

Starting Cloudera Data Visualization in Cloudera Base on premises

Date published: 2020-10-30

Date modified: 2025-09-30



Legal Notice

© Cloudera Inc. 2025. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Support matrix.....	4
Hardware requirements.....	4
Installing Cloudera Data Visualization in Cloudera Base on premises.....	4
Adding Cloudera Data Visualization parcel repository to Cloudera Manager.....	5
Downloading and adding Cloudera Data Visualization CSD files.....	8
Adding Cloudera Data Visualization as a service.....	9
Accessing the Cloudera Data Visualization UI.....	11
Specifying a custom Python interpreter.....	12
Enabling Knox SSO for Cloudera Data Visualization.....	12
Installing additional dependencies for Cloudera Data Visualization in Cloudera Base on premises.....	16
Installing Chromium or Google Chrome for visual export and dashboard previews in Cloudera Data Visualization.....	16
Installing Oracle Client libraries for Cloudera Data Visualization.....	17

Support matrix

Before installation, ensure that your environment meets the following requirements.

Cloudera Data Visualization 8.x supports the following platforms and components:

Cloudera Manager	7.13.1
Cloudera Base on premises	7.3.1, 7.1.9, 7.1.9 SP1
Python	3.11 (bundled with parcel)
Operating systems	RHEL 9.5, 9.4, 9.2, 9.1, 8.10, 8.9, 8.8 Ubuntu 20.4

Cloudera Data Visualization 8.x supports the following databases:

OracleDB	19c
PostgreSQL	16, 15, 14, 13
MySQL	8.0

Related Information

[Specifying a custom Python interpreter](#)

Hardware requirements

Use the following guidelines to determine the minimum memory, CPU, and storage requirements needed to run Cloudera Data Visualization in Cloudera Base on premises.

Service	Memory	CPU	Disk
Cloudera Data Visualization	Minimum: 8 GB	Minimum: 1 Core (for Django)	Minimum: 32 GB total – 16 GB for the application database and 16 GB for the /tmp directory

The required disk space may increase with cluster size and user workloads. Cluster size refers to the number of nodes in the cluster and workload means the number of queries run and the number of concurrent unique users accessing the application in a given time period.

Cloudera Data Visualization relies heavily on its database, which can grow rapidly because of retained query history. A minimum of 16 GB disk space is required for the database, but additional space may be necessary depending on usage. To maintain performance and manage storage effectively, regularly clean up unused or outdated queries and documents.

Installing Cloudera Data Visualization in Cloudera Base on premises

Learn how to deploy Cloudera Data Visualization in Cloudera Base on premises.

Cloudera Data Visualization, when deployed in Cloudera Base on premises, allows you to explore datasets and build interactive visualizations directly within your on-premises Cloudera clusters. One key benefit of this deployment option is that it does not require a container runtime, and can be used without a Kubernetes environment.

This guide provides step-by-step instructions for installing Cloudera Data Visualization on a Cloudera Base on premises cluster.

Limitations

- **OS compatibility:**
Ubuntu 20, as well as RHEL 8 and 9, are supported. For the detailed list of supported OS versions, see the [Support Matrix](#).
- **Data source connections:**
Data connection auto-discovery is not supported. You need to manually set up and configure connections to the data sources.

Prerequisites

- Cloudera Manager and Cloudera Base are installed and running.
 - For installation instructions, see:
 - [Installing Cloudera Manager](#)
 - [Installing Cloudera Runtime](#)
 - For supported versions, see the [Support Matrix](#).
- You have sudo or root access to the relevant cluster nodes.
- You have a MySQL or other relational database configured for metadata storage.



Important:

Install Cloudera Data Visualization on a dedicated host to avoid resource contention with other services. Do not deploy it on nodes already running resource-intensive roles. The Cloudera Manager installation wizard evaluates host hardware to recommend role placements. Ensure that hardware is properly configured before running the wizard to allow appropriate role assignments.

Adding Cloudera Data Visualization parcel repository to Cloudera Manager

Before installing Cloudera Data Visualization, you must set up two components: a Cloudera Data Visualization parcel and a Custom Service Descriptor (CSD).

