

Using YARN Web UI and CLI

Date published: 2020-07-28

Date modified: 2021-12-13



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

Accessing the YARN Web User Interface.....	4
Viewing the Cluster Overview.....	4
Viewing nodes and node details.....	5
Viewing queues and queue details.....	6
Viewing all applications.....	8
Searching applications.....	9
Viewing application details.....	10
UI Tools.....	11
Using the YARN CLI to viewlogs for applications.....	11

Accessing the YARN Web User Interface

Access the YARN Web User Interface to monitor clusters, queues, applications, services, and flow activities.

About this task

You can access the YARN Web User Interface from Cloudera Manager to monitor the cluster.


Procedure

1. In Cloudera Manager, select the YARN service.
 2. Click the Web UI tab.
 3. Select the ResourceManager Web UI.
- The ResourceManager Web UI opens in a new browser window.

Note, that in a ResourceManager High Availability configuration, if you select the non-active ResourceManager the link directs to the web UI of the active ResourceManager.

Viewing the Cluster Overview

The Cluster Overview page shows cluster resource usage by applications and queues, information about finished and running applications, and usage of memory and vCores in the cluster.



Cluster Resource Usage by Applications

Displays the percentage of cluster resources in use by applications and the percentage available for usage.

Cluster Resource Usage by Leaf Queues

Displays the percentage of cluster resources in use by leaf queues and the percentage available for usage.

Finished Apps From All Users

Displays the number of completed, killed, and failed applications.

Monitor Running Apps

Displays the number of pending and running applications.

memory-mb - Usages

Displays the amount of used and available memory.

vcores - Usages

Displays the number of used and available virtual cores.

Monitor Node Managers

Displays the status of the Node Managers under the following categories:

- Active
- Unhealthy
- Decommissioning
- Decommissioned

Viewing nodes and node details

The Nodes page on the YARN Web User Interface enables you to view information about the cluster nodes on which the NodeManagers are running.

The Nodes page displays details under the following headers: Information, Node Status and Nodes Heatmap Chart.

