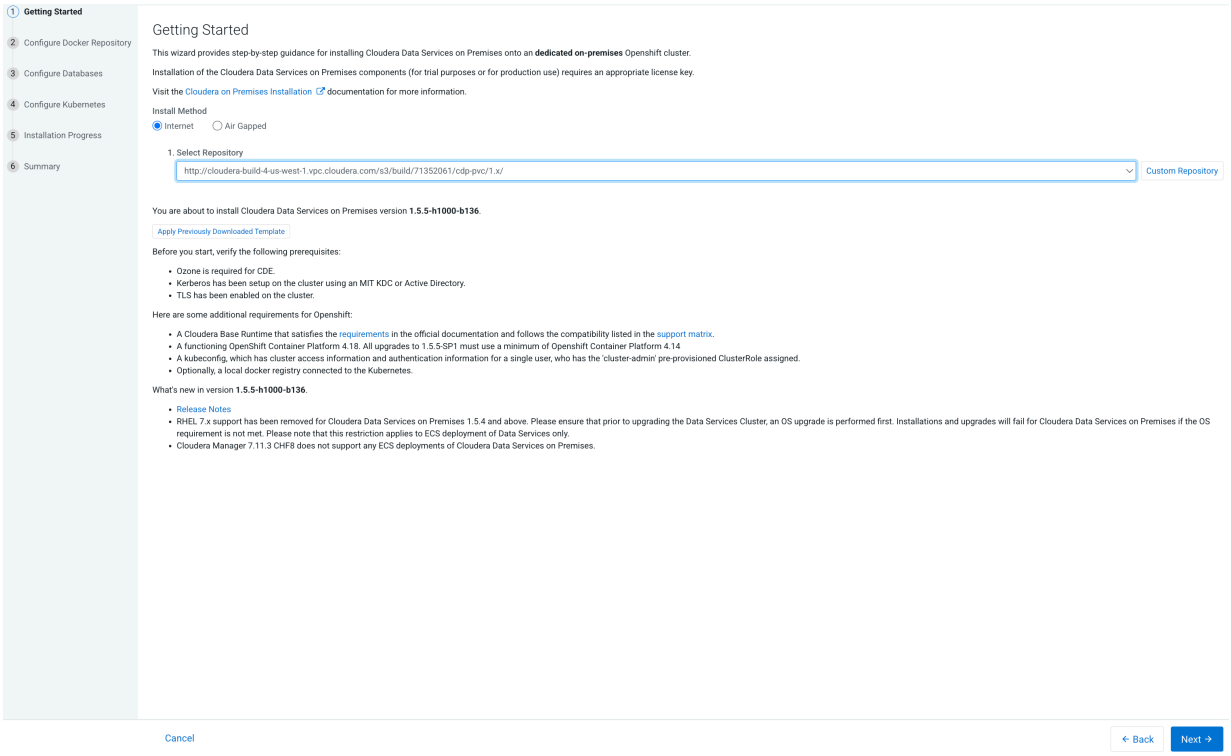


**Note:**

- Verify the prerequisites for the version that you're installing and then click Next.
- You can also apply a template that you may have downloaded during a previous installation. The template contains all the installation configurations. Click Apply Previously Download Template to browse and upload a template stored on your machine.

4. On the Configure Docker Repository page, you must select one of the Docker repository options. If you select Use a custom Docker Repository option, enter your local Docker Repository in the Custom Docker Repository field in the following format:*[*DOCKER REGISTRY*]/[*REPOSITORY NAME*]*. Alternatively, you can use Cloudera default Docker Repository if you are setting up Cloudera on premises in non-production environments.

**Note:**

- Use a custom Docker Repository - Copies all images (Internet or Air Gapped) to the embedded registry
- Use Cloudera default Docker Repository - Copies images from Internet to the embedded registry. This uses the default repository that is in manifest.json. Use Cloudera default Docker Repository option can be selected only if you have selected Internet as the install method.

You can follow these steps to prepare your Docker Repository from a machine that is running Docker locally and has access to all the Docker images either directly from Cloudera or a local HTTP mirror in your network.

a) Click Generate the copy-docker script on the wizard or download the script file.
b) Log in to your custom Docker Registry and run the script using the following commands.

```
docker login <your_custom_registry> -u <user_with_write_access>
```

```
bash copy-docker.txt
```

**Note:** This command downloads 100+ Docker images and it will take some time to download.

c) Enter your Docker user name and password.
d) Click Choose File to upload your Docker certificate.
e) Click Next.