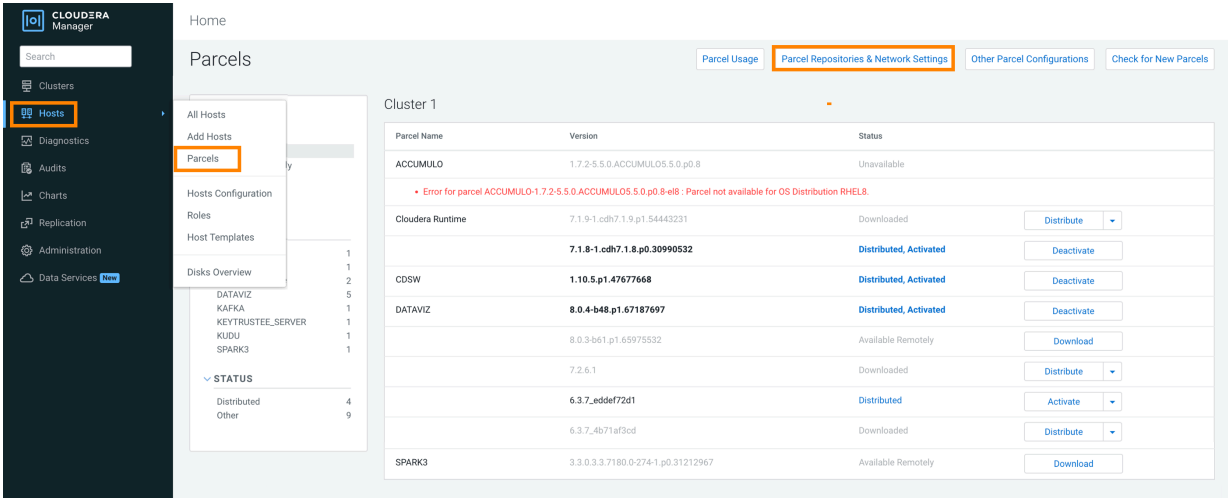
About this task

This guide explains how to add the Cloudera Data Visualization parcel to Cloudera Manager. It involves updating the parcel repository URL, then downloading, distributing, and activating the parcel to make it available for installation.

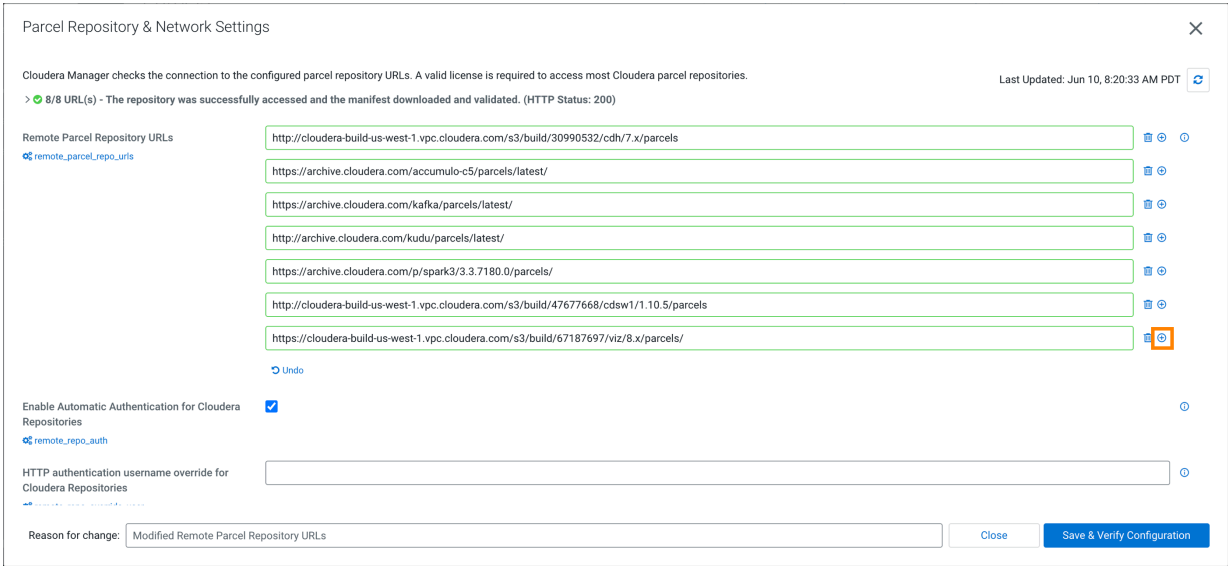
Procedure

1. Log in to Cloudera Manager.
2. From the left navigation bar, go to Hosts Parcels .

3. Select Parcel Repositories & Network Settings on the Parcels page.



4. On the Parcel Repository & Network Settings page, click + to add a new row in the Remote Parcel Repository URLs list.



5. Enter the URL of the Cloudera Data Visualization parcel repository.

Parcel Repository & Network Settings

Cloudera Manager checks the connection to the configured parcel repository URLs. A valid license is required to access most Cloudera parcel repositories. Last Updated: Jun 10, 8:21:47 AM PDT

8/8 URL(s) - The repository was successfully accessed and the manifest downloaded and validated. (HTTP Status: 200)

Remote Parcel Repository URLs

- [remote_parcel_repo_urls](#)
-
-
-
-
-
-
-
-

Enable Automatic Authentication for Cloudera Repositories ☒

[remote_repo_auth](#)

HTTP authentication username and password for:

Close Save & Verify Configuration



Note: Do not include the specific parcel file name in the URL, only provide the directory path. Cloudera Manager will scan the directory and display all available parcels. For example: `https://archive.cloudera.com/p/cdv/8.0.3/parcels/`.

