

## Monitoring

Date published: 2024-01-01

Date modified: 2024-01-01

**CLOUDERA**

# 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.

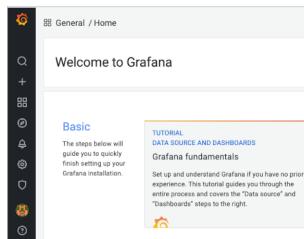
# Contents

<b>Monitoring resources with Grafana.....</b>	<b>4</b>
Grafana in CDW overview.....	4
Limitations of Grafana in CDW.....	6
Getting started in Grafana.....	7
Viewing prebuilt dashboards.....	8
Monitoring HMS.....	9
Monitoring key Hive metrics.....	11
Monitoring Impala executors.....	12
Monitoring Impala catalog.....	14
Monitoring Impala statestore.....	15
Monitoring Impala admission control.....	15
Monitoring Impala coordinators.....	16
Monitoring nodes.....	16
Monitoring Kubernetes Services from Grafana.....	19
Updating a dashboard graph.....	21
Inspecting dashboard data and queries.....	22
Creating a custom dashboard.....	23
Meeting prerequisites to set up alerts.....	24
Creating alerts.....	26
Reviewing alerts and notifications.....	27
<b>Monitoring Kubernetes resources from K8s dashboard.....</b>	<b>28</b>
Activating the K8S dashboard.....	28
Using the K8S dashboard.....	29
<b>Forwarding logs to your observability system.....</b>	<b>33</b>
Providing proxy CA certificates.....	36

# Monitoring resources with Grafana

Grafana is visualization and analytics software for using dashboards to monitor metrics data. You learn how to access pre-built Grafana dashboards to monitor Virtual Warehouses and your compute cluster in Cloudera Data Warehouse (CDW). As a Cluster Operations professional, you visualize Kubernetes and Istio metrics from Grafana to monitor and maintain the cluster.

When you log into Grafana, the Welcome page contains links to documentation about Grafana basics. From the left navigation, you can get started by searching for a particular dashboard or hovering over the navigation icons to familiarize yourself with the UI functions.



As a new user, to familiarize yourself with the pre-built dashboards that chart metrics about Hive, Impala, Hue, Druid, Kubernetes, and Istio in CDW, follow steps in this documentation.

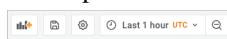
## Related Information

[Grafana documentation](#)

## Grafana in CDW overview

You connect to prebuilt dashboards to view metrics of CDW operations. Cloudera provides prebuilt Grafana dashboards for Hive, Impala, Hue, Druid, Kubernetes, and Istio 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 pre-built dashboards described below. You can view dashboard metrics for different time periods by selecting the period of interest from the dropdown in the horizontal navigation:



You can also create your own dashboards. "Get started with Grafana and Prometheus" describes how to create dashboards of CDP metrics. Describing all the details of how to use Grafana is beyond the scope of this documentation. Grafana described in this documentation is not the enterprise version.

## 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

For example:



## Impala dashboards

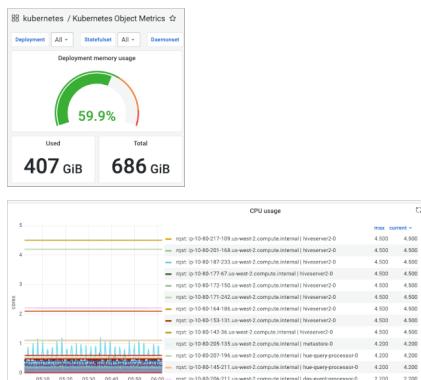
The Impala dashboards include the following operations of the Impala SQL engine in CDW:

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

## Kubernetes dashboards

You can get insight into the operations of your CDP clusters from Kubernetes dashboards. Kubernetes dashboards represent the following metrics:

- CoreDNS: requests and duration of responses
- App Metrics: number of Kubernetes pods, CPU usage relative to request or to limit, memory usage sliced and diced a number of ways
- Object Metrics: Deployment memory and CPU usage, in total, and by node.



## Istio dashboards

To work with Cloudera Support, you might use the Istio dashboards. Istio is an open platform that provides microservice security, connections, and monitoring. The Istio Mesh dashboards cover the following views of the service mesh network of microservices:

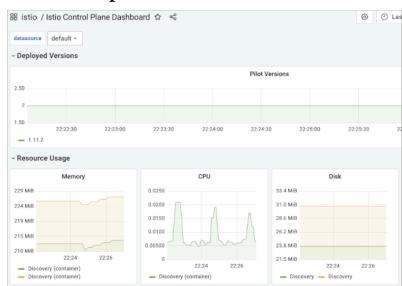
- Istio Mesh summary: describes the network of microservices by HTTP/gRPC and TCP workloads in the Mesh.
- Individual CDP services, such as HiveServer and Impala coordinator: Request and response metrics, such as latency, for each mesh service (HTTP/gRPC and TCP) and client and service workloads metrics.
- Individual workloads using the services: Request and response metrics for each workload in the Mesh (HTTP/gRPC and TCP) and inbound/outbound workload services.

The Istio performance dashboard presents visualizations of the following metrics:

- CDP Usage
- Memory Usage
- Data Rates
- Bytes transferred per second

The Istio control plane dashboard includes memory, CPU, and disk resource usage.

For example:



The sidecar proxy metrics in Istio dashboards reveals the interceptions of network communication between microservices.

The Istio Wasm Plugin extends the Istio proxy capabilities.

## Nodes dashboards

The [Prometheus Node Exporter](#) is used to gather detailed metrics for the AWS and Azure Virtual Machines that host an environment's Kubernetes cluster. These metrics cover both machine and OS level metrics such as CPU, memory, network, processes, time synchronization, disk, and file system. All metric names start with node\_.

Three prebuild dashboards show metrics for AWS/Azure Virtual Machines that host the Kubernetes cluster:

- Cluster Totals
  - Shows CPU/Memory utilization and node counts (shared services/compute for the entire cluster).
- Node Details
  - Shows very detailed metrics for a single node at a time.
- Node Trends
  - Combines CPU, Memory, Disk, and Network metrics for all nodes for node-to-node comparisons.

## Related Information

[Grafana documentation](#)

[Get started with Grafana and Prometheus](#)

[Istio documentation](#)

[Istio in-depth documentation](#)

[Istio github site](#)

[Wasm Plugin](#)

## Limitations of Grafana in CDW

You need to understand the Grafana capabilities in 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.

## Unsupported features

Storing metrics longer than 15 days or consuming more than 90GB of disk space is not supported. Metrics older than 15 days will be deleted. If the stored metrics consume more than 90GB of disk space, metrics will be deleted regardless of the number of days stored.

Custom dashboards you create in Grafana are lost upon restarting or updating the cluster. Only the default dashboards are supported.

## Getting started in Grafana

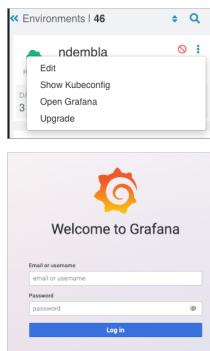
Learn how to log into Grafana in Cloudera Data Warehouse (CDW) Public Cloud.

### Before you begin

- To access Grafana, your CDP administrator must use Management Console to assign the environment to you as a resource.
- You must obtain the DWAdmin role.

### Procedure

- In Data Warehouse Overview , expand the Environment column.
- Search for the environment that is associated with the Virtual Warehouse, click more options  in the environment tile, and select Open Grafana.



- Enter your LDAP user name, or the name provided by your Administrator for logging into Grafana.
- Go back to the environment tile, and select Copy Kubeconf.
- Open a terminal window, paste the contents of your clipboard into a text file, and save the file by the name kube.conf (no extension).
- Export the KUBECONFIG variable equal to the file name.

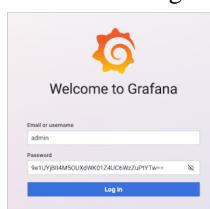
```
export KUBECONFIG=kubeconf
```

- Get the password for Grafana.

For example, use the kubectl command as follows:

```
kubectl get secret grafana -n istio-system -o jsonpath=".data.passphrase" | base64 -D | pbcopy
```

- Go back to the Grafana login dialog, and paste the contents of your clipboard.  
The Grafana login dialog should look something like this:



- Click Log in.

### Related Information

[Assigning resources to users with Management Console User Management](#)