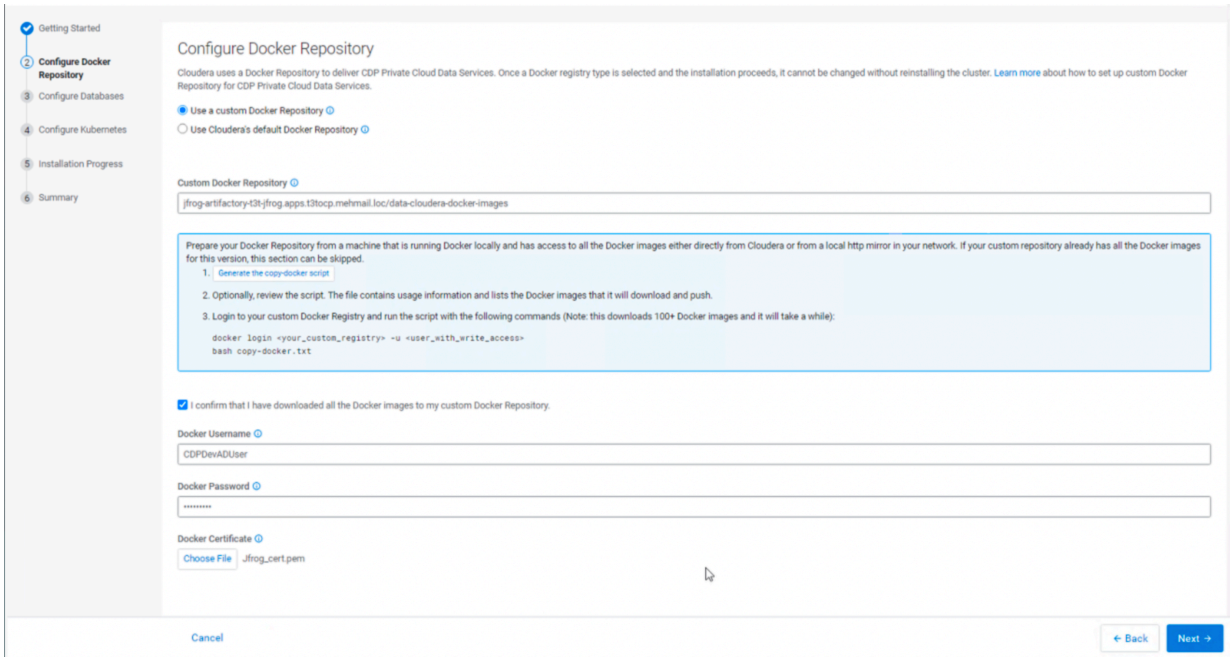


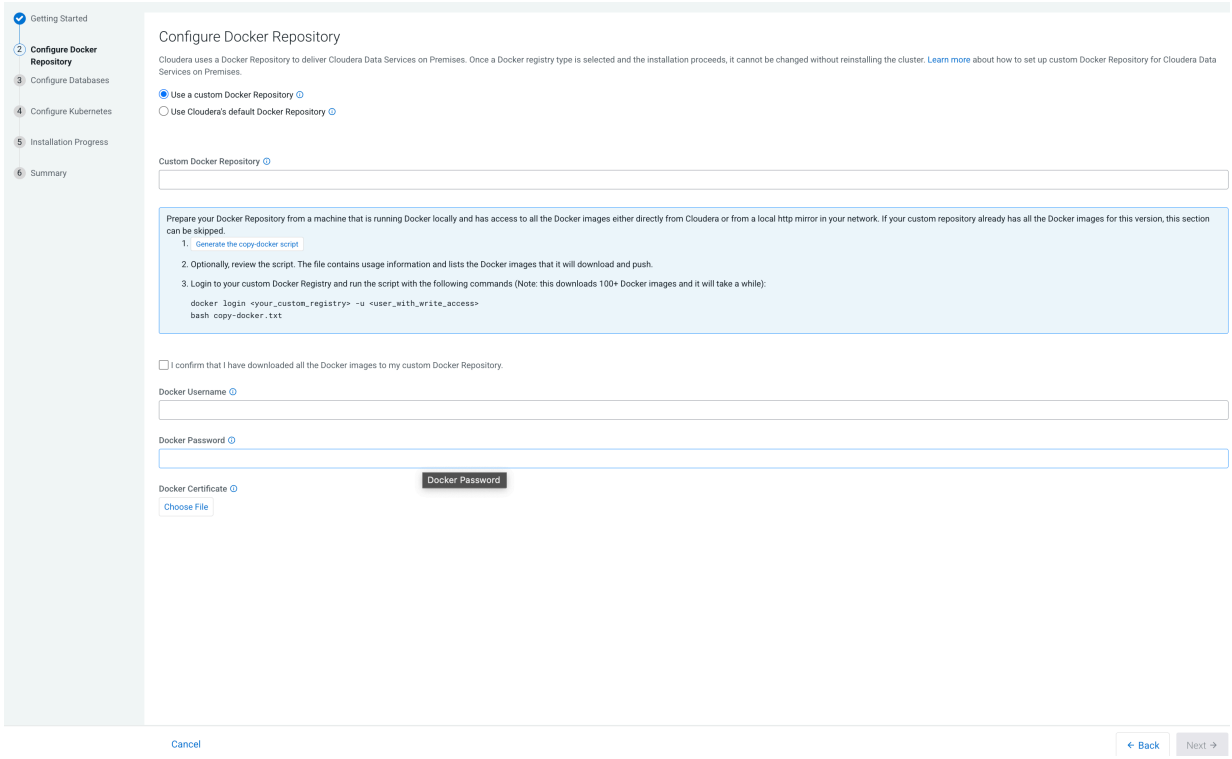

5. On the Configure Databases page, click Next.

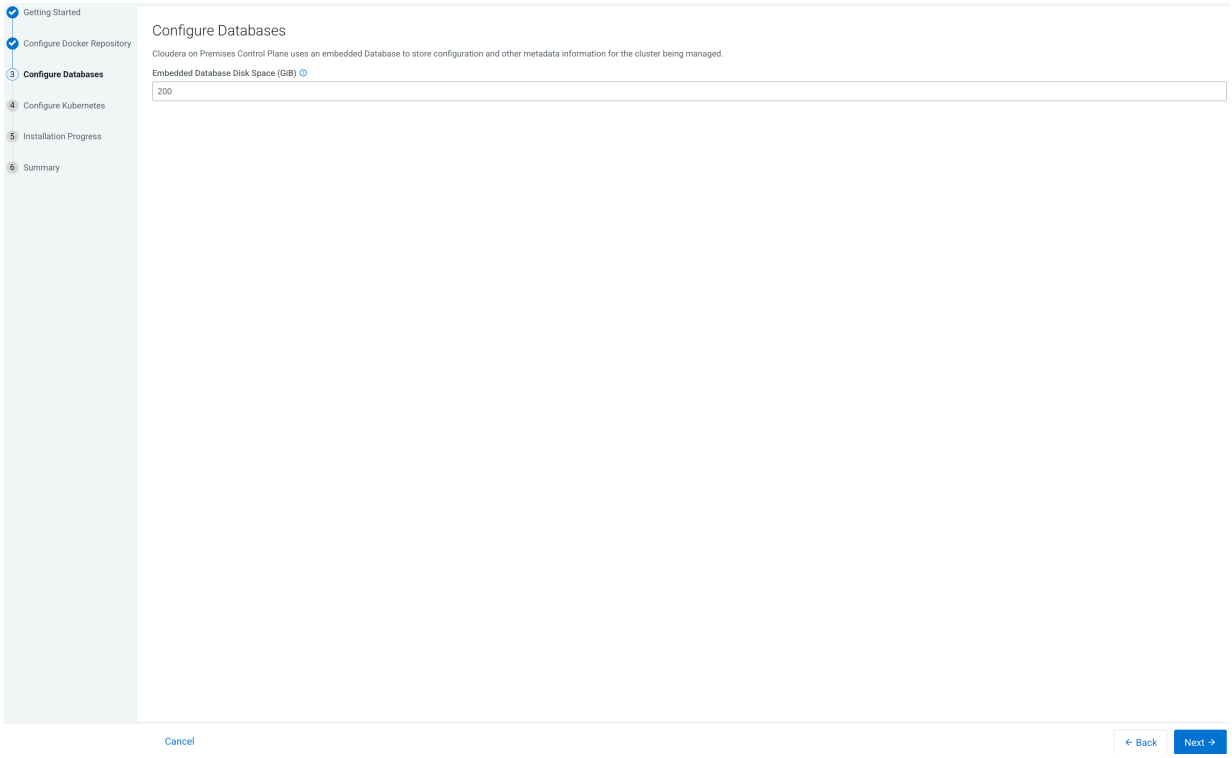


6. On the Configure Kubernetes page, enter your Kubernetes, Docker, database, and vault information.
   a) Upload a Kubernetes configuration (kubeconfig) file from your existing environment. You can obtain this file from your OpenShift Container Platform administrator. Ensure that this kubeconfig has permissions to create Kubernetes namespaces.
   b) In the Kubernetes Namespace field, enter the Kubernetes namespace that you want to use with this Cloudera on premises deployment. Kubernetes virtual clusters are called namespaces. For more information, see Kubernetes namespaces
   c) Enter your Vault information and upload a CA certificate. Cloudera recommends that you use an external Vault for production environments. Enter the Vault address and token, and upload a CA certificate.
   d) Enter a Storage Class to be configured on the Kubernetes cluster. Cloudera on premises uses Persistent Volumes to provision storage. You can leave this field empty if you have a default storage class configured on your Openshift cluster. Click Continue.
   e) Under the Additional Certificates section, click Choose File and add the SSL certificate for your HMS database (MariaDB, MySQL, PostgreSQL, or Oracle). For Cloudera Data Warehouse, it is mandatory to

secure the network connection between the default Database Catalog Hive MetaStore (HMS) in Cloudera Data Warehouse and the relational database hosting the base cluster's HMS.

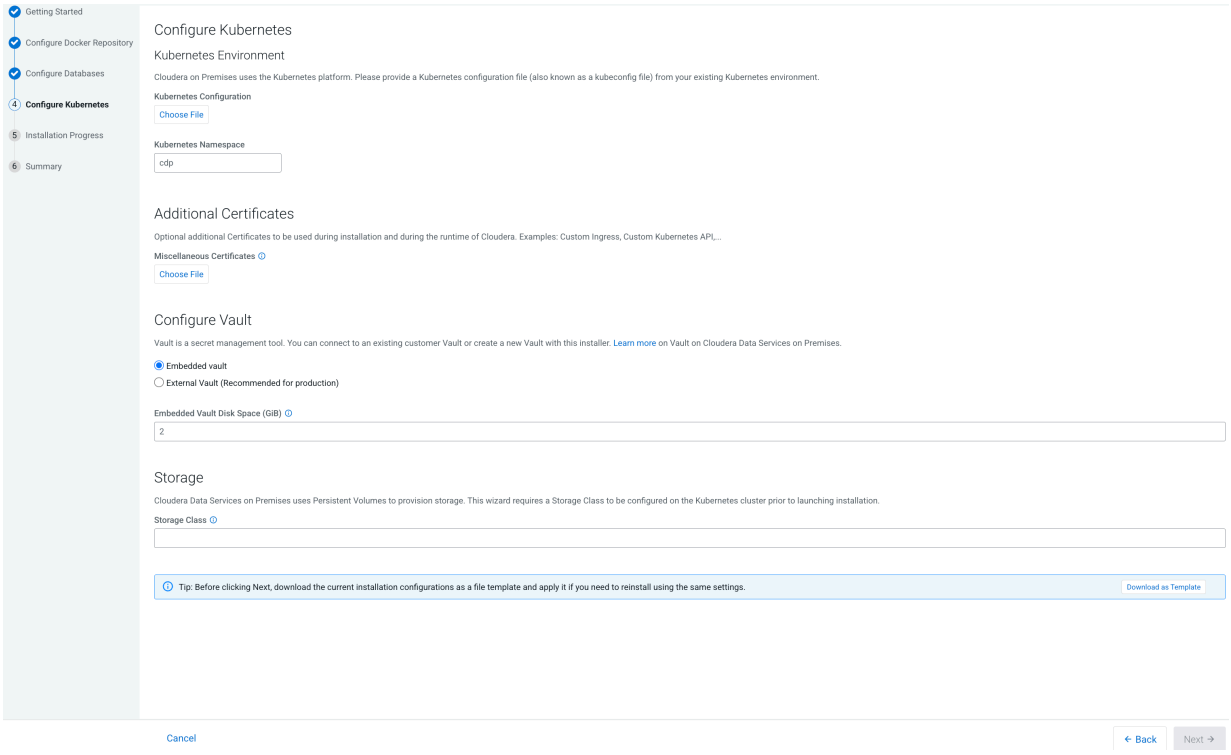


7. If you want to use this installation configuration again to install Cloudera on premises, you have the option to download this information as a template.