- Click Save & Verify Configuration.
- Click Close to return to the Parcels page.
- On the Parcels page, locate the Cloudera Data Visualization parcel in the list.

Home

Parcels

Parcel Usage Parcel Repositories & Network Settings Other Parcel Configurations Check for New Parcels

Cluster 1

Parcel Name	Version	Status
ACCUMULO	1.7.2-5.5.0.ACCUMULOS.5.0.p0.8	Unavailable
Error for parcel ACCUMULO-1.7.2-5.5.0.ACCUMULOS.5.0.p0.8-el8: Parcel not available for OS Distribution RHEL8.		
Cloudera Runtime	7.1.9-1.cdh7.1.9.p1.54443231	Downloaded
	7.1.8-1.cdh7.1.8.p0.30990532	Distributed, Activated
CDSW	1.10.5.p1.47677668	Distributed, Activated
DATAVIZ	8.0.3-b61.p1.65975532	Available Remotely
SPARK3	3.3.0.3.3.7180.0-274-1.p0.31212967	Available Remotely

Location: Cluster 1 Available Remotely

Filters:

PARCEL NAME

- ACCUMULO 1
- CDSW 1
- Cloudera Runtime 2
- KAFKA 1
- KEYTRUSTEE_SERVER 1
- KUDU 1
- SPARK3 1

STATUS

- Distributed 4
- Other 9

- Click Download.
- Once the download is complete, click Distribute.
- After distribution, click Activate.

When prompted, choose whether to restart the service or activate without restarting. Select your preferred option and click OK.

Results

After activation, the parcel becomes available for use when adding the Cloudera Data Visualization service.

What to do next

Once the parcel is activated, you must add the Cloudera Data Visualization CSD files to Cloudera Manager. For instructions, see *Downloading Cloudera Data Visualization CSD files*.

Downloading and adding Cloudera Data Visualization CSD files

A Custom Service Descriptor (CSD) file contains the configuration details for Cloudera Manager to recognize and manage a new service. Follow the steps below to download the required CSD file, apply the correct permissions, and restart Cloudera Manager.

Before you begin

- You must have root or sudo privileges on the Cloudera Manager host.

Procedure

1. Open a terminal and connect to the host where Cloudera Manager is installed.

```
ssh [***USERNAME***]@[***CLOUDERA-MANAGER-HOST***]
```

- Replace [***USERNAME***] with a user that has appropriate permissions (for example: root or a sudo-enabled user).
- Replace [***CLOUDERA-MANAGER-HOST***] with the appropriate hostname or IP address of your Cloudera Manager server.

Example: `ssh root@cloudera-manager.example.com`

2. Once logged in, verify that the default CSD directory exists.

```
ls -lA /opt/cloudera/csd
```



Note: If the directory does not exist, create it before proceeding using `mkdir /opt/cloudera/csd`.

3. Navigate to the CSD directory.

```
cd /opt/cloudera/csd
```

4. Download the Cloudera Data Visualization CSD files to the default CSD directory on the Cloudera Manager host.



Note: The `/opt/cloudera/csd` directory requires root privileges. Use `sudo` to ensure the download succeeds when saving the file directly to this location.

```
sudo wget https://[***YOUR-DOWNLOAD-LOCATION***]/[***YOUR-CSD-FILE***].jar
```

- Replace [***YOUR-DOWNLOAD-LOCATION***] with the download URL provided by your administrator or Cloudera.



Note: Accessing `https://archive.cloudera.com/p/` requires authentication. You can provide your Cloudera credentials using the `--user` and `--password` options.

```
sudo wget --user=<your-username> --password=<your-password> https://...
```

- Replace [***YOUR-CSD-FILE***].jar with the actual filename provided by your administrator or Cloudera.

For example: `https://archive.cloudera.com/p/cdv/8.0.3/redhat8/yum/DATAVIZ-8.0.3-b61.p1.65975532.jar`

5. List the directory contents to confirm that the file has been downloaded.

```
ls -lA
```

You should see the downloaded JAR file listed in the directory.

6. Restart the Cloudera Manager Server to load the new CSD by running the following command on the host where Cloudera Manager is installed.

```
sudo service cloudera-scm-server restart
```


7. Restart Cloudera Management Service in the UI.
 - a) Go to the Status tab.
 - b) Select Actions Restart to restart the Cloudera Management Service.