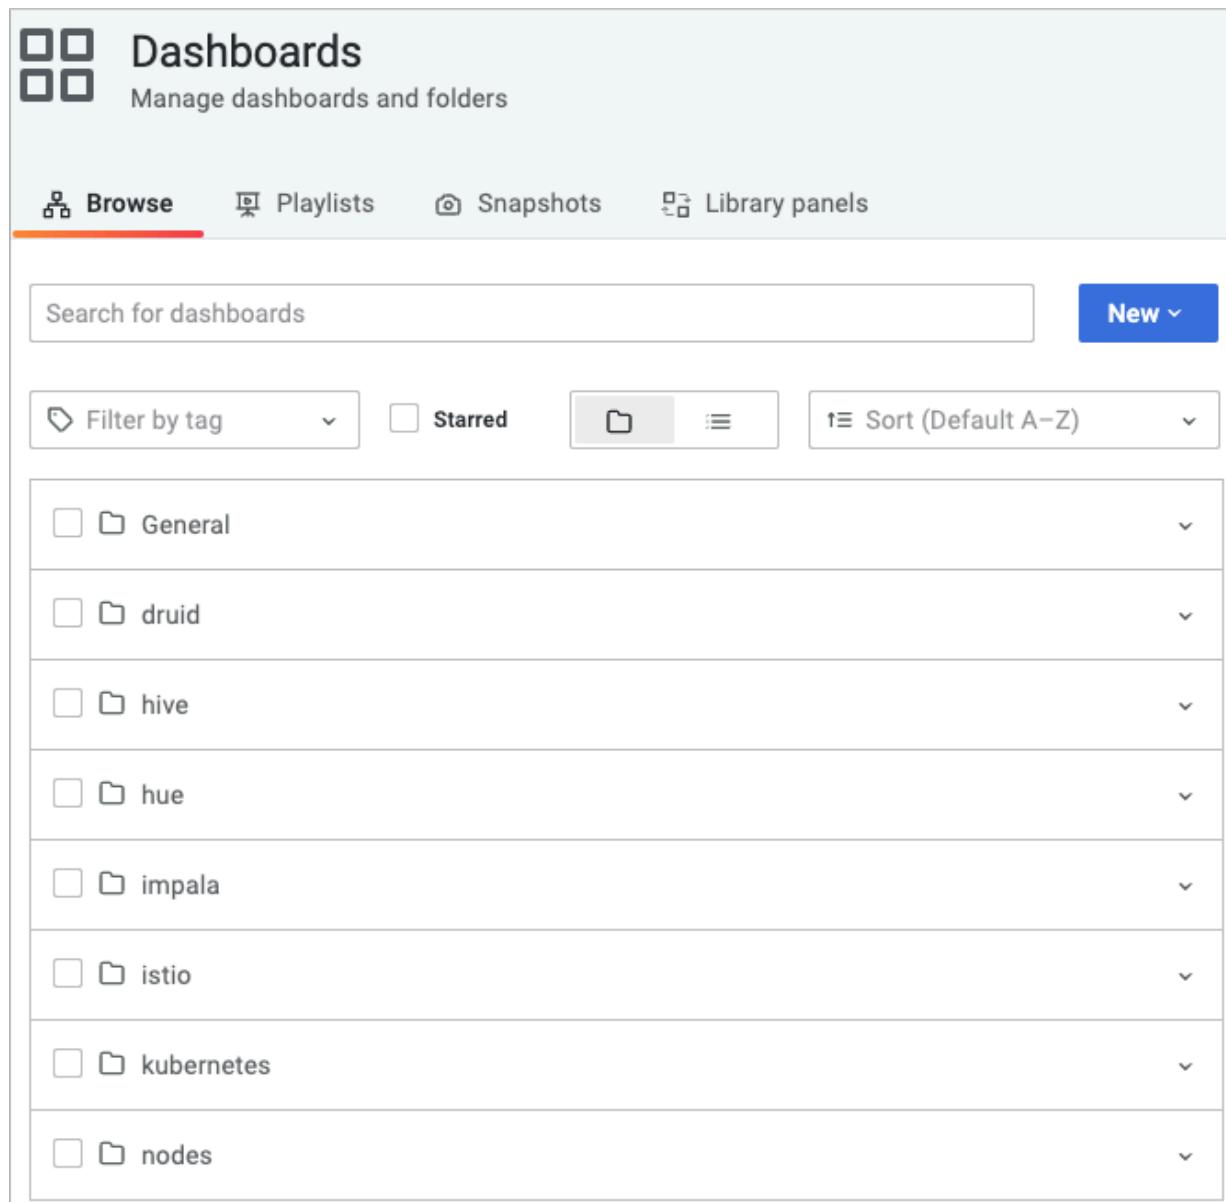
## Viewing prebuilt dashboards

You see how to list dashboard groups, view the Hive dashboard, and see the actual metric data on a point in the X-axis. You can follow the same steps to work with other dashboards.

### Procedure

1. Log into Grafana as described in the previous topic, "Getting Started in Grafana".
2. In the Welcome screen, click grid , and then select Manage.

A list of dashboard groups appears:

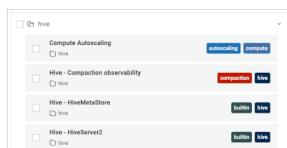


The screenshot shows the Grafana 'Dashboards' management interface. At the top, there are icons for creating a new dashboard () and a folder (), followed by the title 'Dashboards' and the subtitle 'Manage dashboards and folders'. Below this is a navigation bar with four tabs: 'Browse' (which is highlighted with a red underline), 'Playlists', 'Snapshots', and 'Library panels'. A search bar labeled 'Search for dashboards' is positioned below the tabs. To the right of the search bar is a blue 'New' button with a dropdown arrow. Further down are filters for 'Filter by tag' (with a dropdown arrow), 'Starred' (with a checkbox), and 'Sort (Default A-Z)' (with a dropdown arrow). The main content area displays a list of dashboard groups, each represented by a small folder icon and a group name: 'General', 'druid', 'hive', 'hue', 'impala', 'istio', 'kubernetes', and 'nodes'. Each group entry has a small downward arrow to its right.

- Click the name of a dashboard group.

For example, click **hive**.

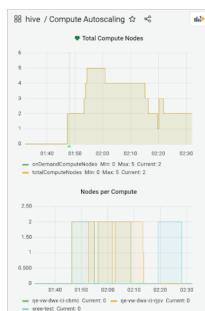
Names of the Hive dashboards in the group appears, a few of which are shown below:



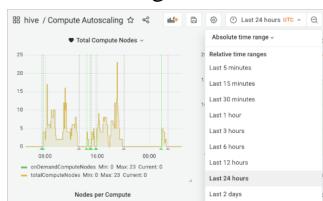
- Click the name of a dashboard group.

For example, click **Compute Autoscaling**.

A number of dashboards appear. Only two are shown below:



- Select viewing of metrics over the last 24 hours.



- Hover over a graphic to reveal metric data for a particular point on the X-axis.



## Related Information

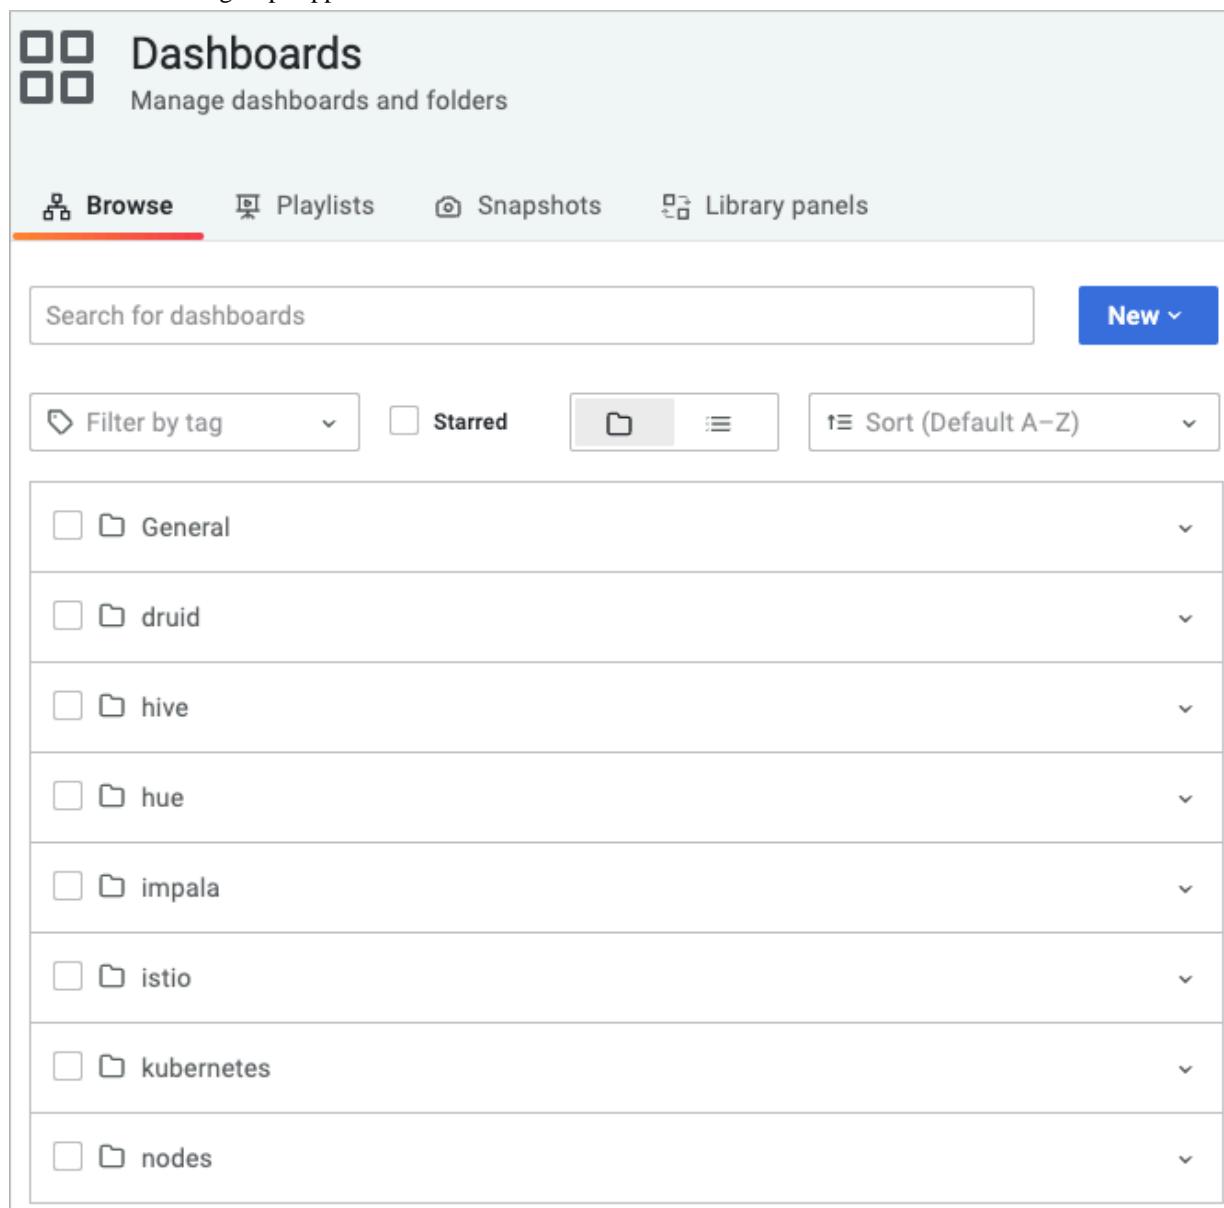
[Assigning resources to users with Management Console User Management](#)

## Monitoring HMS

You can monitor Hive Metastore (HMS), heap usage, and key Hive metrics.