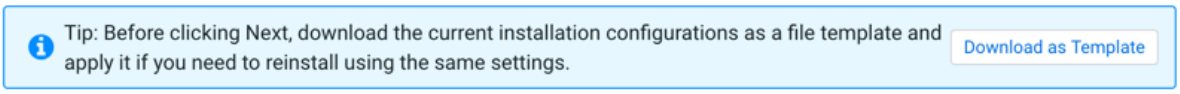


The template file is a text file that contains the database and vault information that you entered for this installation. This template is useful if you will be installing on premises again with the same databases, as the template will populate the fields here automatically. Note that the user password information is not saved in the template.

8. The Installation Progress page appears. When the installation is complete, click Next.

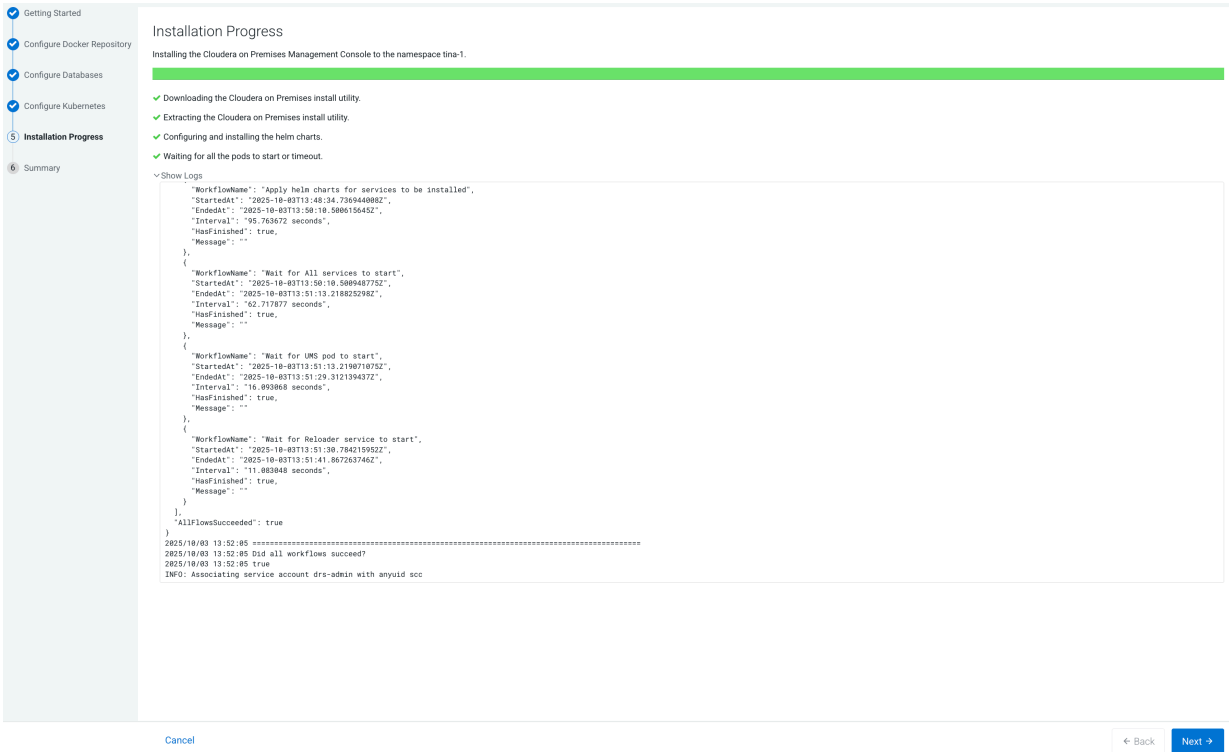


9. The summary message with a link to Launch **Cloudera on premises** appears.

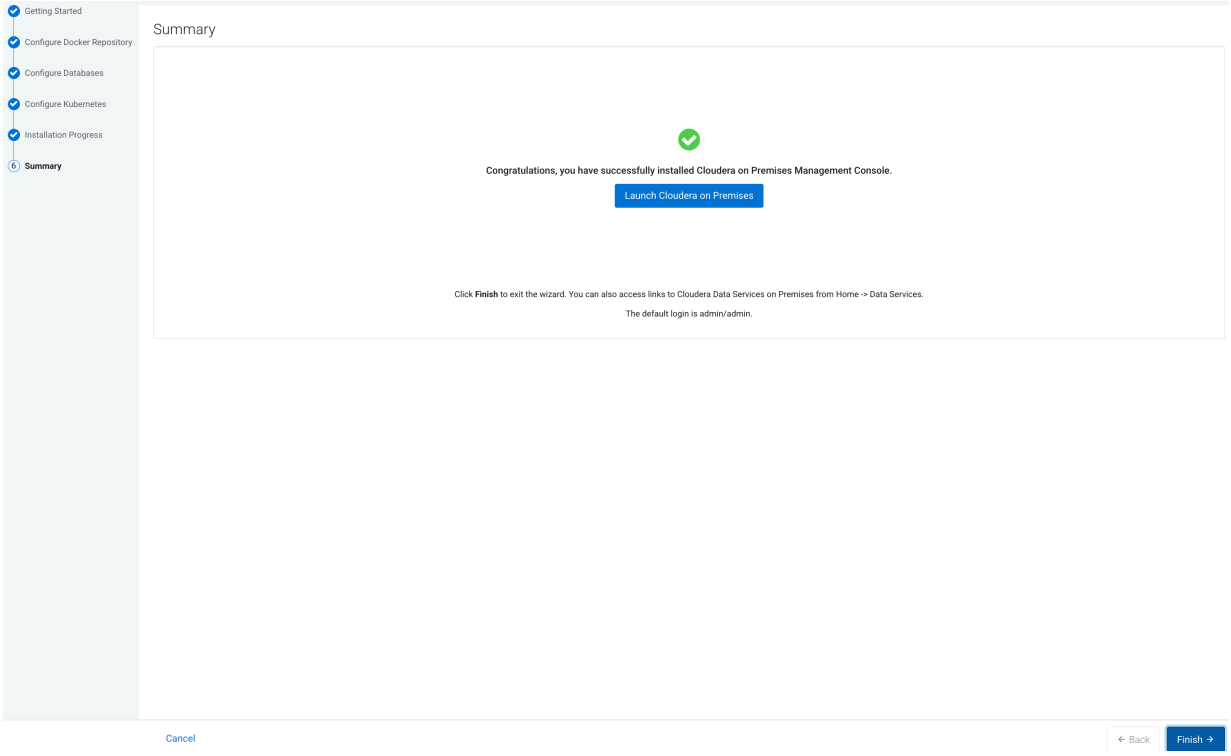


## What to do next

- Click Launch **Cloudera on premises** to launch your Cloudera on premises.
- Log in using the default user name and password admin.

- In the Welcome to **Cloudera on premises** page, click Change Password to change the Local Administrator Account password.
- Set up external authentication using the URL of the LDAP server and a CA certificate of your secure LDAP. Follow the instructions on the Welcome to **Cloudera on premises** page to complete this step.
- Click Test Connection to ensure that you are able to connect to the configured LDAP server.
- Register a Cloudera on premises environment
- Create your first Virtual Warehouse in the Cloudera Data Warehouse Data Services
- Provision an Cloudera AI Workbench in the Cloudera AI Data Services

## Installing in air gap environment

You can launch the Cloudera Data Services on premises installation wizard from Cloudera Manager and follow the steps to install Cloudera Data Services on premises in an air gap environment where your Cloudera Manager instance or your Kubernetes cluster does not have access to the Internet.

### Before you begin

- Ensure that your Kubernetes kubeconfig has permissions to create Kubernetes namespaces.
- You require persistent storage classes defined in your OpenShift cluster. Storage classes can be defined by OpenShift cluster administrators.
- Only TLS-enabled custom Docker Registry is supported. Ensure that you use a TLS certificate to secure the custom Docker Registry. The TLS certificate can be self-signed, or signed by a private or public trusted Certificate Authority (CA).
- Only TLS 1.2 is supported for authentication with Active Directory/LDAP. You require TLS 1.2 to authenticate the Cloudera Control Plane with your LDAP directory service like Active Directory.
- OCP network configurations that restrict pod communication are not supported. For example, multi-tenancy isolation with network policy is not supported.

### About this task

If this Cloudera Manager instance or your Kubernetes cluster does not have connectivity to https://archive.cloudera.com/p/cdp-pvc-ds/, you must mirror the Cloudera archive URL using a local HTTP server.