What to do next

After restarting, Cloudera Manager will detect the Cloudera Data Visualization service. You can now proceed to add and configure it in your cluster.

Adding Cloudera Data Visualization as a service

Use the Add Service wizard in Cloudera Manager to deploy and configure Cloudera Data Visualization on a Cloudera Base on premises cluster.

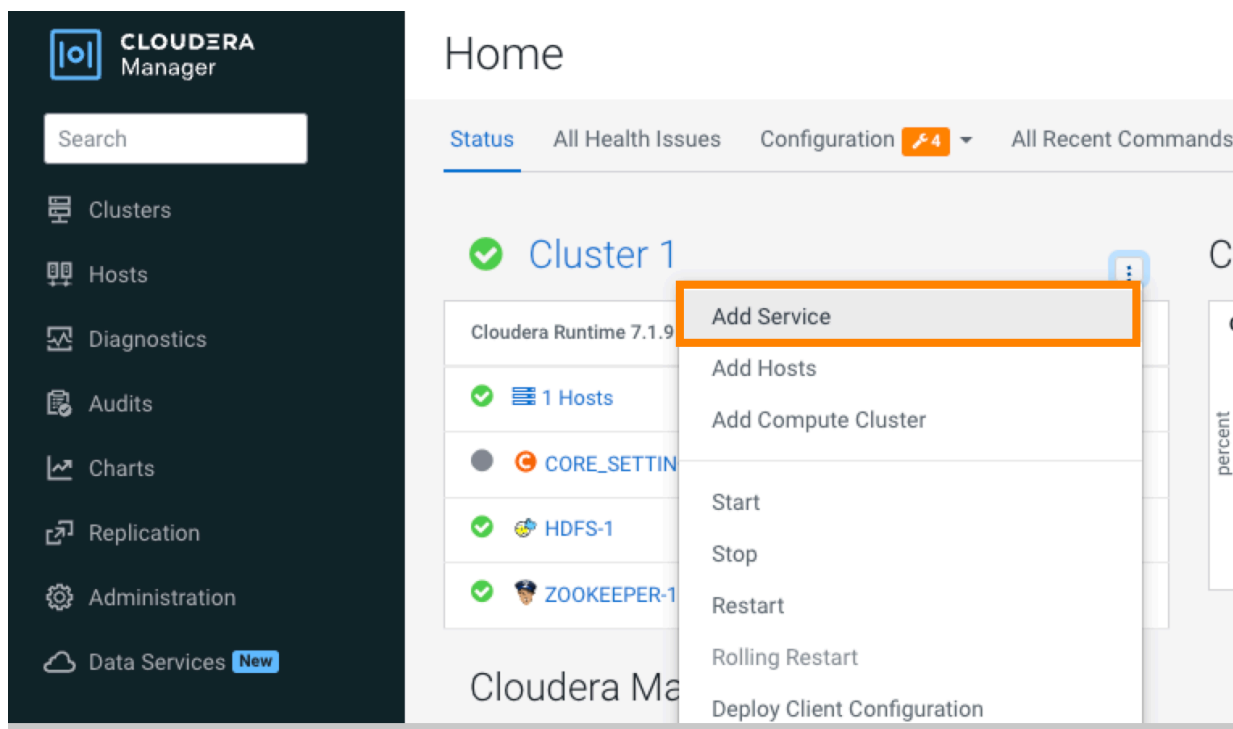
Before you begin

Ensure the following:

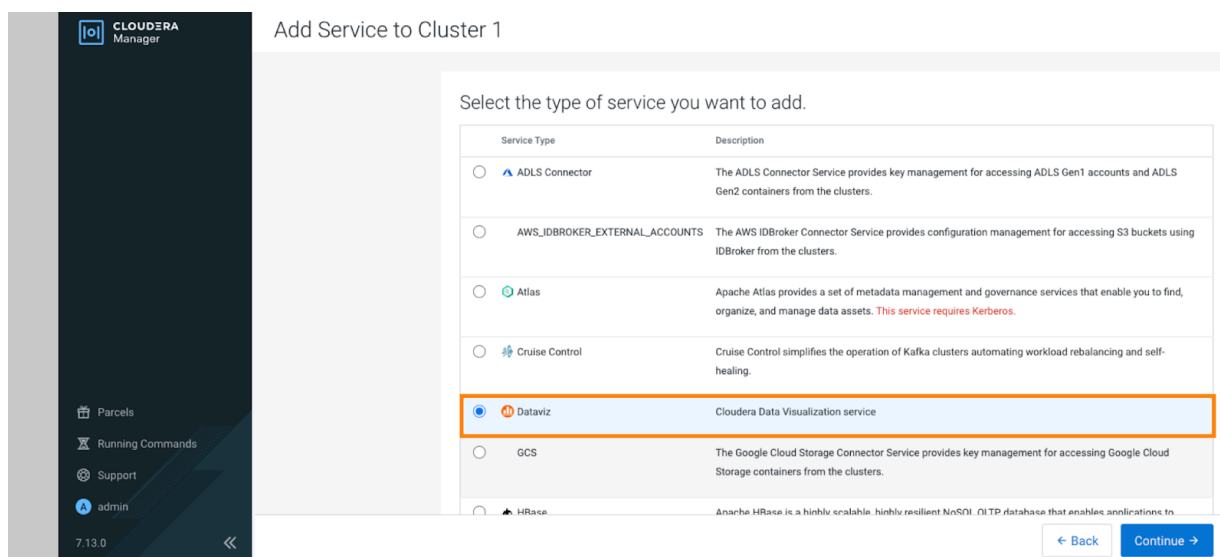
- A Cloudera Base on premises cluster is installed and running.
- You have added the Cloudera Data Visualization CSD files in the /opt/cloudera/csd folder, and Cloudera Manager has been restarted.
- You have added, downloaded, distributed, and activated the Cloudera Data Visualization parcel in Cloudera Manager.