## Procedure

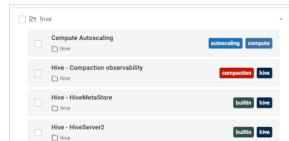
1. In the Welcome screen, click grid , and then select Manage.  
A list of dashboard groups appears:



The screenshot shows the Grafana 'Dashboards' management interface. The 'Browse' tab is selected. At the top, there are icons for 'Dashboards' and 'Folders', and a link to 'Manage dashboards and folders'. Below the tabs are buttons for 'Search for dashboards', 'New', 'Filter by tag', 'Starred', and sorting options ('Sort (Default A-Z)'). The main area lists dashboard groups: General, druid, hive, hue, impala, istio, kubernetes, and nodes. Each group has a folder icon and a dropdown arrow.

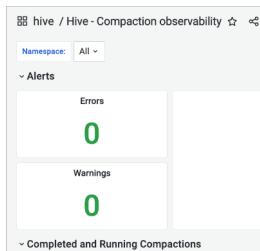
2. Select the **hive** dashboard.

Names of the Hive dashboards in the group appears, a few of which are shown below:

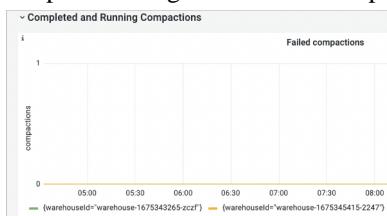


The screenshot shows the 'hive' dashboard group. It contains four sub-dashboards: 'Compute Autoscaling', 'Hive - Compaction observability', 'Hive - HiveMetastore', and 'Hive - HiveServer2'. Each sub-dashboard has a checkbox next to its name.

- Select Hive - Compaction observability, and in Namesparce, select a namespace. In Alerts, the error indicator shows the number of errors and warnings.



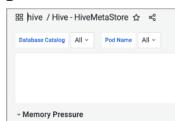
- Expand Completed and Running Compactions. Graphs showing the status of compactions appears.



- Go back to the list of Hive dashboards, and select Hive - HiveMetaStore.

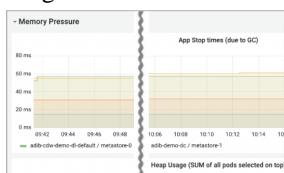


- In Database Catalog, and then in Pod Name, accept the default All, or select a particular Database Catalog and pod.



In Memory Pressure, metrics about the overall status of HMS appear.

- Click each row title to expand and look at various metrics: app stop times due to garbage collection and heap usage.

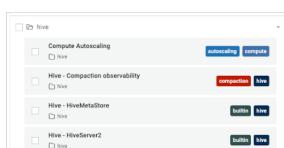


## Monitoring key Hive metrics

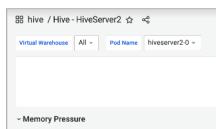
You can get metrics about the status of HiveServer.

### Procedure

- Go back to the list of Hive dashboards, and select Hive - HiveServer2.

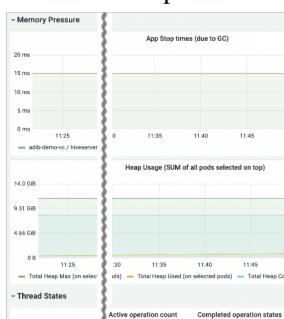


- In Virtual Warehouse, and then in Pod Name, accept the default All, or select a particular Virtual Warehouse and pod.



In Memory Pressure, metrics about the overall status of HiveServer memory usage appear.

- Click each row title to expand and look at various metrics: App stop times due to garbage collection, heap usage, active thread operation count, and completed thread operation states.



- In hive / Hive - Home, you can monitor warehouse growth from metrics on the tables and partitions created hourly.

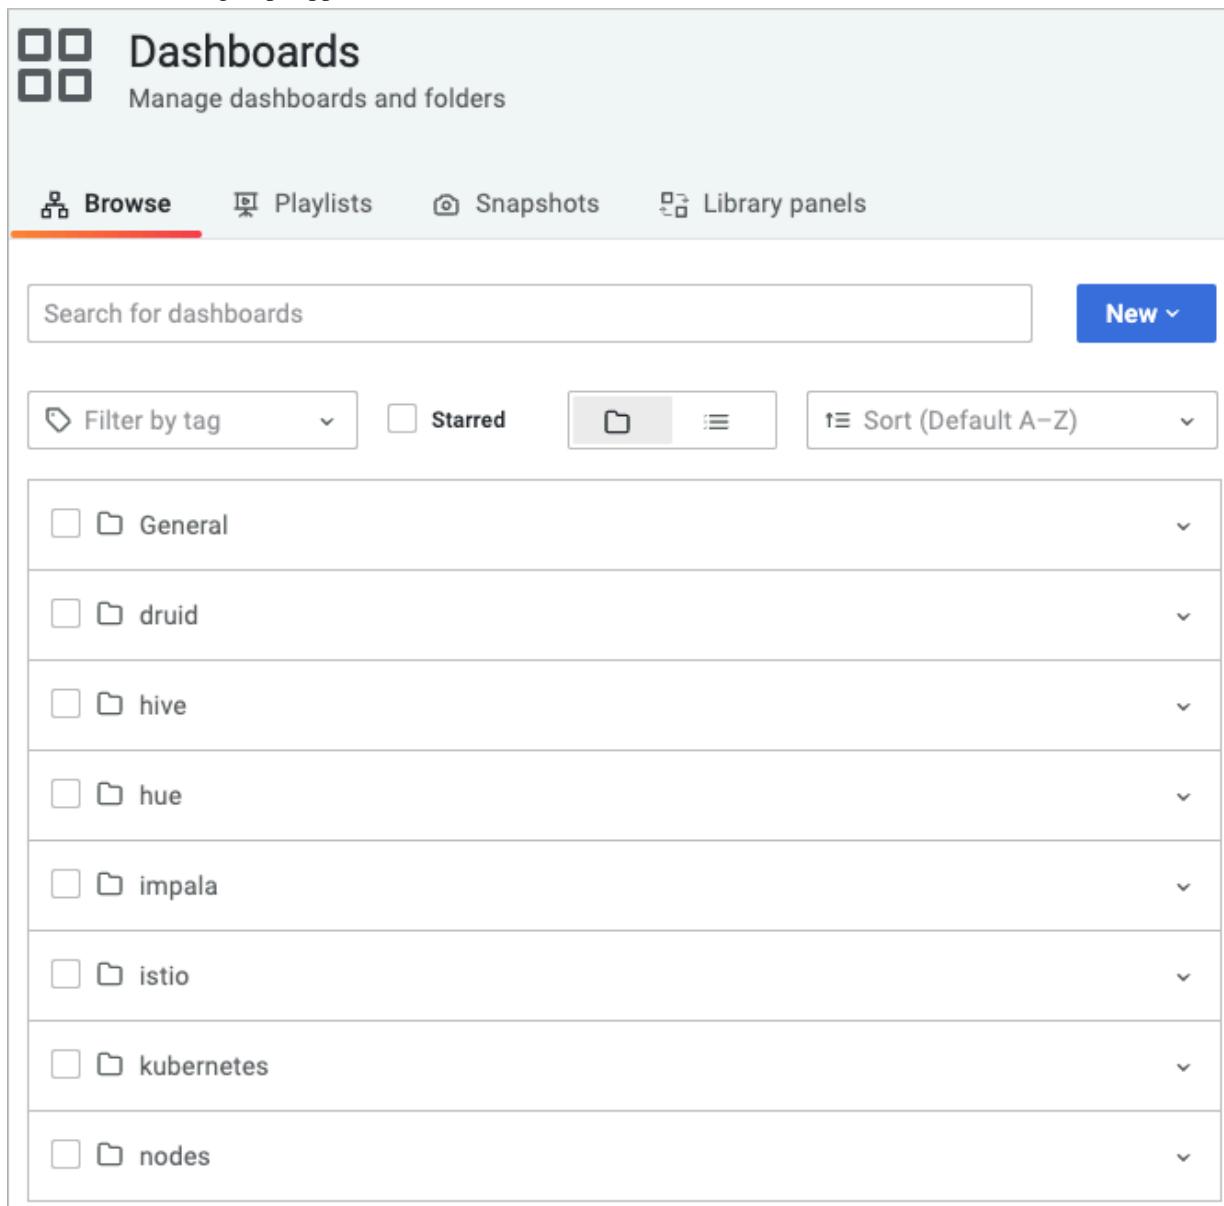


## Monitoring Impala executors

You can monitor Impala executors and get related metrics.

## Procedure

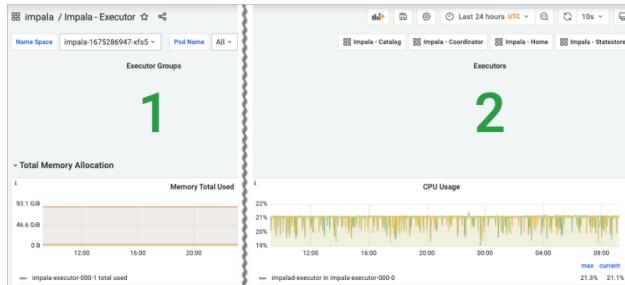
- In the Welcome screen, click grid , and then select Manage. A list of dashboard groups appears:



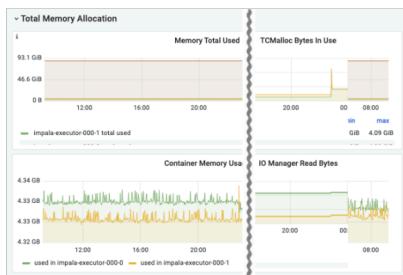
The screenshot shows the Grafana 'Dashboards' management interface. At the top, there are four tabs: 'Browse' (which is highlighted with an orange underline), 'Playlists', 'Saves', and 'Library panels'. Below the tabs is a search bar labeled 'Search for dashboards' and a blue 'New' button. Underneath the search bar are several filter options: 'Filter by tag' (with a dropdown arrow), a 'Starred' checkbox, a folder icon, and a 'Sort (Default A-Z)' dropdown. The main area lists eight dashboard groups, each represented by a folder icon and a group name: 'General', 'druid', 'hive', 'hue', 'impala', 'istio', 'kubernetes', and 'nodes'. Each group entry has a small downward arrow icon to its right.

