

Cloudera Runtime 7.3.2

Installing Metering V2 Manually

Date published: 2020-07-28

Date modified: 2026-03-31

CLOUdera

<https://docs.cloudera.com/>

Legal Notice

© Cloudera Inc. 2026. 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

Creating the Databus credentials.....	4
Installing the Metering V2 Service.....	5

Creating the Databus credentials

Databus credentials are needed to establish secure connections for the Metering V2 service to transmit metrics.

Before you begin

Ensure that you have the following:

- A Cloudera user account that can be promoted to Full Administrator role in Cloudera Manager.
- Grants to create API keys and Machine Users in Cloudera User Management Service (UMS).

Procedure

1. Create a new Machine User by following [Creating a machine user in Cloudera](#).



Note:

- The particular machine user name is not important but will be needed for the following commands.
- Make sure that the machine user is not deleted, to avoid losing metering events, which is registered by the Metering Events Monitor (MEM) service.

2. Retrieve the DatabaseUploader role CRN:

```
cdp iam list-roles | grep "altus:role:DbusUploader"
```

3. Update machine user with DatabaseUploader role:

```
cdp iam assign-machine-user-role --machine-user-name ***USER-NAME-FROM-STEP1*** --role ***CRN-FROM-STEP-2***
```

4. Create API Access Keys

```
cdp iam create-machine-user-access-key --machine-user-name ***USER-NAME-FROM-STEP1***
```

5. Record the accessKeyId and the privateKey for use later:

```
{
  "accessKey": {
    "accessKeyId": "abc",
    "crn": "",
    "actorCrn": "",
    "creationDate": "",
    "status": "ACTIVE",
    "type": "V2"
  },
  "privateKey": "abc"
}
```

6. In the CRN value from the JSON in the previous step, record the region and the account number. For example, the Region (us-west-1) and Tenant ID (215d3903-333d-45ae-8824-6a4a85e17e22) are highlighted in the example CRN:

```
crn:altus:iam:***REGION***:***TENANT-ID***:accesskey:***ACCESS-KEY***
```

```
crn:altus:iam:us-west-1:215d3903-333d-45ae-8824-6a4a85e17e22:accesskey:b940b1ef-3ef7-425f-9140-2709990f544c
```

Related Information

[Cloudera machine user](#)

Installing the Metering V2 Service

Learn how to install and configure the Metering V2 service to monetize API calls to Hive Metastore made by your clients.

Procedure

1. Open Cloudera Manager and open your cluster.

The screenshot displays the Cloudera Manager interface for a cluster. The top navigation bar includes 'Home', 'Status', 'All Health Issues', 'Configuration' (with a red '2' icon), and 'All Recent Commands'. A 'Switch to Table View' button is visible in the top right.

The main content area is divided into two columns. The left column shows the 'dc-datasharing' service status, which is highlighted with an orange box. It is part of 'Cloudera Runtime 7.3.2 (Parcels)' and has a green checkmark. Below it, a list of services is shown with their status and health indicators:

- 2 Hosts (2 red icons)
- Hive Metastore (1 orange icon)
- atlas (green icon)
- core_settings (green icon)
- hbase (green icon)
- hdfs (green icon)
- kafka (green icon)
- knox (green icon)
- ranger (1 orange icon)
- ranger-RANGER_RAZ (green icon)
- solr (green icon)
- zookeeper (green icon)

The right column displays four performance charts under the heading 'Charts'. The time range is set to '30m'.

- Cluster CPU:** A line chart showing 'dc-datasharing_Host CPU Usage Across Hosts' at 4.7%.
- Cluster Disk IO:** A bar chart showing 'Total Disk Bytes Read' at 0 and 'Total Disk Bytes Written' at 654K/s.
- Cluster Network IO:** A line chart showing 'Total Bytes Received' at 7.7K/s and 'Total Bytes Transmitted' at 10.7K/s.
- HDFS IO:** A bar chart showing 'Total Bytes Read' at 1b/s and 'Total Bytes Written' at 0.93b/s.

At the bottom, the 'Cloudera Management Service' section shows 'MGMT' with a green checkmark, 3 orange icons, and a power icon.

2. Click Action Add Service .

Home Switch to Table View

Status All Health Issues Configuration 2 All Recent Commands

✔ dc-datasharing

Cloudera Runtime 7.3.2 Add Service

- ✔ 2 Hosts Start
- ✔ Hive Metastore Stop
- ✔ atlas Restart
- Rolling Restart
- core_settings Deploy Client Configuration
- ✔ hbase Deploy Client Configuration and Refresh
- ✔ hdfs Refresh Cluster
- ✔ kafka Inspect Hosts in Cluster
- Backup Keys from Keytrustee Server
- ✔ knox View Client Configuration URLs
- ✔ ranger Enter Maintenance Mode
- ✔ ranger-RANGER View Maintenance Mode Status
- ✔ solr
- ✔ zookeeper

