

Monitoring

Date published: 2020-08-17

Date modified: 2022-11-18



Legal Notice

© Cloudera Inc. 2024. All rights reserved.

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

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

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

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

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

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

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

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

Contents

Monitoring Data Warehouse service resources with Grafana dashboards.....	4
Connecting to Grafana dashboards in Private Cloud.....	5
Limitations of Grafana in Cloudera Data Warehouse Private Cloud.....	6

Monitoring Data Warehouse service resources with Grafana dashboards

Grafana is visualization and analytics software that enables the development of dashboards to monitor metrics data. You can access pre-built Grafana dashboards to monitor Virtual Warehouses and your compute cluster in Cloudera Data Warehouse (CDW).

You connect to prebuilt dashboards to view metrics of CDW operations. Cloudera provides prebuilt Grafana dashboards for Hive, Impala, and Hue dashboards of metrics data, charts, and other visuals.

Using Grafana, CDP metrics are centralized in a single spot, stored in the Prometheus database, and monitored by Prometheus. Your workload databases are not involved in any way. You can immediately view the following pre-built dashboards:

Hive dashboards

The Hive dashboards cover the following operations of the Hive SQL engine in CDW:


- Auto-scaling
- Hive metastore
- HiveServer
- The Hive service itself (Hive-Home)
- LLAP

Impala dashboards

The Impala dashboards include the following operations of the Impala SQL engine in CDW:

- Catalog server
- Coordinator
- Executor
- Statestore
- The Impala service itself

The following screenshot shows the available scratch and cache disk utilization graphs for the Impala Virtual Warehouse:



You can view dashboard metrics for different time periods by selecting the period of interest from the time range dropdown in the horizontal navigation.

On the Embedded Container Service (ECS) platform, you can view the CPU, memory, network usage, and disk input-output for each CDW node using the `[***ENVIRONMENT-NAME***]-Nodes` option. You can also expand the individual dashboards to see more details, as described in the following table:

Dashboard name	Description	Available metrics
CPU	CPU utilization per node	<ul style="list-style-type: none"> Usage per node Usage per user Usage per system Idle time IO wait
Memory	Memory utilization per node	<ul style="list-style-type: none"> Usage per node Buffer cache Page cache Total, used, and available
Network	Number of bytes and packets sent and received	<ul style="list-style-type: none"> Network transmitted Network received Network transmitted by an interface Network received by an interface
Disk	Disk bytes read and written	<ul style="list-style-type: none"> Bytes written Bytes read IO wait time



Note: Node-level metrics are currently available only for ECS environments.

Related Information

[Grafana documentation](#)

Connecting to Grafana dashboards in Private Cloud

You can access Grafana from the Management Console and view various dashboards related to the Data Warehouse service.


Before you begin

You must have an activated CDW environment to view dashboards for CDW service resources in Grafana.

Procedure

1. Log in to the Management Console as an Administrator.
2. Go to the **Dashboard** page and click Monitoring Dashboard.
- 3.



In the Grafana web interface, click  in the left navigation menu, and select Manage.

A list of available monitoring dashboards is displayed:

CDW service area	Dashboard topics
CDP Control Plane	Alerts generated by Management Console, pod status, including count, restarts, CPU usage, memory usage, and container memory and CPU usage
Data Warehouse compute auto-scaling	Auto-scaling
Hive	Hive MetaStore (HMS), HiveServer2, Hive service (Hive-Home), and several dashboards for LLAP
Impala	Impala components: catalog server, coordinator, executor, statestore, and the overall Impala service (Impala-Home)
Hue	overall Hue service (Hue-Home)

CDW service area	Dashboard topics
Overview	Kubernetes alerts, pod status, pod CPU usage, pod memory usage, app CPU usage, app memory usage, container memory usage, container CPU usage
Nodes	CPU, memory, network usage and disk IO metrics at the node level for a given environment

Limitations of Grafana in Cloudera Data Warehouse Private Cloud

Learn about the Grafana capabilities in Cloudera Data Warehouse (CDW) that Cloudera does not support. Grafana in CDW is intended for use by cluster operations professionals who are familiar with monitoring tools, interpreting metrics, and performing maintenance.

Supported features

- Viewing and organizing Grafana dashboards.

Unsupported features

- Storing metrics longer than 15 days.