### Procedure

1. In Cloudera Manager, on the top right corner, click Add > Add Cluster. The Select Cluster Type page appears.

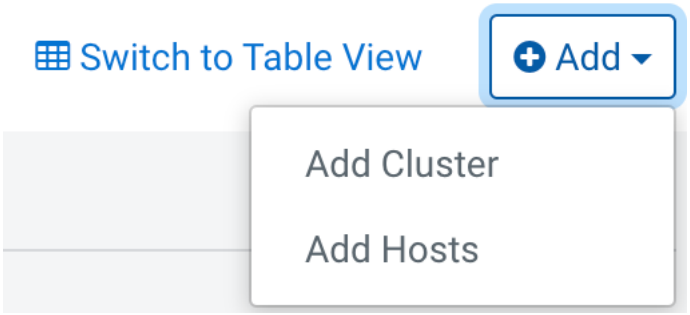

2. On the Select Cluster Type page, select the cluster type as on premises Containerized Cluster. Under Other Options, click here to install Cloudera Data Services on premises. Click Continue.

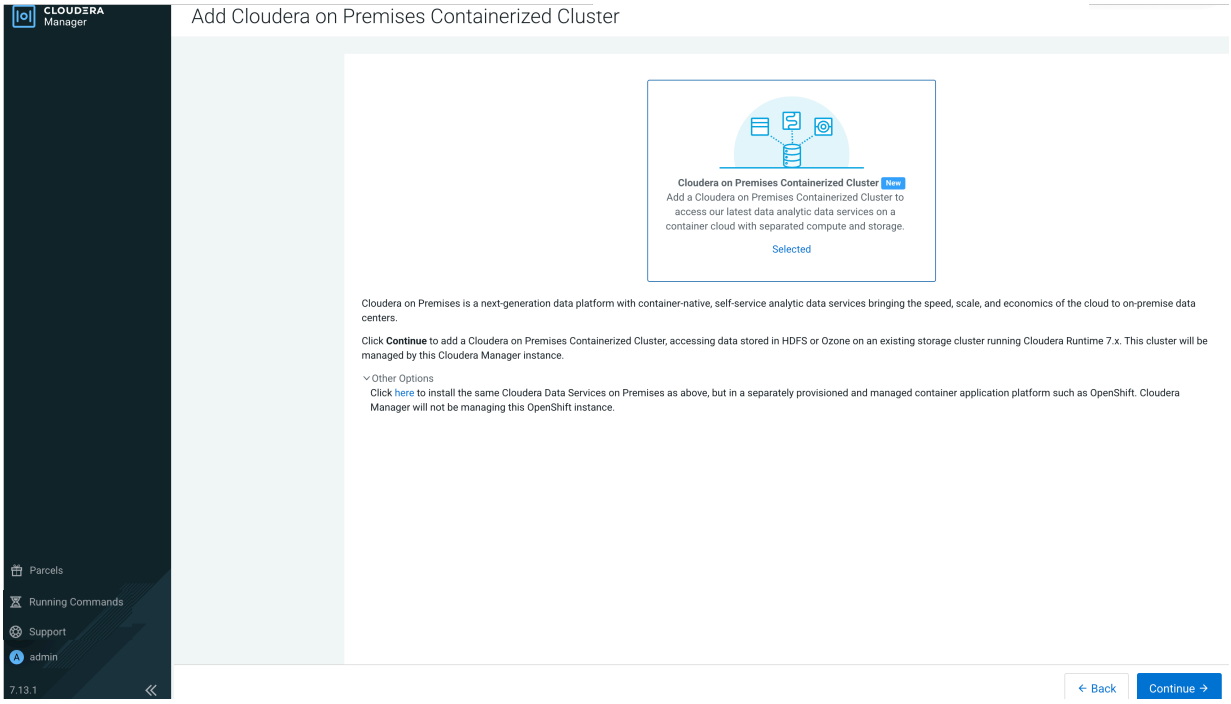


3. On the Getting Started page of the installation wizard, select Air Gapped as the Install Method. When you select the Air Gapped install option, extra steps are displayed. Follow these steps to download the Cloudera Data Services on premises deliverables .

   a. Offline download everything under the release you have been asked to download with the below command:
   LINK=https://<username>:<password>@archive.cloudera.com/p/cdp-pvc-ds/ && echo "$LINK$(curl -I "$LINK/latest/" 2>/dev/null | grep location | awk '{print $2}' | awk -F'/' '{print $4}')/" |xargs wget -l 0 --recursive --no-parent -e robots=off -nH --cut-dirs=2 --reject="index.html*" -t 10

   In the above command, replace `username`, `password` as provided by Cloudera.

   **Note:** If you want to download a specific version, locate the release name and version from https://archive.cloudera.com/p/cdp-pvc-ds/ and replace `'$LINK/latest'` to `'$LINK/<release name and version>'` in the above command. For example, to download 1.5.5-h1000, replace `'$LINK/latest'` with `'$LINK/1.5.5-h1000'` in the above command.

   **Note:** The deliverables for Cloudera Data Services on premises version 1.5.5 SP1 are very large (around 500 GB). Ensure you have enough free storage space before you begin the download. After downloading all the required files, you can reclaim the space by deleting the following items from the download folder:

   - Under 1.5.5-h1000/parcels (or equivalent release folder), retain only the files related to your OS (either RHEL8 or RHEL9).

   b. Once you have all the required files, create a new folder on your local http server (parallel to the folders of the Cloudera Manager and the CDH releases) as follows (example given for the 1.5.5-h1000 release):

```
sudo mkdir -p
/var/www/html/cloudera-repos/cdp–pvc-ds/1.5.5-h1000
```

   c. If required, pass the files to the customer for approval, and once approved, you can move them in the target folder created above (or into the customer organizational repository in case such is used). Once the files were moved into the folder, ensure they have the same permissions as the cm7 files by running (example given for the 1.5.5-h1000 release):

```
sudo chmod -R 755
```

```
/var/www/html/cloudera-repos/cdp-pvc-ds/1.5.5-h1000
```

**d.** Visit the Repository URL: http:///cloudera-repos/ in your browser and verify the files you downloaded are present. If you do not see anything, your web server may have been configured to not show indexes.

**e.** Click the Select Repository drop-down and select http://your_local_repo/cdp-pvc-ds/1.5.5-h1000 (example given for the 1.5.5-h1000 release)

**f.** Click Next.

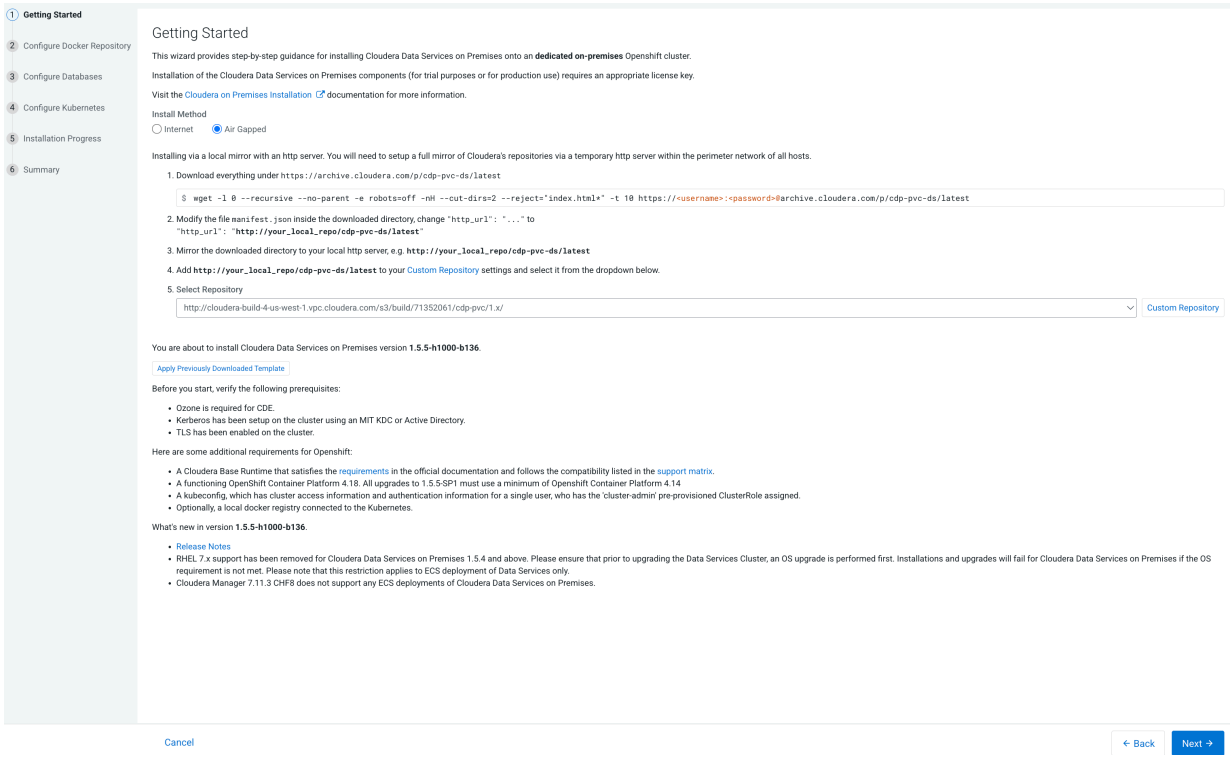


**Note:** You can also apply a template that you may have downloaded during a previous installation. The template contains all the installation configurations. Click Apply Previously Download Template to browse and upload a template stored on your machine.