Procedure

1. Log in to Cloudera Manager.
2. On the Home screen, open the drop-down menu next to your cluster name and select Add Service.



3. Select Dataviz from the list of available services, then click Continue.



The Add Service wizard is displayed.

4. In the Select Dependencies step, leave the default No Optional Dependencies setting and click Continue.
5. Assign roles to hosts.
 - On single-node clusters, roles are automatically assigned, so click Continue to proceed.
 - On multi-node clusters, manually assign roles to the desired hosts and then click Continue to proceed.
6. Configure a metadata database.
 - a) Configure the required database connection details:
 - Type: Select your database type from the dropdown. (For example: MySQL)
 - Database Hostname: Enter the hostname used when creating the database.
 - Database Name: Add the name of the Cloudera Data Visualization database.
 - Username: Add the Cloudera Data Visualization database username.
 - Password: Add the Cloudera Data Visualization database password.
 - b) You can click Test Connection to verify the database settings.
 - c) Click Continue.
7. Review the service configuration summary, then click Continue to begin the installation.
8. On the Command Details screen, monitor the installation progress.
9. Once the installation is complete, click Continue.
10. On the Summary screen, click Finish to complete the setup.

The Cloudera Data Visualization service is now added to your cluster. It will appear in the list of cluster services in Cloudera Manager.

11. Restart the Cloudera Management Service to update stale configurations with new monitoring entities.
12. Check the health status of the new service to verify it has started correctly.

If the Health Status is Good, the service has started successfully.

Results

You have completed your Cloudera Data Visualization installation on Cloudera Base on premises. For further information about adding services to Cloudera Runtime, see *Adding a service*.

What to do next

Open the Cloudera Data Visualization web interface to begin creating dashboards.

Related Information

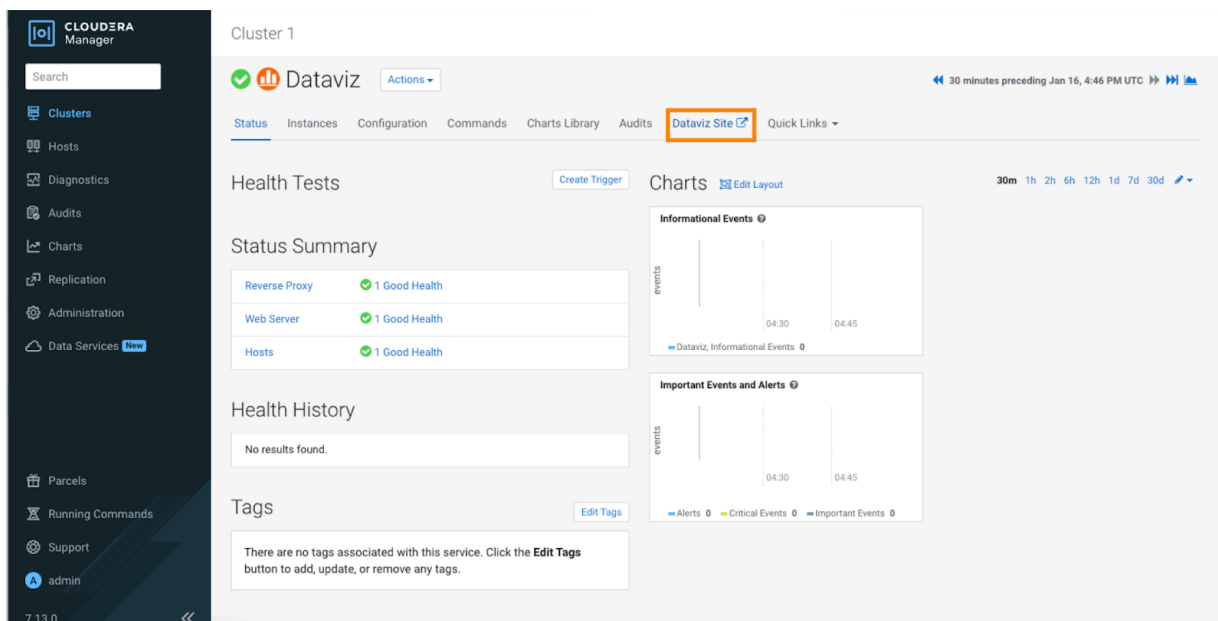
[Adding a service](#)