Information

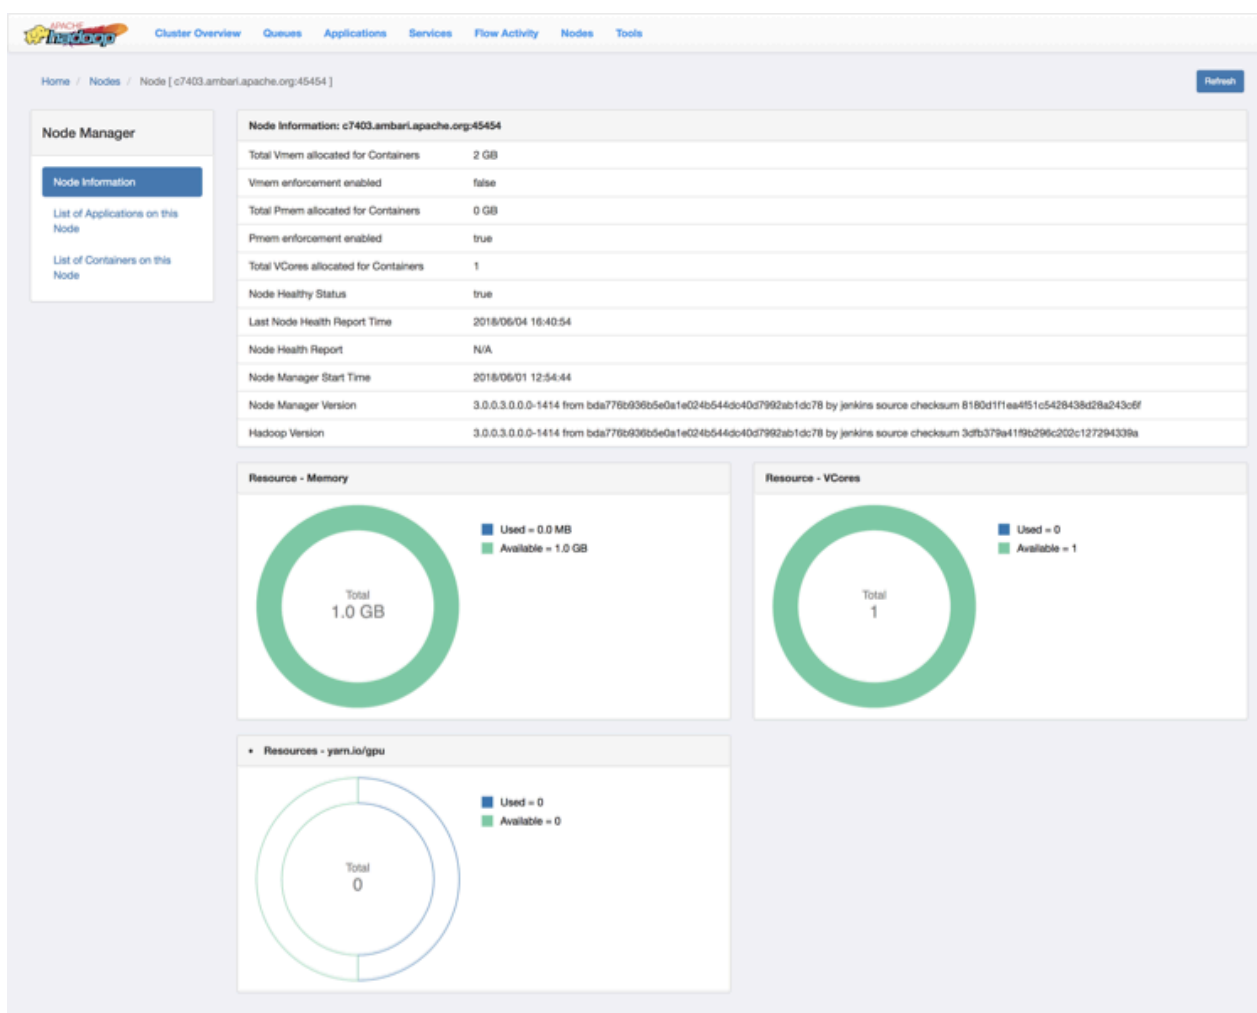
The Information tab displays the node details as shown in the following figure:

Node Label	Rack	Node State	Node Address	Node HTTP Address	Containers	Mem Used
default	/default-rack	RUNNING	c7404.ambari.apache.org:45454	c7404.ambari.apache.org:8042	0	0 B

You can sort through any of the columns to view the details of the required nodes. You can also search to find the specific node labels from the entire list.

Node Manager page

To view the Node Manager page of any node label, click the corresponding Node HTTP address. The Node Manager page displays details of a node as shown in the following figure:



You can also view the resource usage in the following categories in a pie-chart representation:

1. Memory
2. Vcores
3. yarn-io/GPU

Node status

This tab displays the node managers in a pictorial representation. It displays details such as the number of active nodes, number of unhealthy nodes, decommissioning nodes, and the number of decommissioned node managers.


Nodes heatmap chart

This tab graphically displays nodes on the basis of their usage of memory. You can enter host or rack details in the search bar to filter nodes.

Viewing queues and queue details

The Queues page is scheduler dependent and displays details of YARN queues. You can either view queues from all the partitions or filter to view queues of a partition.


Capacity Scheduler



View queue details

In the capacity scheduler view, click the circle that represents a particular queue. The right column of the page gets updated with details of that queue.


The following example shows the details of a queue:



You can double-click the queue to view its details on a separate page. You can also view details of any application submitted to that queue by clicking its corresponding Application ID.

Viewing all applications

You can search for applications and view their details using the YARN Web User Interface.



The screenshot shows the Cloudera Applications page. At the top, there are tabs for Cluster Overview, Queues, Applications (selected), Services, Flow Activity, Nodes, and Tools. Below the tabs, there's a search bar with a dropdown menu set to 'Regex'. The main content area displays a table of applications. The table has columns for Application ID, Application Type, Application Name, User, State, Queue, Progress, Start Time, Elapsed Time, Finished Time, Priority, and %Cluster. Three applications are listed, all in a 'Finished' state with 100% progress.