Cloudera Management Service

✔ MGMT 3 ⏻

Charts 30m 1h 2h 6h 12h 1d 7d 30d

Cluster CPU

Cluster Disk IO

Cluster Network IO

HDFS IO

3. Select Metering V2 in the resulting screen.

<input type="radio"/>	Kudu	Apache Kudu is a data store that enables real-time analytics on fast changing data.
<input type="radio"/>	Livy for Spark 3	Apache Livy for Spark 3 is a REST service used for deploying Spark3 applications.
<input checked="" type="radio"/>	 Metering V2	The Metering V2 service uploads metering data to Cloudera
<input type="radio"/>	Omid	Apache Omid (Optimistically transaction Management In Datastores) is a flexible, reliable, high performant and scalable transactional framework that allows Big Data applications to execute ACID transactions on top of MVCC key/value NoSQL datastores.
<input type="radio"/>	Oozie	Apache Oozie is a workflow coordination service to manage and schedule data processing jobs on your cluster.
<input type="radio"/>	Phoenix	Apache Phoenix is a scale-out relational database that supports OLTP workloads and provides secondary indexes, materialized views, star schema support, and common HBase optimizations. Phoenix uses Apache HBase as the underlying data store.
<input type="radio"/>	Profiler Manager	Profiler Manager is a service to configure or manage profilers. Serves aggregated statistics for assets and enables data discovery for AWS Glue and related services.

- In the **Assign Templates** step (the **Select Dependencies** step will already be done), click the Gateway text box and select the nodes with "master" in their name. Once the hosts have been selected, click the Continue button.



Note: The Metering V2 service should be colocated with the Cloudera Data Sharing service, which in turn is colocated with the HMS Metastores.

Figure 1: Role assignment for Light Duty Data Lakes

Add Metering V2 Service to dc-datasharing

- Select Dependencies
- Assign Templates**
- Review Changes
- Command Details
- Summary

Assign Templates

You can customize where the roles are assigned by assigning role types to host templates.

By choosing a specific host template for a role type, roles of this specific type will be assigned to all the hosts specified by that host template.

For example, if you have two host templates, one called master and one called worker. One host is using the master host template and five hosts are using the worker host template. Using these templates, you can specify which sets of hosts the new roles should be assigned to.

Note: you cannot place the roles onto individual hosts, because all the hosts using a particular template must contain identical roles. Cloudera does not recommend altering existing assignments unless you have specific requirements, such as having pre-selected a specific host for a specific role.

For more information about assigning roles using templates, see, [Assigning Roles using Host Templates](#).

You can also view the role assignments by host and host template. [View By Host Template](#)

Server × 1 New	Gateway × 1 New
<input type="text" value="master"/>	<input type="text" value="master"/>

[Cancel](#) [← Back](#) [Continue →](#)

Figure 2: Role assignment for Enterprise Duty Data Lakes

Add Metering V2 Service to XXXXXXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX

1 Select Dependencies

2 Assign Templates

3 Review Changes

4 Command Details

5 Summary

Assign Templates

You can customize where the roles are assigned by assigning role types to host templates.

By choosing a specific host template for a role type, roles of this specific type will be assigned to all the hosts specified by that host template.

For example, if you have two host templates, one called master and one called worker. One host is using the master host template and five hosts are using the worker host template. Using these templates, you can specify which sets of hosts the new roles should be assigned to.

Note: you cannot place the roles onto individual hosts, because all the hosts using a particular template must contain identical roles. Cloudera does not recommend altering existing assignments unless you have specific requirements, such as having pre-selected a specific host for a specific role.

For more information about assigning roles using templates, see, [Assigning Roles using Host Templates](#).

You can also view the role assignments by host and host template. [View By Host Template](#)