Accessing the Cloudera Data Visualization UI

After adding the service, follow these steps to access the Cloudera Data Visualization user interface.

Procedure

1. In Cloudera Manager, go to the Dataviz service details page.
2. Click Dataviz Site.



This opens the Cloudera Data Visualization interface in a new browser tab.

3. Use the default credentials to log in:

- Username: vizapps_admin
- Password: vizapps_admin



Note:

It is recommended to change the default password after first login to help maintain the account security. For more information, see *Changing user passwords*.

Results

You have now successfully added the Cloudera Data Visualization service and logged into the user interface. You can begin:

- Connecting to data sources
- Creating dashboards and visualizations

Related Information

[Changing user passwords](#)

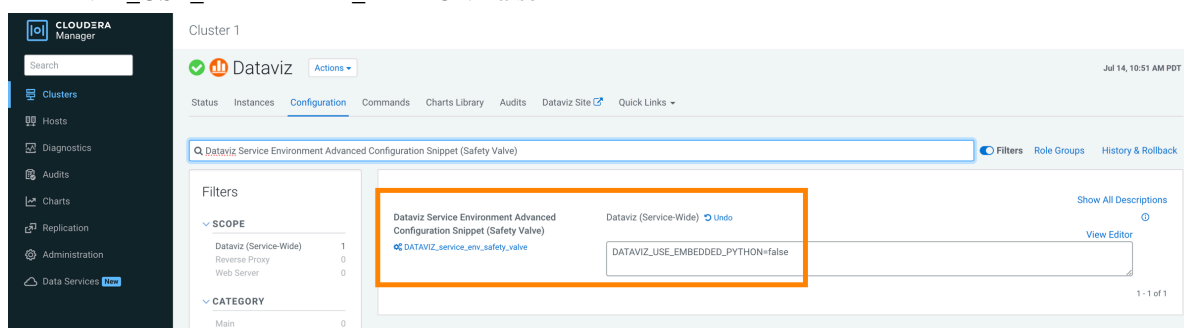
Specifying a custom Python interpreter

Cloudera Data Visualization includes an embedded Python interpreter by default. However, you can configure the system to use the OS-provided Python or a custom Python build instead. This provides greater flexibility and can improve performance when compatible Python versions are available on your system.

Procedure

1. To use a non-embedded Python interpreter:
 - a) In Cloudera Manager, go to Clusters and select the Dataviz service.
 - b) Click the Configuration tab.
 - c) Set the following environment variable in the Dataviz Service Environment Advanced Configuration Snippet (Safety Valve) configuration to disable the embedded Python.

`DATAVIZ_USE_EMBEDDED_PYTHON=false`



This enables Cloudera Data Visualization to use the system's default, OS-provided Python interpreter.

- d) Restart the Dataviz service to apply the changes.
2. To use a custom Python build:

If you want to use a specific Python installation (for example one you compiled from source or installed with `pyenv`), follow these additional steps:

 - a) Set the `PYTHONHOME` environment variable to point to the root directory of the Python installation you want to use.
 - b) Configure the library paths by making sure that the following directories are accessible:

- `python/lib`
- `sqlite/lib` (only if you compiled SQLite with this Python build)

You can do this by:

- Adding them to the `LD_LIBRARY_PATH` environment variable, **or**
- Registering them with `ldconfig` if you are using dynamically linked executables.

Enabling Knox SSO for Cloudera Data Visualization

You can configure Cloudera Data Visualization to use Knox Single Sign-On (SSO) if Knox is available in your cluster. The setup requires a Kerberized environment and includes configuring both Cloudera Data Visualization and Knox.

Before you begin

- The cluster must be Kerberized.

- Cloudera Data Visualization must be installed and Knox must be available on your cluster.
- Kerberos authentication must be enabled for Cloudera Data Visualization.

If not already configured, enable Kerberos authentication for Cloudera Data Visualization.