Application ID	Application Type	Application Name	User	State	Queue	Progress	Start Time	Elapsed Time	Finished Time	Priority	%Cluster
application_1521192804348_0003	MAPREDUCE	word count	ambari-qa	Finished	default	100%	2018/03/16 15:08:03	23s 473ms	2018/03/16 15:08:26	0	0
application_1521192804348_0002	YARN	DistributedShell	ambari-qa	Finished	default	100%	2018/03/16 15:08:41	7s 789ms	2018/03/16 15:08:49	0	0
application_1521192804348_0001	TEZ	OrderedWordCount	ambari-qa	Finished	default	100%	2018/03/16 15:08:57	19s 51ms	2018/03/16 15:09:17	0	0

The Applications page displays details of the YARN applications in a tabular form.

- **Application ID:** The identifier for the application.
- **Application Type:** Specifies the application type for Mapreduce, YARN, TEZ, or other applications.
- **Application Name:** Provides the name of the application
- **User:** The name of the user who is the owner of the application.
- **State:** The running state of the application.
- **Queue:** Specifies the name of the queue to which the application belongs.
- **Progress:** The progress of the application in a percentage display.
- **Start Time:** The time when an application run started.
- **Elapsed Time:** The time taken for the application run.
- **Finished Time:** The time of completion of the application run.
- **Priority:** The priority of running the application.
- **%Cluster:** The percentage of cluster resources used by the application run.

Searching applications

The Applications page displays the list of YARN applications in a tabular form. You can apply search filters on this list to show only those applications that match the search criteria.

About this task

You can specify the search criteria either as regular expressions or SQL statements.

Procedure

1. On the Applications page, select either Regex or SQL from the drop-down list depending on the type of search you want to perform.
2. In the Search box, specify the search criteria.

Search Criteria	Description
Regex	<p>Lists the applications whose details match the search criterion specified as a regular expression.</p> <p>For example, if you want to view application runs longer than an hour, mention the regular expression <code>^hso</code> that the YARN UI shows only those applications that mention the Elapsed Time value in terms of hours, minutes, and seconds.</p>
SQL	<p>Lists the applications whose details match the search criterion specified as a SQL statement.</p> <p>For example, if you want to view all the applications submitted by the user yarn, specify <code>"User"='yarn'</code> as the search criterion.</p>

- Click Search to view details of the applications matching the search criteria.

**Note:**

- Apart from specifying search criteria to filter the list of applications, you can also select or clear the State and Queue check-boxes to view a specific set of applications depending on your requirements.
- You can sort the application entries in ascending or descending order by clicking the corresponding arrow next to any column header in the table.

Viewing application details

Clicking a YARN application on the Applications page displays its details.

You can view the following details for the selected application:

- Application Attempts
- Resource Usage
- Diagnostics
- Logs

The screenshot shows the 'word count' application details page. At the top, it indicates the application is 'SUCCEEDED' with ID 'application_1527837750131_0002'. Below this, it shows the state 'Finished' and the user 'ambari-qa' with a completion time of '2018/06/01 12:56:35'. On the right, there are settings for 'default', 'Priority 0', and a 'History' link. A navigation bar contains tabs for 'Attempts List', 'Resource Usage', 'Diagnostics', and 'Logs'. The 'Attempts List' tab is active, showing a 'Graph View' and a 'Grid View'. The 'Grid View' displays a table of application attempts. The first attempt, 'appattempt_1527837750131_0002_000001', is highlighted. To the right of the table, a detailed view for this attempt is shown, including fields for Application Attempt Id, Started Time, Finished Time, Elapsed Time, AM Container Id, AM Node Id, AM Node Web UI, and a Log link.