4. On the Configure Docker Repository page, you must select one of the Docker repository options.

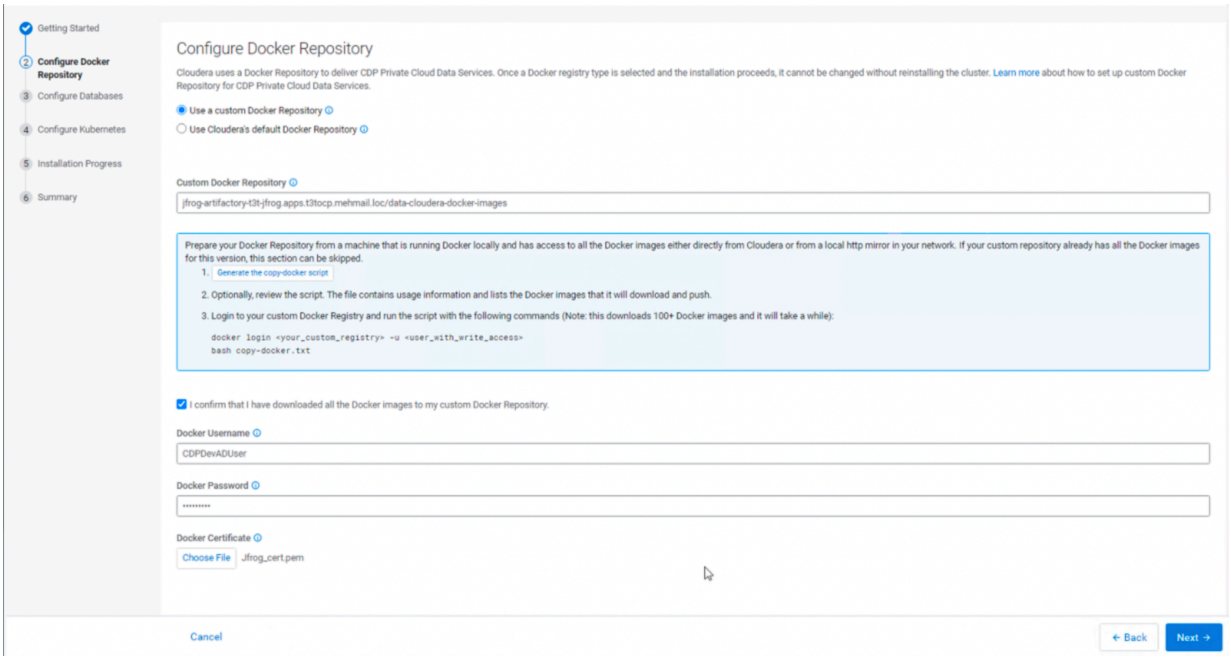


- Use a custom Docker Repository - Selecting this option, copies all images (Internet or Air Gapped) to the embedded registry. If you select Use a custom Docker Repository option, then enter your local Docker Repository in the Custom Docker Repository field in the following format:*[*DOCKER REGISTRY*]/[*REPOSITORY NAME*]*.
- Use **Cloudera** default Docker Repository - Selecting this option, copies images from Internet to the embedded registry. This uses the default repository that is in manifest.json. This option can be selected only if you have selected Internet as the install method. Use **Cloudera** default Docker Repository if you are setting up Cloudera on premises in non-production environments.

Follow the steps below to prepare your Docker Repository from a machine that is running Docker locally and has access to all the Docker images either directly from Cloudera or a local HTTP mirror in your network:

a) Click Generate the copy-docker script on the wizard or download the script file.
b) The script file will be downloaded on your personal computer.
c) Edit the script file as follows:

   For example, locate the original push command string:

```
$PODMAN_OR_DOCKER push "$DOCKER_REGISTRY_DEST/$imagePathTag"
```

   and replace with the modified command, adding the "`--tls-verify=false`" flag:

```
$PODMAN_OR_DOCKER push --tls-verify=false "$DOCKER_REGISTRY_DEST/$imageP
athTag"
```

d) Upload this edited script file back to the server that has the connection to the artifactory.
e) Log in to your custom Docker Registry and run the script using the following commands (as root):

   Replace the usename and `<PASSWORD>` with your credentials entitled to write into the remote artifactory. Note to run this as root user.

```
root@server:/home/cdp # podman login https://jfrog-artifactory-
jfrog.example.com/artifactory/data-cloudera-docker-images -u CDPPrdAD
User -p <PASSWORD> -v --tls-verify=false
Used: /run/containers/0/auth.json
Login Succeeded!
```

```
root@server:/home/cdp
#execute the script using:
# bash /home/cdp/copy-docker.txt
```

**Note:** This command downloads 100+ Docker images and it will take some time to download. You can speed up the process by opening additional parallel CLI sessions to the same node from which you executed the above script and then run the same script again. The scripts identifies already picked sources and will parallely push the remaining images.

f) Enter your Docker user name and password.
g) Click Choose File to upload your Docker certificate.
h) Click Continue.

5. On the Configure Databases page, click Next.

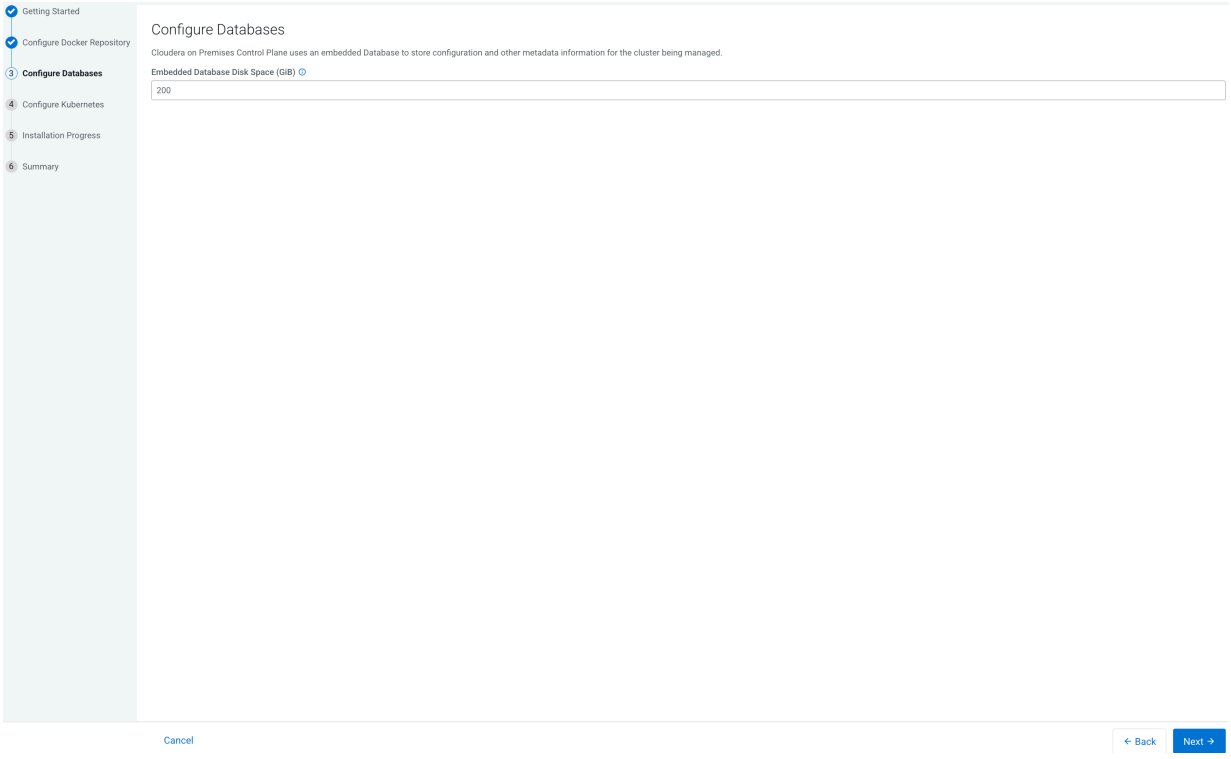

6. On the Configure Kubernetes page, enter your Kubernetes, Docker, database, and vault information.
   a) Upload a Kubernetes configuration (kubeconfig) file from your existing environment. You can obtain this file from your OpenShift Container Platform administrator. Ensure that this kubeconfig has permissions to create Kubernetes namespaces.
   b) In the Kubernetes Namespace field, enter the Kubernetes namespace that you want to use with this Cloudera on premises deployment. Kubernetes virtual clusters are called namespaces. For more information, see Kubernetes namespaces
   c) Enter your Vault information and upload a CA certificate. Cloudera recommends that you use an external Vault for production environments. Enter the Vault address and token, and upload a CA certificate.
   d) Enter a Block Storage Class to be configured on the Kubernetes cluster. Cloudera on premises uses Persistent Volumes to provision storage. You can leave this field empty if you have a default storage class configured on your Openshift cluster. Click Continue.