- Click the impala dashboard group.
- In the list of Impala dashboards, and select Impala - Executor.

The number of Executor Groups and Executors appear.



#### 4. View graphs and metrics of memory usage.



#### 5. Get information about Executor node CPU usage.



## Monitoring Impala catalog

You can monitor Impala catalog from Grafana.

### Procedure

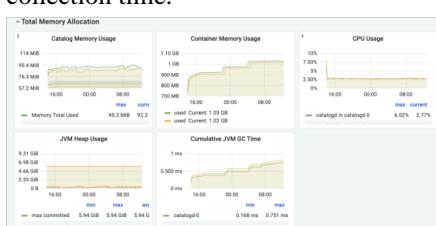
#### 1. In the list of Impala dashboards, and select Impala - Catalog.

The number of databases and tables appear.

#### 2. Select a Name Space, and from the Last 24 hours dropdown, select an time interval to monitor.



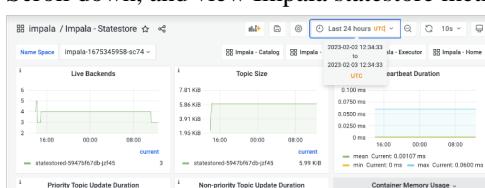
#### 3. View catalog memory usage, container memory usage, cpu usage, JVM heap usage, and cumulative JVM garbage collection time.



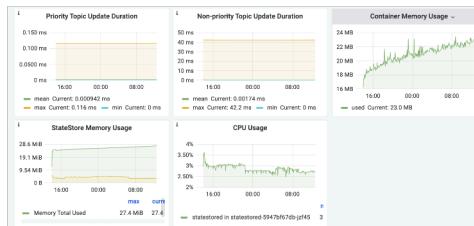
#### 4. Scroll down, and view remote procedure call metrics.



#### 5. Scroll down, and view Impala statestore metrics.



6. Scroll down, and view update, memory, and CPU metrics.

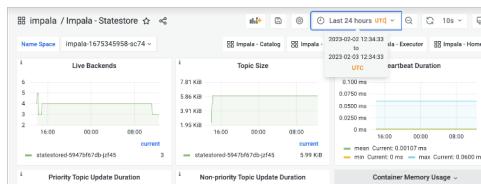


## Monitoring Impala statestore

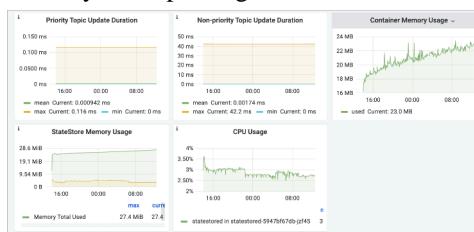
You can monitor Impala statestore from Grafana.

### Procedure

1. In the list of Impala dashboards, and select Impala - Statestore.
2. Scroll down, and view live backends, topic size, and heartbeat duration.



3. Scroll down, and view priority and non-priority topic update duration, container memory usage, statestore memory, and cpu usage.



## Monitoring Impala admission control

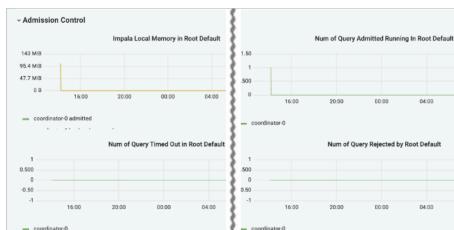
You can monitor Impala admission control from Grafana.

### Procedure

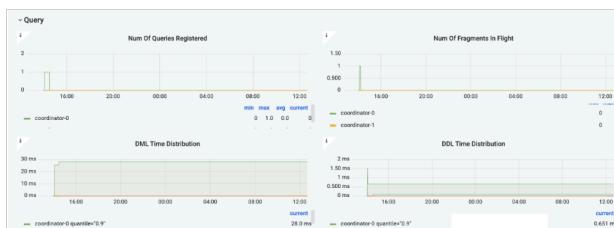
1. In the list of Impala dashboards, and select Impala - Home.
- In Impala Overview, the number of databases, tables, catalogd status, and other metrics appear.



- In Admission Control, view the Impala local memory in root default, the number of queries admitted that are running in root default, and the queries in root default (not shown).



- View the number of queries timed out and rejected in root default.
- Scroll down and in Query view the number of registered queries, fragments in flight, number of open files (not shown), data manipulation language (DML) time distribution, and data definition language (DDL) time distribution.



## Monitoring Impala coordinators

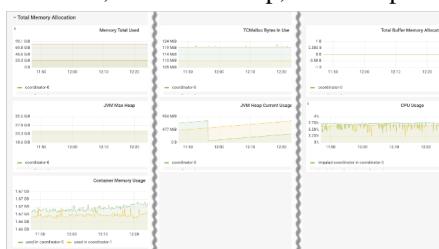
You can monitor Impala coordinators caching, resource use, such as files handles hit, memory consumption, and more.

### Procedure

- In the list of Impala dashboards, click Impala - Coordinator.
- Select a name space and pods, and view cache hit status, remote data cache miss bytes, catalog cache hit rate (not shown), file handles hit count, and file handles miss count.



- Scroll down to Total Memory Allocation and view memory used, TCMalloc bytes in use, total buffer memory allocated, JVM max heap, JVM heap current usage, and CPU usage.

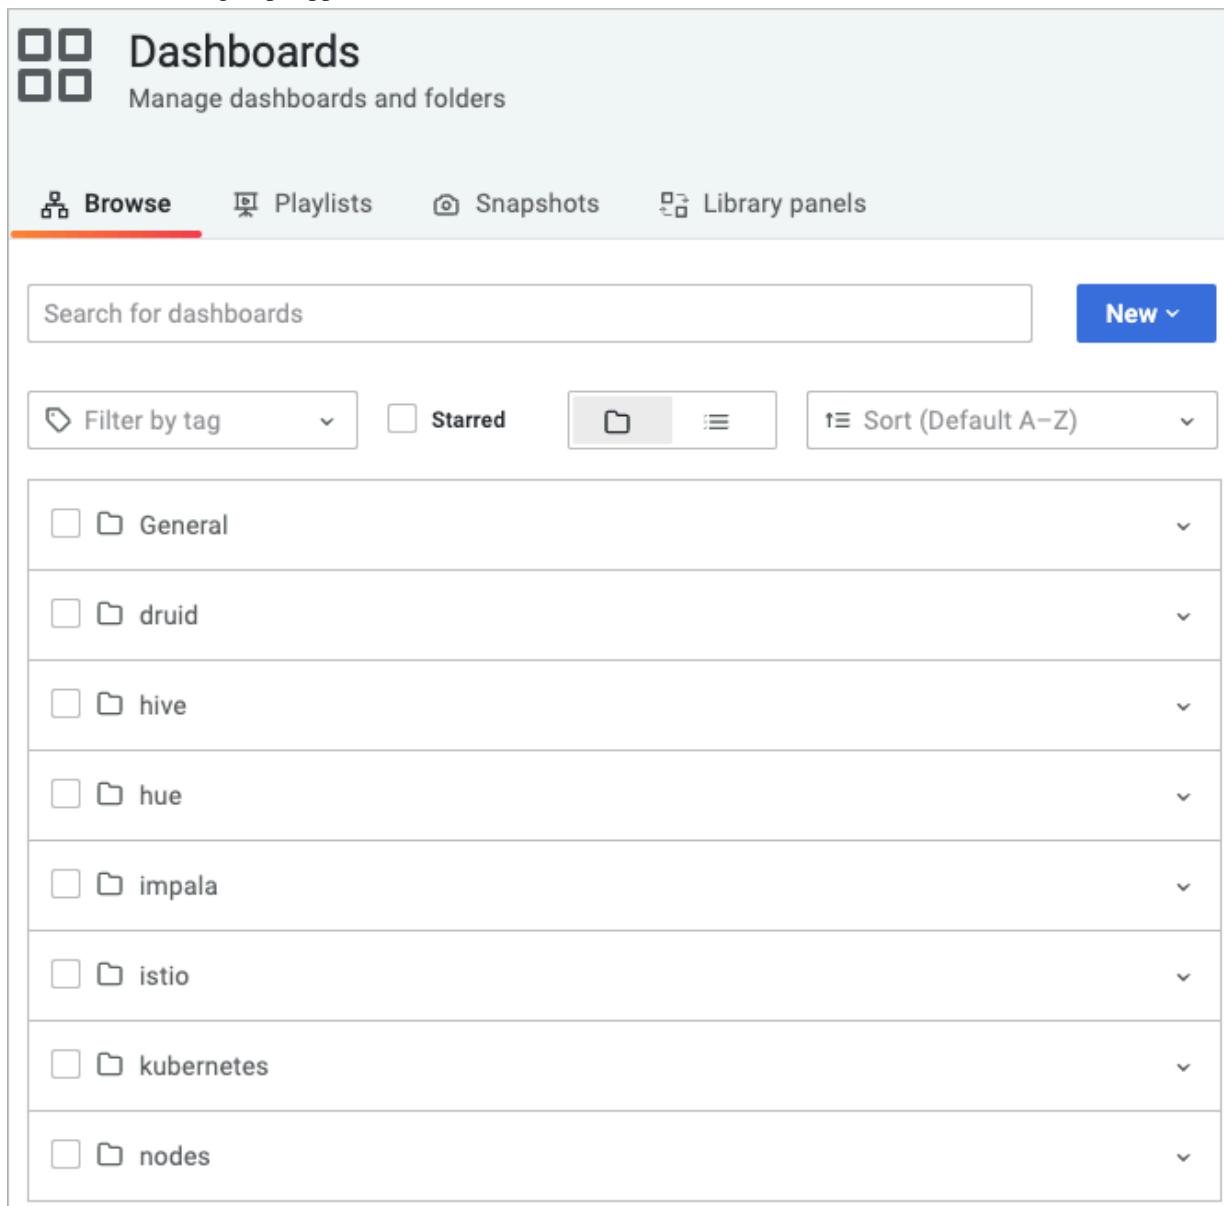


## Monitoring nodes

You can get Kubernetes, cluster health, and node health information by monitoring CDW nodes.