Application Attempts	
attempt_1	

appattempt_1527837750131_0002_000001	
Application Attempt Id	appattempt_1527837750131_0002_000001
Started Time	2018/06/01 12:56:00
Finished Time	2018/06/01 12:56:35
Elapsed Time	35 Secs
AM Container Id	container_1527837750131_0002_01_000...
AM Node Id	c7403.ambari.apache.org:45454
AM Node Web UI	c7403.ambari.apache.org:8042
Log	Link

Application Attempts

You can view the attempts in Graph View and Grid View.

Graph View

A graph displays the start time and finish time of the attempt. You can also find the details of the attempt such as application attempt ID, started time, finished time, elapsed time, AM Container ID, and AM Node ID in the form of a table. You can access the Node UI using the AM Node Web UI. You can also view the log by clicking on the log link.

Grid View

A table displays the details of the application attempts. You can find the details of the attempt such as application attempt ID, started time, finished time, elapsed time, AM Container ID, and AM Node ID. You can access the Node UI using the AM Node Web UI. You can also view the log by clicking on the log link.

Resource Usage

This tab displays the resources used by the application attempts.

Diagnostics

Use this tab to view the diagnostics details of the application attempts. You can view any outstanding resource requests, and scheduling information.

Logs

Use this tab to view logs specific to containers. Select an attempt from the dropdown list and select the specific container to view the desired logs.

UI Tools

You can view the YARN configuration and YARN Daemon logs on the Tools page.

YARN Configuration

You can see the values of the properties defined in the following configuration files:

1. Core Configuration: Details of properties defined in the core-default.xml file.
2. YARN Configuration: Details of properties defined in the yarn-default.xml file.
3. MapReduce Configuration: Details of the properties defined in the mapred-default.xml file.

YARN Daemon Logs

You can view the list of log files. Click on a log file to view its contents in another tab of your browser.

Using the YARN CLI to viewlogs for applications

Use the YARN CLI to view logs for running application. Configure the log aggregation to aggregate and write out logs for all containers belonging to a single Application grouped by NodeManagers to single log files at a configured location in the file system.

You can access container log files using the YARN ResourceManager web UI, but more options are available when you use the yarn logs CLI command.

View all log files for an application

Use the following command format to view all logs for an application:

```
yarn logs -applicationId <Application ID>
```

View a specific log type for an application

Use the following command format to view all logs of a particular type for an application:

```
yarn logs -applicationId <Application ID> -log_files <log_file_type>
```

For example, to view only the stderr error logs:

```
yarn logs -applicationId <Application ID> -log_files stderr
```

The -logFiles option also supports Java regular expressions, so the following format would return all types of log files:

```
yarn logs -applicationId <Application ID> -log_files .*
```

View ApplicationMaster log files

Use the following command format to view all ApplicationMaster container log files for an application:

```
yarn logs -applicationId <Application ID> -am ALL
```

Use the following command format to view only the first ApplicationMaster container log files:

```
yarn logs -applicationId <Application ID> -am 1
```

List container IDs

Use the following command format to list all container IDs for an application:

```
yarn logs -applicationId <Application ID> -show_application_log_info
```

View log files for one container

Once you have the container IDs, you can use the following command format to list the log files for a particular container:

```
yarn logs -applicationId <Application ID> -containerId <Container ID>
```

Show container log file information

Use the following command format to list all of the container log file names (types) for an application:

```
yarn logs -applicationId <Application ID> -show_container_log_info
```

You can then use the `-logFiles` option to view a particular log type.

View a portion of the log files for one container

For large container log files, you can use the following command format to list only a portion of the log files for a particular container:

```
yarn logs -applicationId <Application ID> -containerId <Container ID> -size  
<bytes>
```

To view the first 1000 bytes:

```
yarn logs -applicationId <Application ID> -containerId <Container ID> -size  
1000
```

To view the last 1000 bytes:

```
yarn logs -applicationId <Application ID> -containerId <Container ID> -size  
-1000
```

Download logs for an application

Use the following command format to download logs to a local folder:

```
yarn logs -applicationId <Application ID> -out <path_to_local_folder>
```

The container log files are organized in parent folders labeled with the applicable node ID.

Display help for YARN logs

To display Help for yarn logs, run the following command:

```
yarn logs -help
```