Server × 2 New Gateway × 2 New

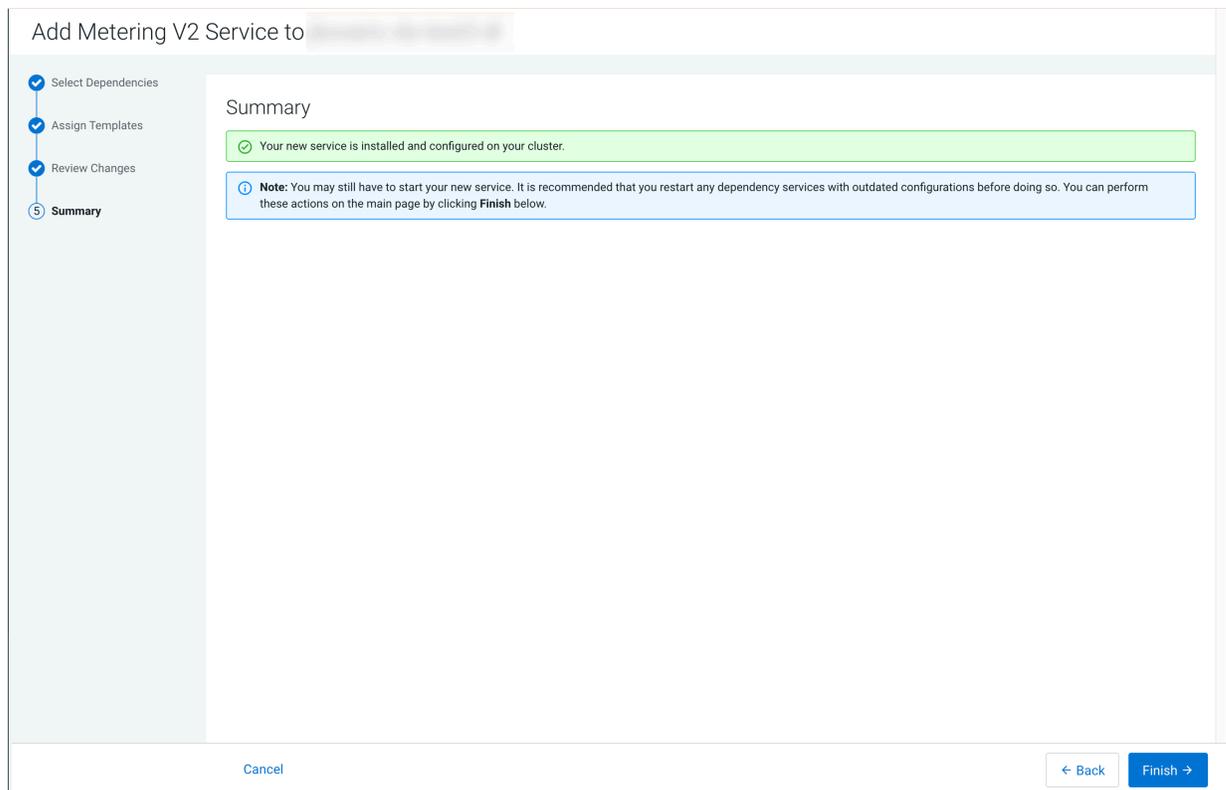
gateway gateway

[Cancel](#) [← Back](#) [Continue →](#)

5. In **Review Changes**, click **Continue**.

6. The **Command Details** step will automatically switch to **Summary** if no problems are detected with the first commands.

7. In **Summary**, click the **Finish**.



The screenshot shows a web interface for adding a Metering V2 Service. The page title is "Add Metering V2 Service to [redacted]". On the left, a vertical navigation pane shows four steps: "Select Dependencies", "Assign Templates", "Review Changes", and "Summary" (which is highlighted with a circled '5'). The main content area is titled "Summary" and contains a green success message: "Your new service is installed and configured on your cluster." Below this is a blue note box with an information icon: "Note: You may still have to start your new service. It is recommended that you restart any dependency services with outdated configurations before doing so. You can perform these actions on the main page by clicking **Finish** below." At the bottom of the page, there are three buttons: "Cancel", "← Back", and "Finish →".



Attention: The installation fails after clicking finish as you need to provide the DBUS information for correct operation after the first installation attempt. Continue with the next step.

8. Go to Clusters Metering V2 Configuration and fill in the following information:

The screenshot shows the Cloudera Metering V2 Configuration page. The page has a search bar at the top, a navigation menu with 'Status', 'Instances', 'Configuration', 'Commands', 'Charts Library', and 'Quick Links'. The 'Configuration' tab is active. On the left, there is a 'Filters' panel with sections for 'SCOPE', 'CATEGORY', and 'STATUS'. The 'SCOPE' section is expanded, showing 'Metering V2 (Service-Wide)' selected. The main configuration area displays several configuration items with their values and a 'Show All Descriptions' link. The 'Base data directory - DO NOT CHANGE' is set to '/var/metering'. 'Alternatives Priority' is set to '50'. 'DBUS Access Secret Key', 'DBUS Access Secret Key Algorithm', 'DBUS Access Key ID', and 'DBUS Host' are all marked as 'This field is required'.

- **DBUS Access Secret Key** - The value from the privateKey attribute when the DBUS machine user was created.
- **DBUS Access Secret Key Algorithm** - Set it to Ed25519.
- **DBUS Access Key ID** - The value from the accessKeyId attribute when the DBUS machine user was created.
- **DBUS Host** - dbusapi.us-west-1.sigma.altus.cloudera.com



Note: The DBUS Host will change depending on the region:

- For the EU region (eu-1) use api.eu-1.cdp.cloudera.com.
- For the AP region (ap-1) use api.ap-1.cdp.cloudera.com.

- **DBUS Stream** - manowar_prod-mow-prod-MeteringV2



Note: The DBUS Stream will change depending on the region:

- For the EU region (eu-1) use manowar_prod_euc1-mow-prod-eu-central-1-MeteringV2.
- For the AP region (ap-1) use manowar_prod_apse2-mow-prod-ap-southeast-2-MeteringV2.

- **DBUS App Name** - manowar_prod-mow-prod-MeteringV2



Note: The DBUS App Name will change depending on the region:

- For the EU region (eu-1) use manowar_prod_euc1-mow-prod-eu-central-1-MeteringV2.
- For the AP region (ap-1) use manowar_prod_apse2-mow-prod-ap-southeast-2-MeteringV2.

- **DBUS Partition Key** - Enter the Tenant ID that was recorded when the DBUS machine user was created.
- **Region** - Enter the Region that was recorded when the DBUS machine user was created.

9. Click Save Changes.

10. Click Clusters Metering V2 Actions Restart .