

Monitoring

Date published: 2019-12-17

Date modified: 2022-05-05



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

Flink Dashboard.....	4
Streams Messaging Manager integration.....	4
Enabling Flink DEBUG logging.....	4

Flink Dashboard

The Flink Dashboard is a built-in monitoring interface for Flink applications in Cloudera Streaming Analytics. You can monitor your running, completed and stopped Flink jobs on the dashboard. You reach the Flink Dashboard through Cloudera Manager.

After deploying Flink and the required components, you can configure and monitor each component individually, or the whole cluster with Cloudera Manager. For the general use of Cloudera Manager, see the [Cloudera Manager documentation](#).

The Flink Dashboard acts as a single UI for monitoring all the jobs running on the YARN cluster. It shows all the running, failed, and finished jobs.



Note: The Flink Dashboard is an updated version of the Flink HistoryServer.

You can also use the dashboard to navigate between the different Flink clusters from a central place.

The screenshot displays the Flink Dashboard interface. On the left is a dark sidebar with navigation links: Overview, Jobs, Running Jobs, and Completed Jobs. The main content area has a top status bar showing 'Version: 1.10.0-csa1.2.0.0-SNAPSHOT', 'Commit: 44ac4bd @ 20.04.2020 @ 16:54:12 CEST', and a 'Message' icon. Below this, there are two summary cards: 'Available Task Slots' showing '2' (Total Task Slots: 16, Task Managers: 5) and 'Running Jobs' showing '4' (Finished: 0, Canceled: 4, Failed: 3). The 'Running Job List' table shows four jobs: 'default: select * from KioskIntrusionSFO', 'SFO Predictive Maintenance', 'Kiosk Outlier Classifier', and 'Kiosk Alerts', all with status 'RUNNING'. The 'Completed Job List' table shows two jobs: 'TXL Predictive Maintenance' with status 'FINISHED' and 'default: select * from KioskIntrusionSFO' with status 'CANCELED'.

Job Name	Start Time	Duration	End Time	Tasks	Status	Cluster
default: select * from KioskIntrusionSFO	2020-04-20 22:44:26	2m 20s	-	1 1	RUNNING	application_1587365700987_0019
SFO Predictive Maintenance	2020-04-20 22:24:13	22m 32s	-	16 16	RUNNING	application_1587365700987_0015
Kiosk Outlier Classifier	2020-04-20 22:22:00	24m 45s	-	16 16	RUNNING	application_1587365700987_0014
Kiosk Alerts	2020-04-20 22:20:03	26m 42s	-	20 20	RUNNING	application_1587365700987_0013

Job Name	Start Time	Duration	End Time	Tasks	Status	Cluster
TXL Predictive Maintenance	2020-04-20 22:25:19	19m 46s	2020-04-20 22:45:06	12 1	FINISHED	
default: select * from KioskIntrusionSFO	2020-04-20 22:43:07	1m 1s	2020-04-20 22:44:09	1 1	CANCELED	application_1587365700987_0019

Streams Messaging Manager integration

You can use the Streams Messaging Manager (SMM) UI to monitor end-to-end latency of your Flink application when using Kafka as a datastream connector.

End-to-end latency throughout the pipeline can be monitored using SMM. To use SMM with Flink, interceptors need to be enabled for Kafka in the Flink connectors.

For more information about enabling interceptors, see the [SMM documentation](#).

Enabling Flink DEBUG logging

You can review the log text files of the Flink jobs when an error is detected during the processes. When you set the log level of Flink to DEBUG, you can easily trace the log file for errors.

About this task

A log file is created for every Flink process that contains messages for the different events happening in the given process. You can use these log files to solve the errors and problems that can occur during Flink processes. You can access the Flink logs using the Flink Dashboard.

Procedure

1. Navigate to the **Configuration** page in Cloudera Manager.
 - a) Go to your cluster in Cloudera Manager.
 - b) Select Flink from the list of services.
 - c) Click Configuration.
2. Search for Flink Client Advanced Configuration Snippet (Safety Valve) for flink-conf/log4j.properties configuration.
3. Add the following parameters to the Safety Valve:

```
logger.flink.name = org.apache.flink
logger.flink.level = DEBUG
```

4. Click Save Changes.
5. Restart the Flink service with Action > Restart .
6. Access the YARN Resource Manager user interface to stop the YARN job of the Flink application.
 - a) Go back to your cluster in Cloudera Manager.
 - b) Select YARN from the list of Services.
 - c) Select Applications.

You are redirected to the Resource Manager page, and the running Flink applications are displayed.
7. Select the application you need to stop.
8. Click Settings.
9. Select Kill application.
10. Navigate to **Flink Dashboard** and review the log level for the running job.
 - a) Go back to Cloudera Manager.
 - b) Select Flink from the list of services.
 - c) Click Flink Dashboard.
11. Select Task Managers from the main menu.
12. Select the previously submitted job.
13. Click Logs.