1. In Cloudera Manager, go to Clusters and select the Dataviz service.
2. Click the Configuration tab.
3. In the Search bar, type `kerberos` to filter the relevant settings.
4. Find the Enable Kerberos Authentication property.
5. Check the Dataviz (Service-Wide) box next to Enable Kerberos Authentication.
6. For Kerberos Principal, enter the Kerberos username (principal short name) the Cloudera Data Visualization service should use.
7. Click Save Changes.
8. Restart the Dataviz service for the changes to take effect.

The screenshot shows the Cloudera Manager interface for the Dataviz service configuration. The search bar at the top contains the text "Kerberos". On the left, a "Filters" sidebar shows the "SCOPE" filter with "Dataviz (Service-Wide)" selected, and the "CATEGORY" filter with "Main" selected. The main configuration area displays two properties:

- Enable Kerberos Authentication:** The property `KERBEROS_AUTH_ENABLE` (alias `kerberos.auth.enable`) is shown. The "Dataviz (Service-Wide)" checkbox is checked, and a description states "Indicates whether Kerberos is enabled."
- Kerberos Principal:** The property `kerberos Princ Name` (alias `kerberos.princ.name`) is shown. The "Dataviz (Service-Wide)" checkbox is checked, and the text input field contains the value "dataviz". A description states "Kerberos principal short name used by all roles of this service."

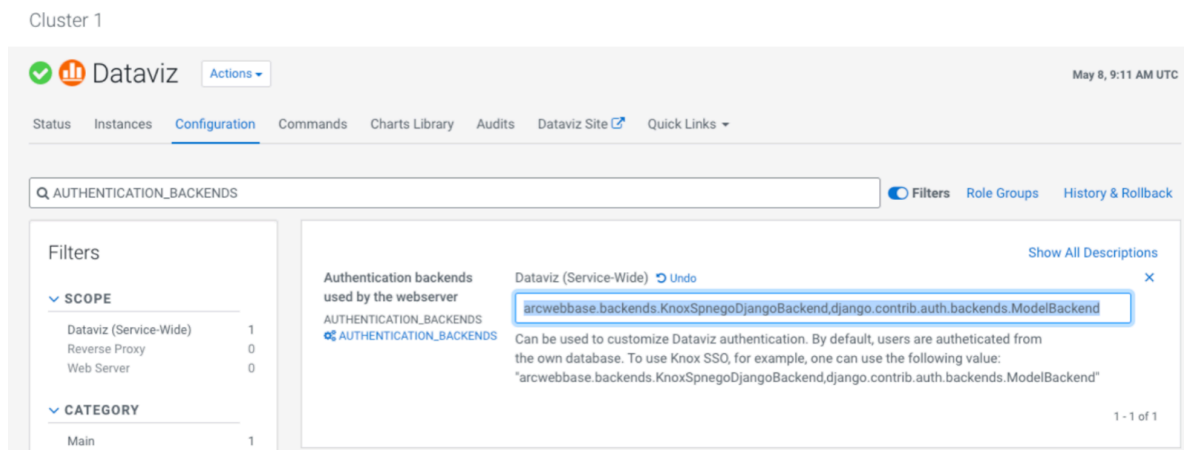
At the bottom right of the configuration area, it indicates "1 - 2 of 2" items.

Kerberos is now enabled for Cloudera Data Visualization, which is a prerequisite for enabling Knox SSO.

Procedure

1. Configure Cloudera Data Visualization for Knox SSO.

- In Cloudera Manager, go to Clusters and select the Dataviz service.
- Click the Configuration tab.
- Search for the Authentication backends used by the webserver property and set its value to:
`arcwebbase.backends.KnoxSpnegoDjangoBackend,django.contrib.auth.backends.ModelBackend`



Note:

This configuration enables Knox SSO while still allowing fallback login at `/arc/apps/login`. To enforce Knox authentication only, remove the fallback by setting `AUTHENTICATION_BACKENDS` to `arcwebbase.backends.KnoxSpnegoDjangoBackend`

- Restart the Dataviz service to apply the changes.
- ### 2. Make the Dataviz service definition available to Knox.
- Locate `DATAVIZ-KNOX-[***VERSION***].tgz` on the Cloudera archive site and download it to the node where the Knox service is running.

This file is co-located with the Cloudera Data Visualization CSD on the archive site and must be manually downloaded to each Knox host.

For example, to download `DATAVIZ-KNOX-8.0.4-b47.p1.67141340.tgz` to `/tmp`, run the following command:

```
cd /tmp;
wget https://archive.cloudera.com/p/cdv/8.0.4/redhat8/yum/DATAVIZ-KNOX-8.0.4-b47.p1.67141340.tgz
```

- SSH to the Knox node.

```
ssh [***KNOX-NODE***]
```

Replace `[***KNOX-NODE***]` with the hostname or IP address of your Knox node.