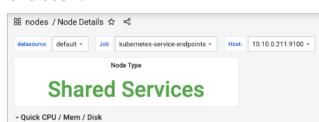
## Procedure

- In the Welcome screen, click grid , and then select Manage. A list of dashboard groups appears:



The screenshot shows the Grafana 'Dashboards' management interface. At the top, there are icons for creating new dashboards and folders, followed by the title 'Dashboards' and the subtitle 'Manage dashboards and folders'. Below this is a navigation bar with four tabs: 'Browse' (which is highlighted with an orange underline), 'Playlists', 'Saves', and 'Library panels'. A search bar labeled 'Search for dashboards' is positioned below the tabs. To the right of the search bar is a blue 'New' button. Underneath the search bar are several filter options: 'Filter by tag' (with a dropdown arrow), a checkbox for 'Starred', and buttons for 'File' and 'List' view. There is also a 'Sort (Default A-Z)' dropdown. The main content area displays a list of dashboard groups, each represented by a folder icon and a group name. The groups listed are: General, druid, hive, hue, impala, istio, kubernetes, and nodes. Each group entry has a small dropdown arrow to its right.

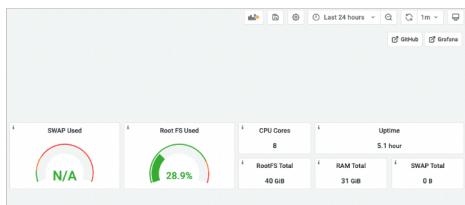
- Click the nodes dashboard group.  
Names of the nodes dashboards in the group appear:
- In the list of nodes dashboards, click Node Details, and then set the node of interest: In datasource, set default for example; in Job, set kubernetes-service-endpoints for example; in Host, set the IP address of a host in the CDW cluster.



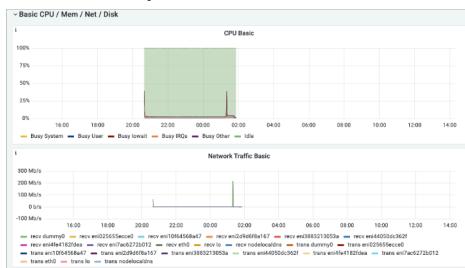
The screenshot shows the 'nodes / Node Details' dashboard. At the top, there are dropdown menus for 'datasource' (set to 'default'), 'Job' (set to 'kubernetes-service-endpoints'), and 'Host' (set to '10.10.0.211:9100'). Below these is a section titled 'Node Type' with a dropdown menu. The main content area features a card titled 'Shared Services' with the sub-section 'Quick CPU / Mem / Disk'.

- Expand Quick CPU/Mem/Disk, and view the percentage of busy CPU, the percentage of average load over the past 5 minutes and over the last 15 minutes, and RAM used.

5. On the right side of the dashboard, with Last 24 hours selected, for example, view the SWAP space used, root file system used, CPU cores used, and uptime of the node.



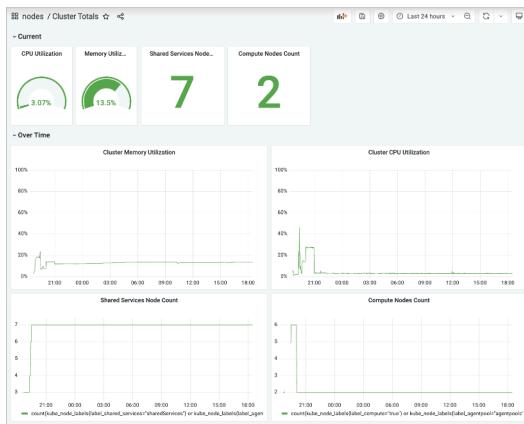
6. Scroll down, expand Basic CPU/Mem/Net/Disk, and see graphs on the left showing metrics about high activity from various sources, such as a busy user or busy interrupt requests (IRQ), and a graph of packets received and transmitted by network traffic.



7. On the right, see graphs of basic memory and disk space use.



8. In the list of nodes dashboards, click Cluster Totals to view cluster health in graphs and metrics over the past 24 hours, as indicated by Cluster Memory Utilization, Cluster CPU Utilization, Shared Services Node Count, and Compute Nodes Count.



9. In the list of nodes dashboards, click Node Trends to view node health over the past 24 hours, as indicated by CPU Utilization, Load Average and CPU Count, Memory Utilization, and Disk Utilization.



# Monitoring Kubernetes Services from Grafana

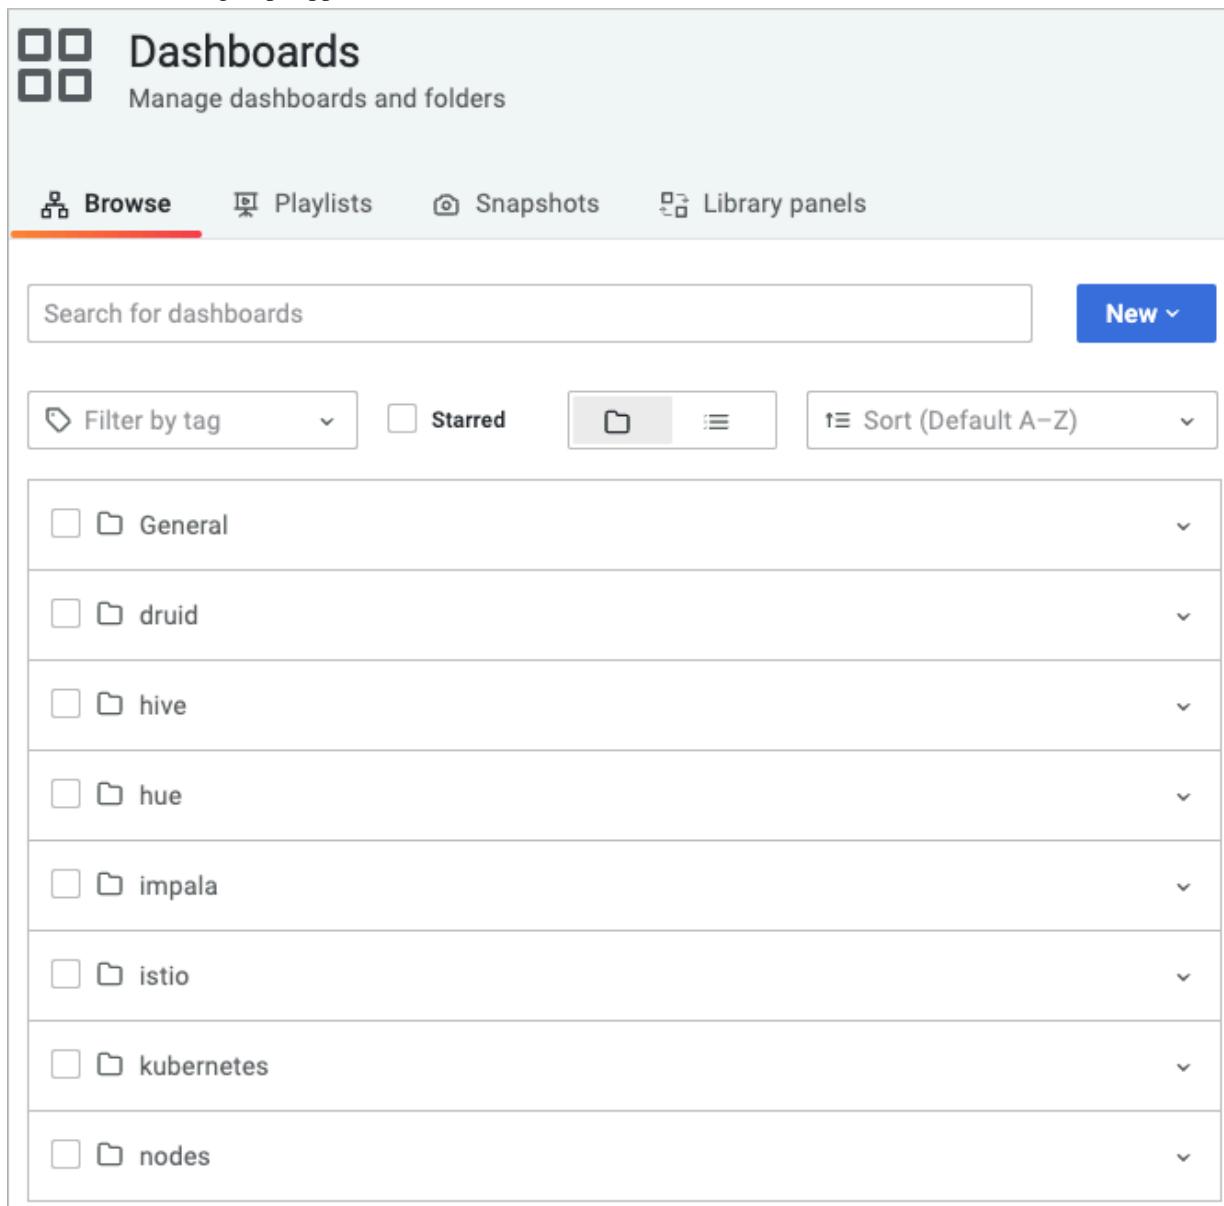
You can monitor the Amazon Elastic Kubernetes Service (EKS) or the Azure Kubernetes Service (AKS) used by your CDW cluster from Grafana.

## About this task

You can also monitor the Amazon Elastic Kubernetes Service (EKS) from the [Cloudera K8s dashboard](#).