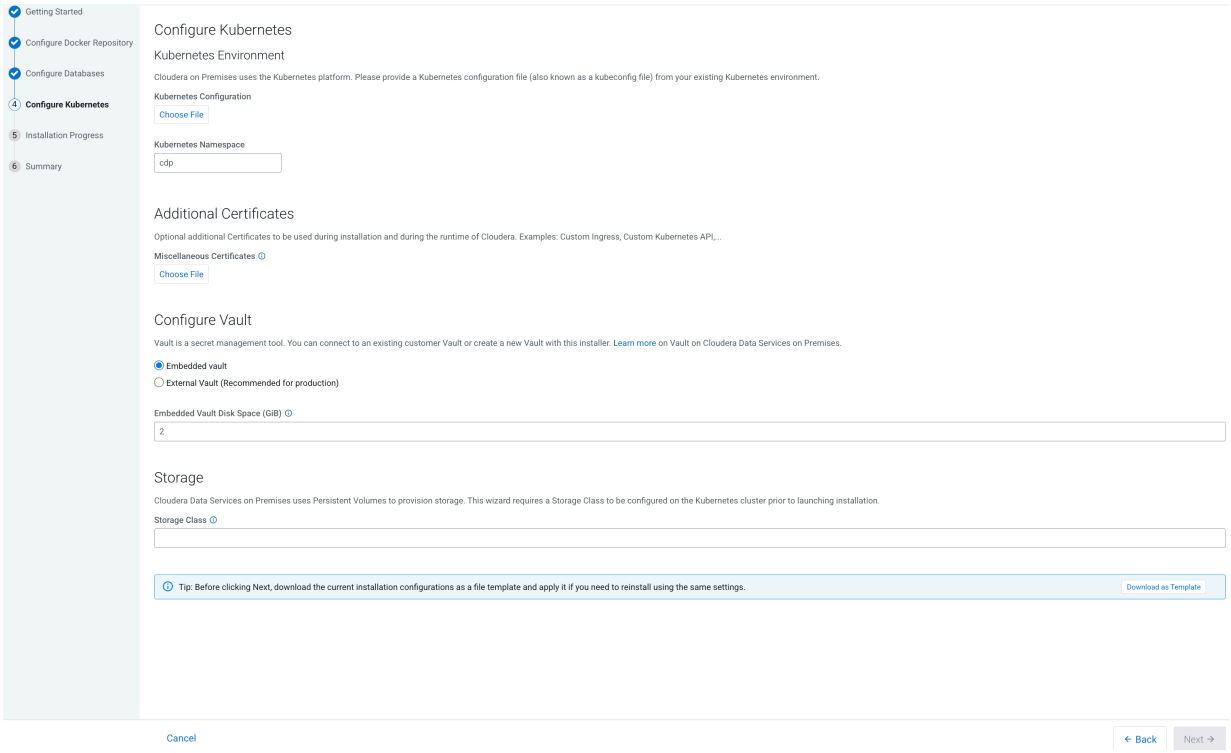


7. If you want to use this installation configuration again to install Cloudera on premises, you have the option to download this information as a template.

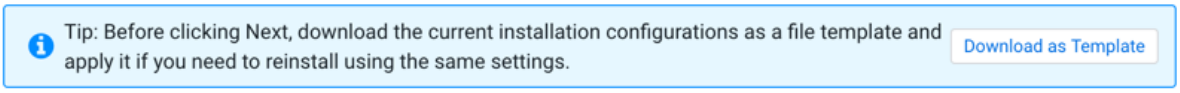


The template file is a text file that contains the database and vault information that you entered for this installation. This template is useful if you will be installing Cloudera on premises again with the same databases, as the template will populate the fields here automatically. Note that the user password information is not saved in the template.