- Navigate to the temp directory and extract the archive to the Knox services directory.

```
cd /tmp
tar xzf /tmp/DATAVIZ-KNOX-[***VERSION***].tgz -C /opt/cloudera/parcels/CDH/lib/knox/data/services
```

Replace `[***VERSION***]` with the version number of the DATAVIZ-KNOX archive you downloaded.

- Verify that the Cloudera Data Visualization service was extracted correctly.

```
ls -la /opt/cloudera/parcels/CDH/lib/knox/data/services | grep dataviz
```

3. Update the Knox cdp-proxy topology.

- In Cloudera Manager, go to Clusters and select the Knox service.
- Click the Configuration tab.
- Search for the Knox Simplified Topology Management - cdp-proxy property.
- Edit the cdp-proxy topology by adding a new service entry for Cloudera Data Visualization.


```
DATAVIZ:url=[ ***DATAVIZ-SERVICE-URL*** ]
```

- In Cloudera Manager, locate Knox Simplified Topology Management .
- Save the topology changes.
- Restart the Knox service.

Results

Once configuration is complete, Cloudera Data Visualization will be accessible through Knox at: `https://[***KNOX-HOST**]:[***PORT**]/gateway/cdp-proxy/dataviz/`

Alternatively, you can also access Cloudera Data Visualization using the Knox Gateway home page.


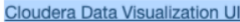



Welcome knoxui

– General Proxy Information


Knox Version	1.3.0.7.1.9.0-387 (hash=9883cc1a2c0846945cd042940f0eaf98ad07e1ea)
TLS Public Certificate	PEM JKS
Integration Tokens	Token Management Token Generation

– Topologies

–cdp-proxy ↗

Cloudera Manager Admin Console	Cloudera Data Visualization UI	HDFS Namenode UI	Impala UI
--------------------------------	--------------------------------	------------------	-----------



Impala UI
 +cdp-proxy-api ⚙️
 +cdp-proxy-token ⚙️
 +tokenexchange ⚙️

Access will be granted based on the identity authenticated by Knox.

Installing additional dependencies for Cloudera Data Visualization in Cloudera Base on premises

To enable certain features and integration, Cloudera Data Visualization in Cloudera Base on premises requires additional system dependencies to be installed manually on the hosts running Cloudera Data Visualization.

This section describes how to install and configure the following dependencies:

- Chromium – required for exporting visual content and displaying dashboard thumbnails.
- Oracle Client Libraries – required if you plan to use Oracle as the external metadata store.

Installing Chromium or Google Chrome for visual export and dashboard previews in Cloudera Data Visualization

Cloudera Data Visualization relies on Chromium to render and export visual outputs.

This enables:

- Downloading dashboards as PNG or PDF files
- Displaying thumbnail previews of dashboards

By default Cloudera Data Visualization expects the Chromium binary to be located at: `/usr/lib64/chromium-browser/headless_shell`

If Chromium is available at this location, no additional configuration is required. If not, you need to install Chromium on the hosts running Cloudera Data Visualization.

Installing Chromium on Red Hat–like systems

Use the following command to install Chromium on RHEL, CentOS, or similar systems: `sudo yum install chromium-headless`

This installs Chromium at the expected default location used by Cloudera Data Visualization.

Installing Google Chrome on Ubuntu

You can install Google Chrome from the official Google repository by following [Google's installation guide](#).

After installation, the binary is located at: `/usr/bin/google-chrome`

Using a custom Chromium path

If Chromium is installed at a different location, configure Cloudera Data Visualization to use that path by setting the `CHROMIUM_PATH` environment variable.

1. In Cloudera Manager, go to Clusters and select the Dataviz service.
2. Go to the Configuration tab.
3. Search for Dataviz Service Environment Advanced Configuration Snippet (Safety Valve).
4. Switch to View as Text and add the following line.

```
CHROMIUM_PATH=/path/to/your/custom/headless_shell
```

5. Save the changes and restart the Dataviz service for the changes to take effect.

Installing Oracle Client libraries for Cloudera Data Visualization

Cloudera Data Visualization supports the use of external databases as a metastore. When using Oracle as the metastore, additional setup steps are required beyond the ones needed for MySQL/MariaDB or PostgreSQL, as Cloudera Data Visualization does not include an Oracle database connector.

About this task

To enable connectivity to an Oracle metastore, you must install the Oracle client libraries on each host running the Cloudera Data Visualization webserver.

Procedure

1. Download the appropriate Oracle Instant Client libraries for your operating system from the [official Oracle website](#).
2. Follow the installation instructions provided on the [Installing Oracle Instant Client](#) page.