## Procedure

- In the Welcome screen, click grid , and then select Manage. A list of dashboard groups appears:



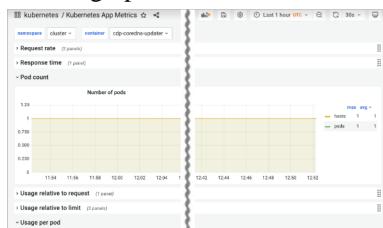
The screenshot shows the Grafana 'Dashboards' management interface. At the top, there are icons for creating new dashboards and folders, followed by the title 'Dashboards' and the subtitle 'Manage dashboards and folders'. Below this is a navigation bar with four tabs: 'Browse' (which is highlighted with an orange underline), 'Playlists', 'Saves', and 'Library panels'. A search bar labeled 'Search for dashboards' is positioned below the tabs. To the right of the search bar is a blue 'New' button. Underneath the search bar are several filter options: 'Filter by tag' (with a dropdown arrow), a checkbox for 'Starred', and buttons for 'File' and 'List' view. There is also a 'Sort (Default A-Z)' dropdown. The main content area displays a list of dashboard groups, each represented by a folder icon and a group name. The groups listed are: General, druid, hive, hue, impala, istio, kubernetes, and nodes. Each group entry has a small downward arrow icon to its right.

- Click the kubernetes dashboard group.

Names of the kubernetes dashboards in the group appear:

- Click the Kubernetes App Metrics, and then select a namespace, for example cluster, and container, for example a container named cdp-coredns-updater.

View a graph and metrics of the number of pods.



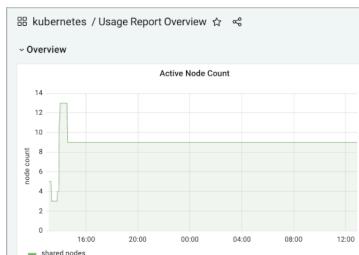
#### 4. View the CPU usage per pod.



#### 5. On the right, view the memory usage per pod.



#### 6. From the kubernetes dashboard group, click Usage Report Overview, and see the active nodes on the X axis of the Active Node Count graph.

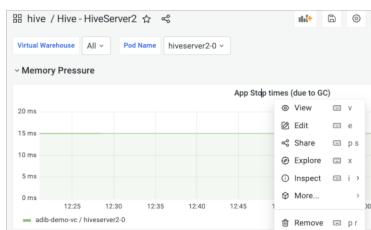


## Updating a dashboard graph

You can view, edit, share, explore, or remove a dashboard graph, and more.

### Procedure

#### 1. Open a dashboard, for example the HiveServer2 dashboard, and click the title in a row, for example, the **App Stop times** title.



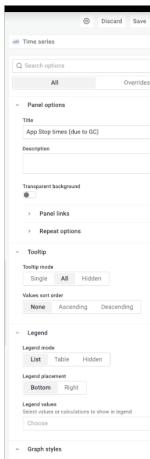
#### 2. Click Edit.

The edit panel appears.



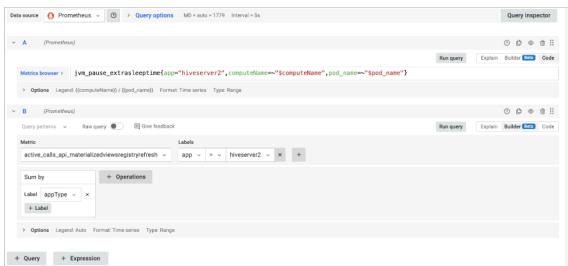
#### 3. At the top left side of the edit panel, change the Virtual Warehouse and Pod on which the graph is based.

4. At the top right side of the edit panel, click  to add variables and permissions to the panel.
5. On the right side of the edit panel, select a graph type, such as Time Series, and specify panel options, such as the title of the graph, tooltip mode, legend, graph styles, and more.



6. Scroll down the edit panel, and change query options.

For example, in Metrics browser, change the metrics to be graphed from jvm\_pause\_extrasleepetime.



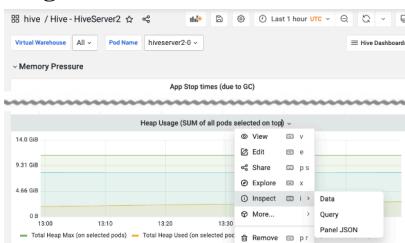
7. On the right side of the edit panel, use the Builder code option  to frame the query and Explain to understand the query.
8. At the bottom of the edit panel, click + Query to add a query.
9. At the top-right, click Apply to save changes.

## Inspecting dashboard data and queries

You can drill down into details of resource usage of Hive or Impala, get statistics about usage, and run queries on the details.

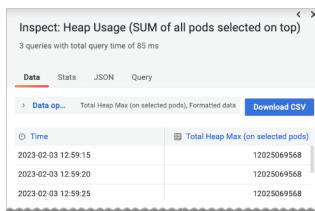
### Procedure

1. Open a dashboard, for example the HiveServer2 dashboard, and click the title in a row, for example, the **Heap usage** title.



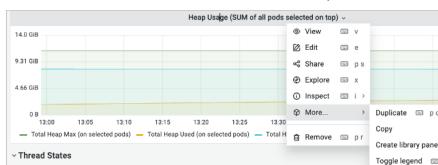
**2. Click Inspect Data .**

The total max heap usage over time appears.



**3. Click Stats or Query to view statistics or queries.**

**4. In the HiveServer2 dashboard, and click the Heap usage title More Duplicate panel content.**



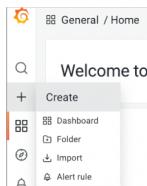
**5. In the HiveServer2 dashboard, and click the Heap usage title More Copy panel content to another dashboard.**

## Creating a custom dashboard

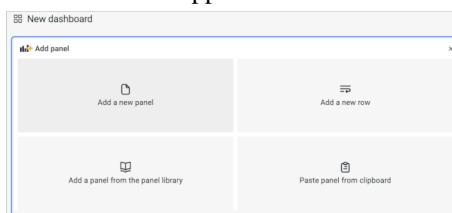
You can create a custom dashboard in a few steps.

### Procedure

**1. Log into Grafana, and in the Welcome screen, click + Create**

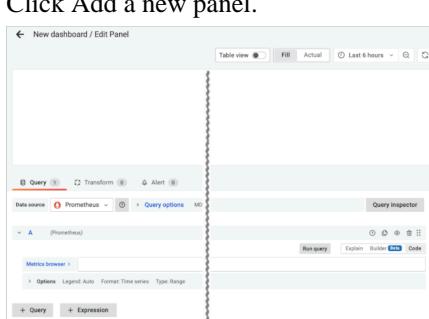


New dashboard appears.

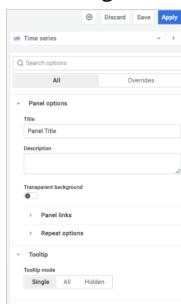


**2. At the top right side of the edit panel, click to add variables, permissions, and names to the new dashboard.**

**3. Click Add a new panel.**



4. On the right side of the New edit panel, click to add variables and permissions to the new dashboard.



5. Edit the panel and rename it test1.

6. Click Apply.

Test1 is added to the dashboard.

7. Edit the panel, rename it test2, and click Apply.

Test2 is added to the dashboard with Test1.



## Meeting prerequisites to set up alerts

Before you can set up an alert triggered by a dashboard event, you must configure an SMTP (Simple Mail Transfer Protocol) server, an alert recipient, and a notification policy.

### Procedure

1. Configure grafana using kubectl.

```
kubectl edit configmap grafana -n istio-system
```

For more information, see "[Granting remote access to Kubernetes](#)".

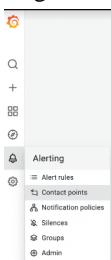
2. Update the grafana.ini to configure SMTP, setting user to the alert account email id and the password to the app password generated in your gmail account.

```
...
[smtp]
enabled = true
host = smtp.gmail.com:465
user = youremail@gmail.com
password = substitute_your_password
;cert_file=
;key_file=
skip_verify = true
from_address = alertgrafan@gmail.com
from_name = Grafana
* EHLO identity in SMTP dialog (defaults to instance_name)
;ehlo_identity = dashboard.example.com
;startTLS_policy = NoStartTLS
```

3. Refresh the grafana pod.

```
kubectl get pod -A grep -i grafana
kubectl delete pod grafana-9bcb5d5c-w8r6c -n istio-system
```

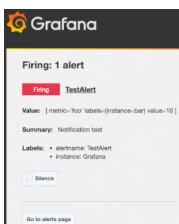
4. Log into Grafana, and click Alerting Contact points .



5. In Contact point type, select an alert receiver type.
6. In Addresses, enter an email address, and then in Message, compose an email message body. For example, enter an email message body test grafana email.

A screenshot of the 'Contact points' configuration page in Grafana. The top navigation bar shows 'Alerting' selected. The main form is titled 'Update contact point'. It includes fields for 'Name' (set to 'grafana-default-email'), 'Contact point type' (set to 'Email'), 'Address' (set to 'amit.mishra@cloudera.com'), and an optional 'Message' field containing 'test grafana email'. At the bottom, there are 'Optional Email settings' checkboxes for 'Single email' and 'Send a single email to all recipients', and a 'Notification settings' checkbox for 'Disable resolved message'.

7. Click Test to verify that a test alert arrives in your inbox.

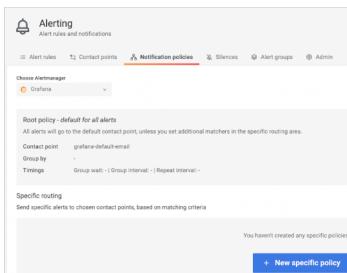


8. Scroll up and in Name, enter the contact point name.

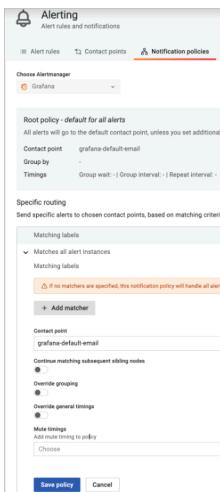
9. Log into Grafana, and click Alerting Notification policies .



10. Click New specific policy.



11. In Contact Point, select the contact point to receive notification, and save the policy.

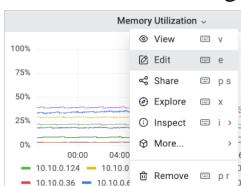


## Creating alerts

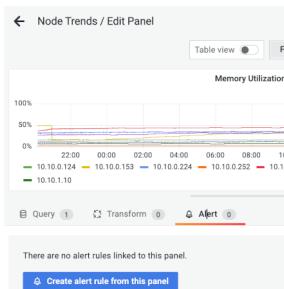
After meeting prerequisites to set up alerts, you can create an alert for an existing dashboard.