8. The Installation Progress page appears. Click Continue.

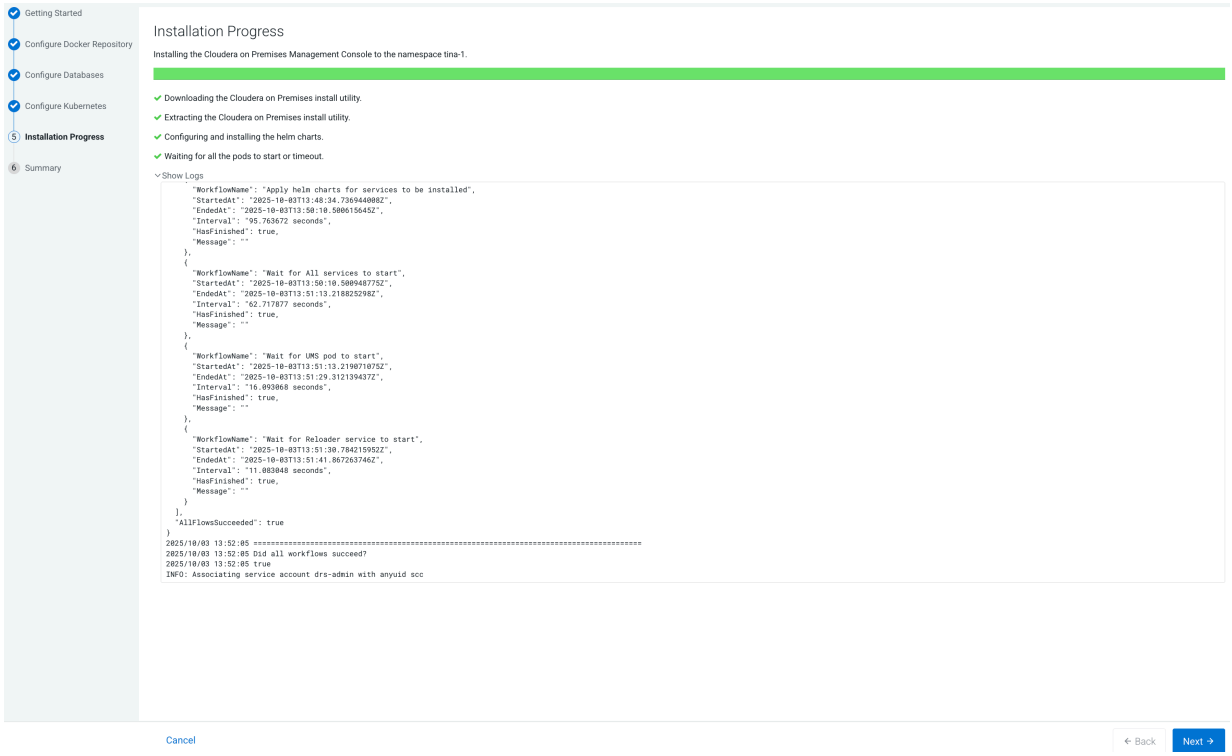


9. The summary message with a link to Launch **Cloudera on premises** appears.

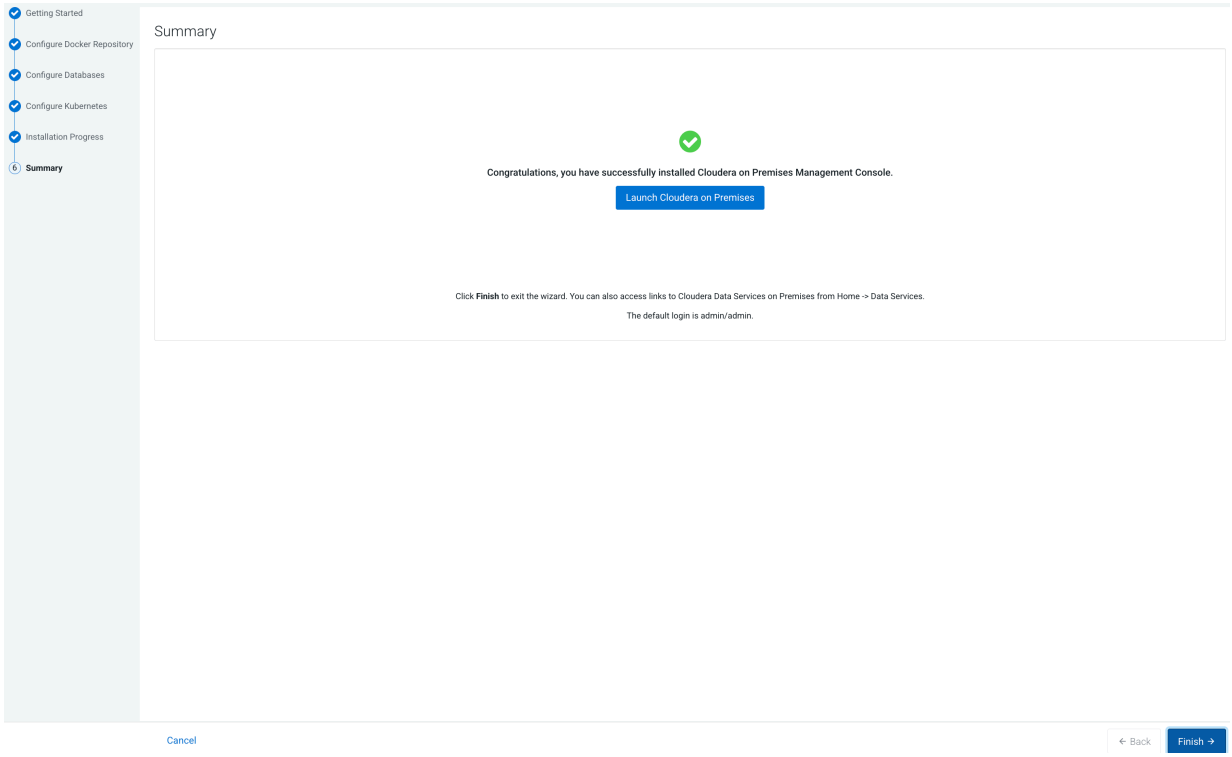


## What to do next

1. Click Launch **Cloudera on premises** to launch your Cloudera Data Services on premises.
2. Log in using the default user name and password admin/admin.

3. In the Welcome to **Cloudera on premises** page, click Change Password to change the Local Administrator Account password.
4. Set up external authentication using the URL of the LDAP server and a CA certificate of your secure LDAP. Follow the instructions on the Welcome to **Cloudera on premises** page to complete this step.
5. Click Test Connection to ensure that you can connect to the configured LDAP server.
6. Register a Cloudera Data Services on premises environment.
7. Create your first Virtual Warehouse in the Cloudera Data Warehouse Data Services and/or Provision an Cloudera AI Workbench in the Cloudera AI Data Services.

## Uninstalling Cloudera Data Services on premises

You can uninstall Cloudera Data Services on premises from your Cloudera Base on premises using the Cloudera Manager.

### Before you begin

Before you uninstall Cloudera Data Services on premises, ensure that you have deleted all the Cloudera on premises environments registered in your Cloudera Data Services on premises. You can delete your registered environments using Cloudera Management Console.

### Procedure

1. In Cloudera Manager, navigate to Cloudera Data Services on premises and click  . Click Uninstall.

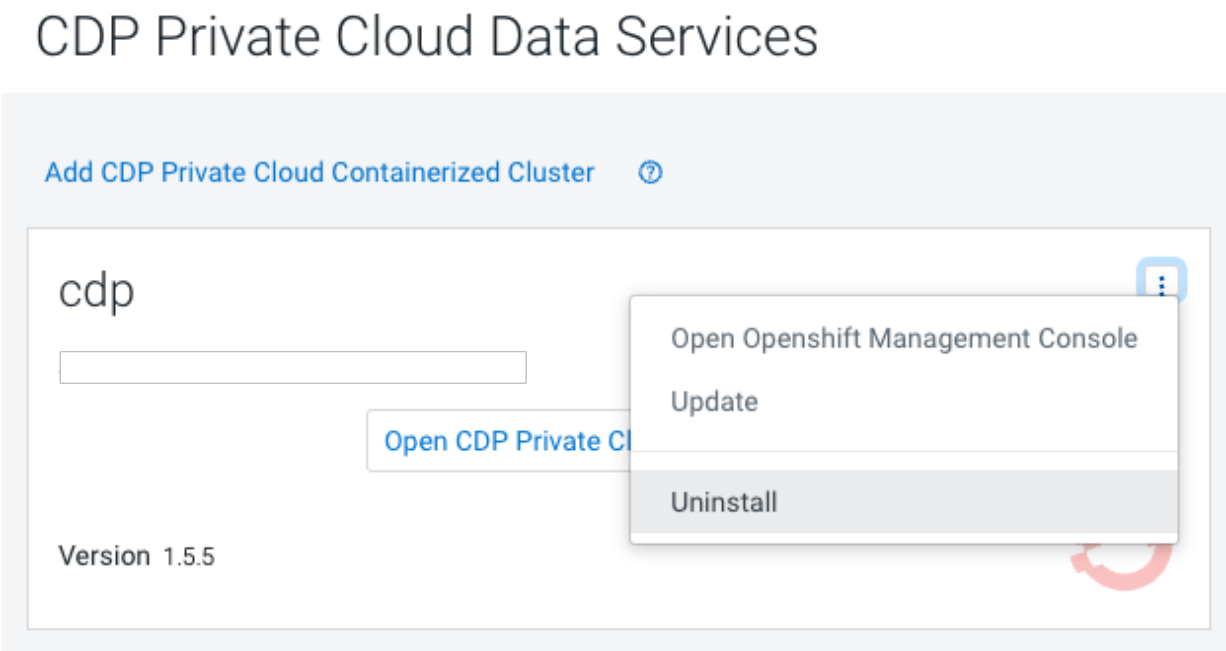


2. The Collect Information page appears. You must select the checkbox associated with your Cloudera on premises Environments. Click Choose File to upload your kubeconfig file associated with your Kubernetes cluster.

3. Select Keep shared artifacts if you have other Cloudera Data Services on premises instances running in your Kubernetes cluster, or select Delete shared artifacts to remove any cluster global security policies or objects associated with this Kubernetes namespace.