### Procedure

1. Click the title of the graph that you want to fire an alert, and select Edit.



2. In the edit panel that appears, click Alert Create alert rule from this panel .



Create alert rule appears. The Rule name is set to Memory Utilization.

3. In Create alert rule scroll down, and in Group, enter a meaningful name for the alert group, for example, test alert..

4. In B, define a classic condition to trigger the alert, for example when last() of A is above 20.

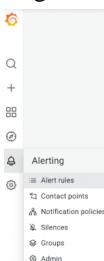
5. In Define alert conditions, select B for the expression that triggers the alert.

6. Click Preview alerts, and then scroll up to the top right of Create alert rule, and click Save.

## Reviewing alerts and notifications

### Procedure

1. Log into Grafana, and click Alerting Alert rules .



- In Alert rules, in the **hive** directory, click Firing to review the alert in the UI.

The top screenshot shows the 'Alerting' interface with 'Alert rules' selected. It displays a search bar for 'Search by data source' (All data sources) and 'Search by label' (State: Firing). Below this, it shows '24 rules: 2 firing, 22 normal'. Under the 'Grafana' section, there's a tree view with 'hive' expanded, showing a table with 'State' (Firing) and 'Name' (HiveServer2 Memory - Max). Another tree view for 'nodes' shows a table with 'State' (Firing) and 'Name' (TCP Connections). The bottom screenshot shows the Grafana interface with a single 'Firing' alert for 'TCP Connections'. The alert details show multiple metrics with labels like 'instance\_only' and values ranging from -1.0 to 1.0.

## Monitoring Kubernetes resources from K8s dashboard

As a Cloudera Data Warehouse (CDW) administrator, you can monitor Kubernetes resources in your CDW cluster. On the K8S dashboard, you can view the state of the resources, such as CPU and memory usage, see the status of pods, and download logs.

The dashboard can provide insights into the performance and health of a CDW cluster. From the dashboard, if authorized, you can monitor the environments of any CDW cluster efficiently. You do not need to copy/paste kubeconfig files to switch to monitoring another environment. Monitoring the dashboard can help keep your cluster running smoothly and efficiently.

By default, the dashboard is disabled. You follow instructions in the next topic to activate and use the dashboard. Using the dashboard incurs some cloud cost, and is designed to time out automatically after 4 hours to prevent wasting resources. Cloudera recommends deactivating the dashboard when not in use to reduce cloud expenses.

### Prerequisites

- You have an AWS or Azure environment in Cloudera Data Warehouse Public Cloud.
- You have [activated](#) your environment.
- You obtained the CDW Admin role.

## Activating the K8S dashboard

From your environment, you can flip a switch to activate, or deactivate the K8S dashboard. You can then click a link to view the K8S dashboard in your web browser.

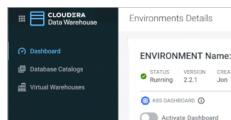
### Before you begin

You must meet the prerequisites listed in the previous topic.

## Procedure

1. In the Data Warehouse service, expand Environments by clicking More....
2. In Environments, search for and locate the environment that you want to view.
3. Click Edit.

In Environment Details, you see Activate Dashboard.



4. Toggle Activate Dashboard to on.
5. Wait for the K8S Dashboard Activated success message to appear.  
A link to the K8S dashboard appears in your environment tile.

## What to do next

- Follow steps in the next topic to familiarize yourself with the K8S dashboard.
- Deactivate the dashboard when not in use because an activated dashboard incurs a cost.

## Using the K8S dashboard

You see how to view the dashboard and get ideas about the types of K8S metrics, charts, and other visualizations that appear on the dashboard. You might use the dashboard to keep your cluster running efficiently and troubleshoot problems.

### Before you begin

- You have met the prerequisites listed in Monitoring Kubernetes resources.
- You have activated the K8S dashboard.

## Procedure

1. Log in to the CDP web interface, navigate to Data Warehouse Overview , and select your environment. Environment Details shows the K8S dashboard is activated.

STATUS	VERSION	CREATED BY
Running	1.6.3-b241	John

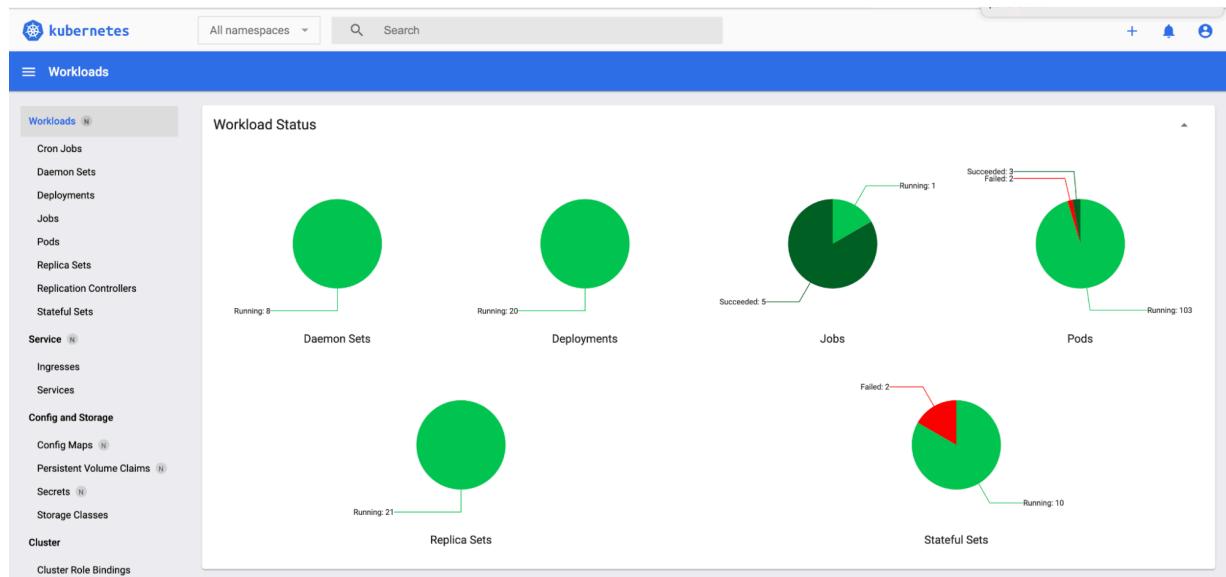
Activate Dashboard [Open](#)

2. In Environment Details, click the Open link.

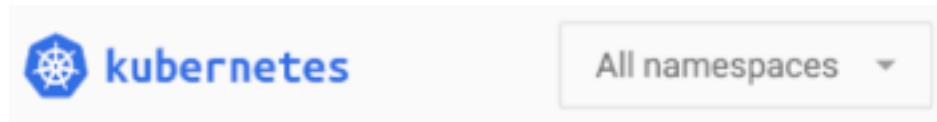
 **Activate Dashboard** [Open](#)

The K8S dashboard appears.

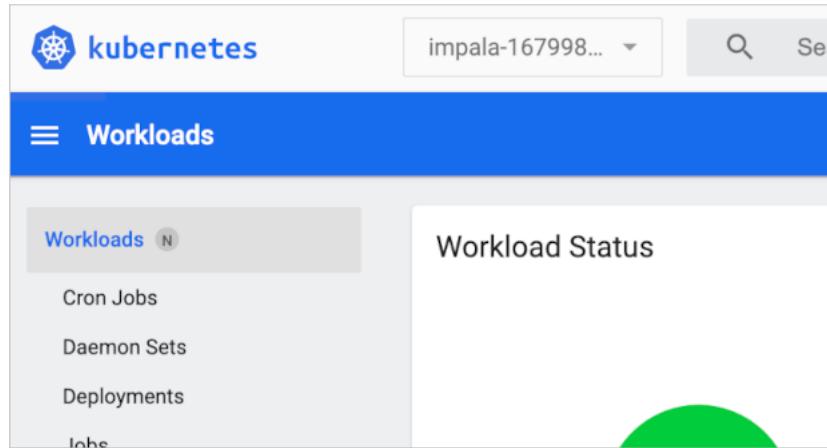
3. View the status of your containerized applications.  
For example:



4. Click All Namespaces and select a specific Namespace.

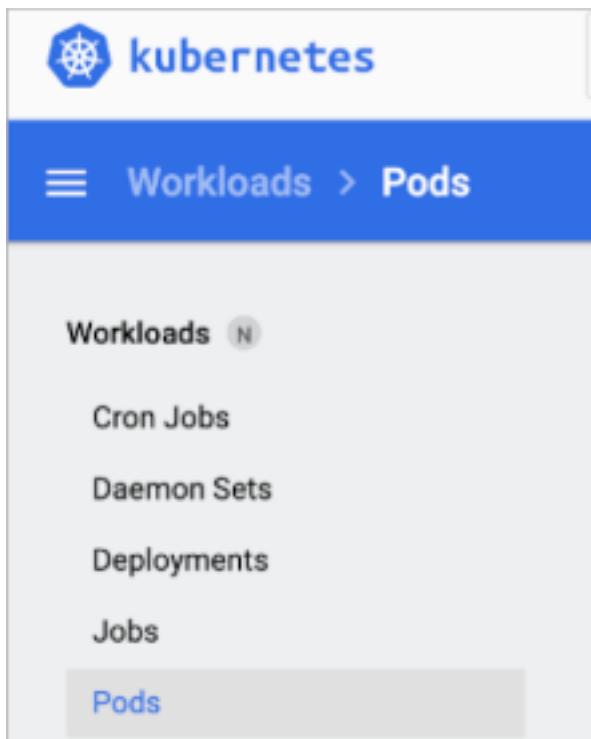


The workload status of the namespace, impala-167998, in this example appears.



5. Kubernetes dashboard features are available and documented on the [Kubernetes site](#).

6. To download logs, click Pods.



The dashboard of pods appears.

Name	Namespace	Images	Labels	No
kubernetes-dashboard-6422902b-75665a1e-8658c5787c-xgrtd	kubernetes-dashboard	container-dev.repo.cloudera.com/cloudera_thir dparty/kubernetesui/dashbaord:v2.5.1 container-dev.repo.cloudera.com/cloudera_thir dparty/istio/proxyv2:1.1 5.1	app.kubernetes.io/component: kubernetes-dashboard app.kubernetes.io/instance: kubernetes-dashboard-6422902b-75665a1 e app.kubernetes.io/managed-by: Helm	ip-2.c

7. Click Options  , and select Logs.  
For example:

Name	Namespace	Images	Created	
kubernetes-dashboard-6422902b-75665a1e-8658c5787c-xgrtd	kubernetes-dashboard	containerd.io/dparty/kubernetes-dashboard: containerd.io/dparty/istio:5.1	12 minutes ago	<span>⋮</span> <div> <a href="#">Logs</a>   <a href="#">Exec</a>   <a href="#">Edit</a>   <a href="#">Delete</a> </div>
uninstall-k8s-dashboard-45pwt	kubernetes-dashboard	containerd.io/dparty/binary:1.23 containerd.io/dparty/k8s-uninstall:1.23	13 ago	<span>⋮</span>

Logs appear.

Logs from kubernetes-da... in kubernetes-da... [Download logs](#) ⋮

```
2023/03/28 07:40:24 Getting list of all jobs in the cluster
2023/03/28 07:40:24 [2023-03-28T07:40:24Z] Outcoming response to 127.0.0.1:8080
2023/03/28 07:40:24 [2023-03-28T07:40:24Z] Incoming HTTP/1.1 GET /api/v1/namespaces/default/jobs?labelSelector=&page=1&sortBy=d,creationTimestamp
2023/03/28 07:40:24 (MISSING)itemsPerPage=10&sortBy=d,creationTimestamp request :
2023/03/28 07:40:24 Getting list of all deployments in the cluster
2023/03/28 07:40:24 [2023-03-28T07:40:24Z] Incoming HTTP/1.1 GET /api/v1/namespaces/default/deployments?labelSelector=&page=1&sortBy=d,creationTimestamp
2023/03/28 07:40:24 127.0.0.1:8080 - "GET /api/v1/namespaces/default/cronjobs?labelSelector=&page=1&sortBy=d,creationTimestamp" 200 1034
2023/03/28 07:40:24 [2023-03-28T07:40:24Z] Incoming HTTP/1.1 GET /api/v1/namespaces/default/cronjobs?labelSelector=&page=1&sortBy=d,creationTimestamp
2023/03/28 07:40:24 Getting list of all cron jobs in the cluster
2023/03/28 07:40:24 [2023-03-28T07:40:24Z] Incoming HTTP/1.1 GET /api/v1/namespaces/default/pods?labelSelector=&page=1&sortBy=d,creationTimestamp
2023/03/28 07:40:24 127.0.0.1:8080 - "GET /api/v1/namespaces/default/pods?labelSelector=&page=1&sortBy=d,creationTimestamp" 200 1034
2023/03/28 07:40:24 [2023-03-28T07:40:24Z] Incoming HTTP/1.1 GET /api/v1/namespaces/default/pods?labelSelector=&page=1&sortBy=d,creationTimestamp
2023/03/28 07:40:24 Getting list of all pods in the cluster
W0328 07:40:24.827364       1 warnings.go:70] batch/v1beta1 CronJob
```

8. Click Download logs

# Forwarding logs to your observability system

You can forward logs from environments activated in Cloudera Data Warehouse (CDW) to observability and monitoring systems such as Datadog, New Relic, or Splunk. You learn how to configure a CDW environment for these systems.

## About this task

After configuring log forwarding as described in this task, logs flow from CDW to your system automatically. You enjoy the convenience of sorting, searching, and viewing logs on your own system instead of grepping logs from diagnostic bundles on S3 or ABFS. In addition to configuring the log forwarding, you configure removal of debug logs and text strings from the logs. You can configure log forwarding to one of the following observability systems:

- Datadog — <https://github.com/DataDog/fluent-plugin-datadog>
- Honeycomb.io — <https://docs.honeycomb.io/getting-data-in/logs/log-collectors/fluentd/>
- New Relic — <https://github.com/newrelic/newrelic-fluentd-output>
- Splunk — <https://github.com/splunk/fluent-plugin-splunk-hec> (covers both Splunk-HEC and Splunk-SCS)

You create the log forwarding configuration in valid fluentd format. The configuration is inserted into a larger fluentd configuration. All fluentd events are copied and relabeled with the new label @cloudera\_cdw. Your custom configuration is then inserted between <label> tags:

```
<label @cloudera_cdw>  
customer config goes here  
</label>
```

You can use any of the built-in fluentd filter, formatter, parser, or output plugins to build the custom config.

## Before you begin

- Before configuring log forwarding you must [activate an AWS environment](#) or [activate an Azure environment](#) in CDW.
- You must be [familiar with fluentd](#) and accept the responsibility of configuring log forwarding to your observability systems.

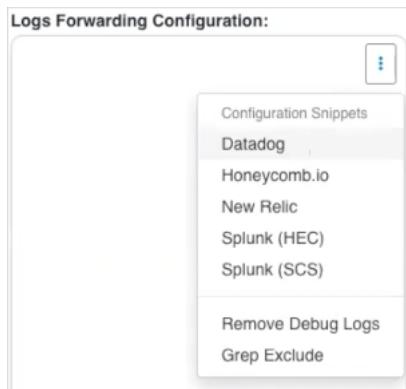
## Procedure

1. In the Data Warehouse service, expand the Environments column by clicking More....
2. In Environments, locate your environment, and click Options  Edit Observability tab .



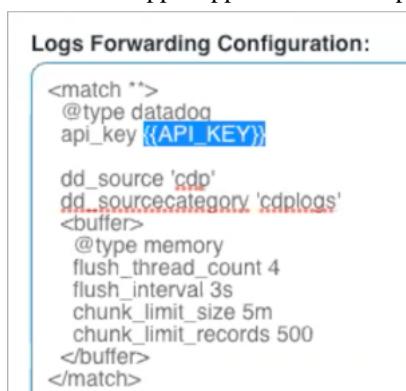
3. Decide how you want to create the fluentd config.
  - Write your own fluentd config from the ground up.
  - Use a Cloudera-provided snippet as a template to write your fluentd config.

4. In Log Forwarding Configuration, click Options .

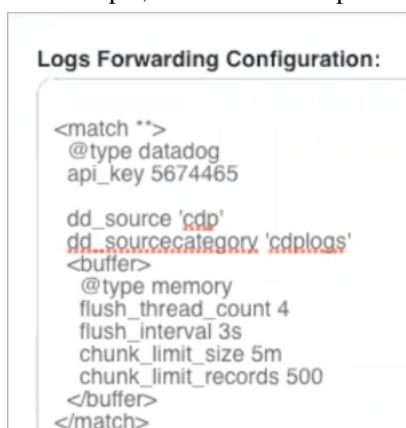


5. Select one of the systems, such as Datadog, to configure.

A fluentd snippet appears. For example, the Datadog snippet appears:



6. Replace the snippet with the fluentd config you wrote from the ground up, or customize the provided snippet. For example, to customize the provided snippet replace the placeholder {{API Key}} with the actual key.



7. (Optional) If debug level log messages are not desired, add a fluentd filter to remove them: In the environment,

click Options , and select Remove Debug Logs.

The fluentd snippet appears for removing debug logs. For example:



```
<filter **>
@type grep
<exclude>
key log
pattern /debug/
</exclude>
</filter>
|
<match **>
@type datadog
api_key 5674465

dd_source 'cdp'
dd_sourcecategory 'cdplogs'
<buffer>
@type memory
flush_thread_count 4
flush_interval 3s
chunk_limit_size 5m
chunk_limit_records 500
</buffer>
</match>
```

No user customization is necessary to remove debug logs.

8. (Optional) If certain log messages do not provide value for you, remove them with a fluentd grep exclude filter: In

the environment, click Options , select Grep Exclude, and replace {{PATTERN}} with the grep expression that matches the phrase you want to exclude.



```
<filter **>
@type grep
<exclude>
key log
pattern /debug/
</exclude>
</filter>
<filter **>
@type grep
<exclude>
key log
pattern /ldontshowup/
</exclude>
</filter>

<match **>
@type datadog
api_key 5674465

dd_source 'cdp'
dd_sourcecategory 'cdplogs'
<buffer>
@type memory
flush_thread_count 4
flush_interval 3s
chunk_limit_size 5m
</buffer>
```

For more information about using Grep Exclude, see <https://docs.fluentd.org/filter/grep>.

9. If you use a proxy server for outbound traffic, provide the proxy server's CA certificates PEM bundle as described in the next task.

## 10. Click Apply.

CDW tests the log forwarding configuration and proxy CA certificates bundle, and saves the configuration if both are valid. An invalid log forwarding config error message appears in the event of a configuration problem. For example:

If your configuration is valid, CDW initiates a restart of fluentd to apply the updated config. You see the following indicators of success:

- The environment Running indicator changes, blinks Updating, and then once again says Running.
- You see logs appearing in your observability system.

Many factors affect how long it takes for forwarding to begin, but generally, the bigger your CDW environment, the longer it takes.

## Providing proxy CA certificates

If you use a TLS-terminating proxy server to inspect outbound internet traffic, you need to provide the proxy server's CA certificates bundle in PEM bundle format when you configure log forwarding.

### About this task

You learn how to use the Observability tab in CDW Environment Details to configure the Proxy CA Certificates (PEM Bundle) field.

### Before you begin

Before you apply the proxy CA certificate to a configuration of log forwarding, you must provide a configuration in the Logs Forwarding Configuration section of the Observability tab.

### Procedure

1. In the Data Warehouse service, expand the Environments column by clicking More....
2. In Environments, locate your environment, and click Options Edit Observability tab . Environment details include a UI for configuring log forwarding.



3. Obtain and copy your proxy server's CA certificates PEM bundle.
4. In Proxy CA Certificates, paste the copy of the PEM bundle.
5. Click Apply.

If the certificate and log forwarding configuration are valid, log forwarding begins. If the certificates are invalid, an error message occurs.

The log forwarding configuration and certificates are not saved.