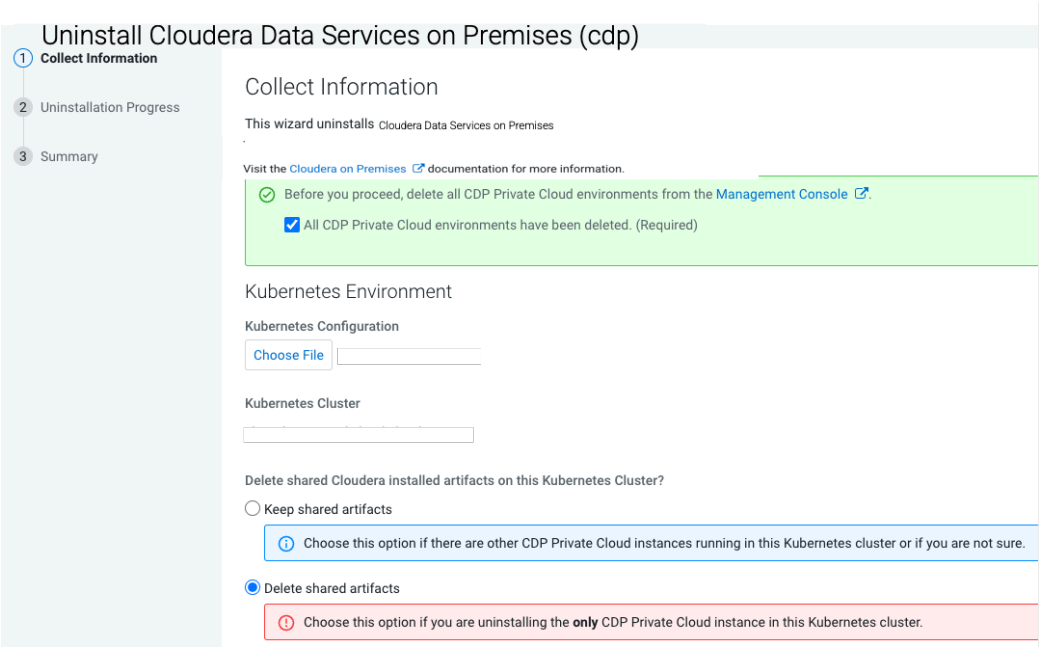

4. Click Continue to complete the process.

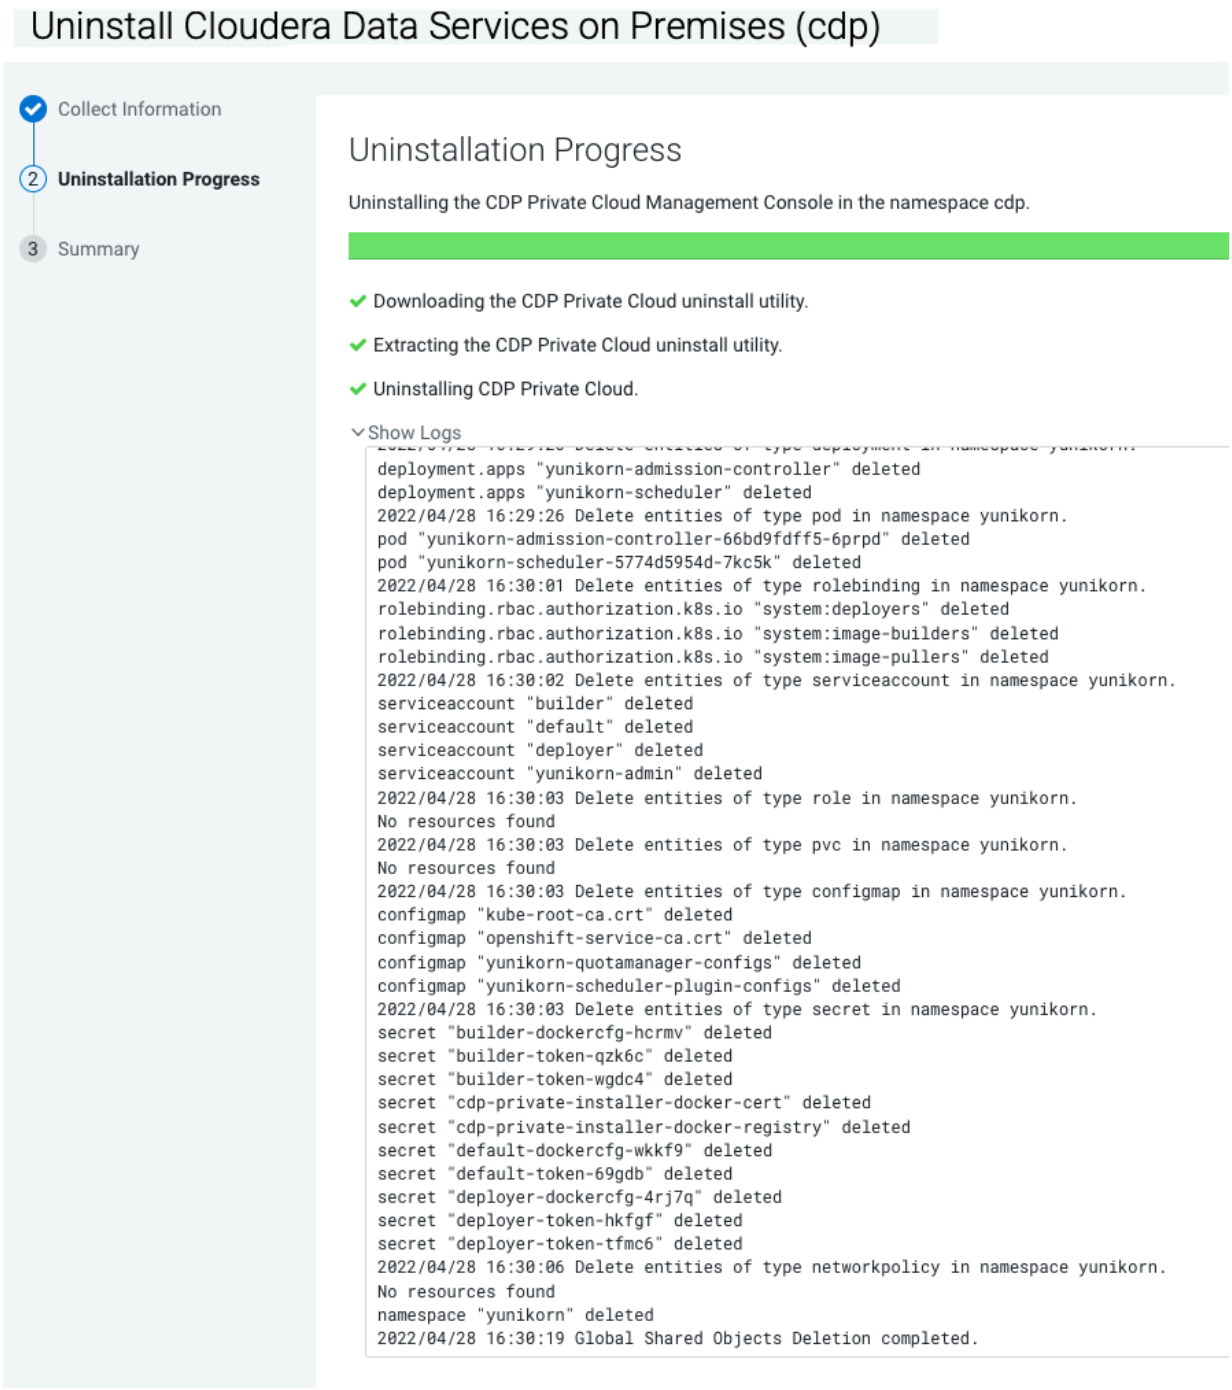


You will now see that Cloudera on premises has been uninstalled.


# Dedicating OCP nodes for specific workloads

You can use the kubectl taint command to dedicate OCP cluster nodes for specific workloads. You can dedicate GPU nodes for Cloudera AI Workbench, and NVME nodes for Cloudera Data Warehouse workloads.

## About this task

Run the following command to get a list of all of the cluster nodes:

```
kubectl get nodes
```

Run the following command to list information about a specific cluster node:

```
kubectl describe node <node_name>
```

In the returned output, look for the Taints field.

## Dedicate a GPU node for Cloudera AI Workbench

1. Run the following command to dedicate a GPU node for Cloudera AI Workbench:

   ```
   kubectl taint nodes <node_name> nvidia.com/gpu=true:NoSchedule
   ```

   No other workload pods will be allowed to run on the tainted node.
2. Run the following command to confirm that the taint has been successfully applied:

   ```
   kubectl describe node <node_name>
   ```

   In the returned output, look for the Taints field.

- To remove the taint, run the following command:

  ```
  kubectl taint nodes <node_name> nvidia.com/gpu=true:NoSchedule-
  ```

  This command returns:

  ```
  node/<node_name> untainted
  ```

## Dedicate a SSD node for Cloudera Data Warehouse workloads

1. Run the following command to dedicate a NVME node for Cloudera Data Warehouse workloads:

   ```
   kubectl taint nodes <node_name> ssd/nvme=true:NoSchedule
   ```

   No other workload pods will be allowed to run on the tainted node.

2. Run the following command:
   - To confirm that the taint has been successfully applied:

     ```
     kubectl describe node <node_name>
     ```

     In the returned output, look for the Taints field.
   - To remove the taint, run the following command:

     ```
     kubectl taint nodes <node_name> ssd/nvme=true:NoSchedule-
     ```

     This command returns:

     ```
     node/<node_name> untainted
     ```
3. Run the following command:
   - To label the dedicated NVME node for Cloudera Data Warehouse workloads:

     ```
     kubectl label node <node name> cdw.cloudera.com/dedicated-executor-host=
     <any value>
     ```
   - To unlabel the dedicated NVME node for Cloudera Data Warehouse workloads:

     ```
     kubectl label node <node name> cdw.cloudera.com/dedicated-executor-host-
     ```

### Additional Notes

**Note:**

To taint the node of an existing cluster which already has Cloudera AI and Cloudera Data Warehouse workspaces running, you must also run the following commands:

```
kubectl drain <node_name> --ignore-daemonsets --delete-emptydir-data --
timeout=600s
kubectl uncordon <node_name>
```

**Note:**

You cannot apply a taint to a master node, or to a single-node cluster.

### Related Information

Taints and Tolerations

## Configuring GPU node labeling on OCP

You can use NVIDIA Feature Discovery to generate labels for the set of GPUs available on OCP nodes. You can use these node labels to assign workloads to specific GPU devices.

### Configuring GPU node labeling on OCP nodes

1. Review the prerequisites listed on the NVIDIA GPU feature discovery page.
2. Use the instructions under Deployment via helm to deploy the GPU node labeling feature.
3. Information about using GPU node labeling is also available on the NVIDIA GPU feature discovery page.

### Known Issues and Limitations

- GPU node labeling is only supported for GPU cards manufactured by NVIDIA.