

Cloudera Runtime 7.1.1

Configuring Apache Atlas

Date published: 2020-05-30

Date modified:

CLOUdera

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

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

Configuring and Monitoring Atlas.....	4
Showing Atlas Server status.....	4
Accessing Atlas logs.....	5

Configuring and Monitoring Atlas


Cloudera Manager manages Atlas as a service, including monitoring its health and collecting its logs. Use Cloudera Manager to configure additional health checks, set up extractors, and set other Atlas properties.

Typically, you would need to make changes to Atlas configuration for the following reasons:

- Expanding resources for the Atlas server. For example in CDP Data Center, the initial memory allocation for Atlas is 2 MB.
- Changing Atlas security settings to add security levels or update parameters for Knox, LDAP, TLS, and Kerberos.
- Updating Atlas for other environment changes such as log directory locations.

To start or stop metadata collection from a service, make the configuration change in the properties for that service.

Atlas collects metadata for services in a given cluster managed by Cloudera Manager. If Cloudera Manager manages more than one cluster, you may have multiple instances of Atlas, one in each cluster. When you check the health of Atlas or configure its properties, make sure you are working in the correct cluster.



Here are some of the Atlas operations you can perform from Cloudera Manager:

Related Information

[Changing the Configuration of a Service or Role Instance](#)

Showing Atlas Server status

Cloudera Manager shows the status of the Atlas Server role running in a given cluster. To understand if the Atlas UI is up and working for users, check the health of the other services on which Atlas depends: HBase, Kafka, and Solr.

The high-level status that Cloudera Manager shows for Atlas indicates the health of the Atlas Service. If you are experiencing problems with the Atlas UI or with metadata not showing up as expected, check to make sure the dependent services are also healthy.

To view the high-level status of Atlas and its dependent services:


1. Open the Cloudera Manager Admin Console.
 - From CDP Public Cloud, [access Cloudera Manager](#) through your Data Lake.
 - From CDP Data Center, [access Cloudera Manager](#) directly in your browser.

- Choose Clusters <cluster name>. The cluster detail page shows the high-level status of each service. Check Atlas, HBase, Kafka, and Solr.

The screenshot shows the Cloudera Manager interface for 'Cluster 1'. The left sidebar contains navigation options like Clusters, Hosts, Diagnostics, Audits, Charts, Replication, and Administration. The main panel displays the 'Status' of various services. A list of services is shown with their respective status icons (green checkmarks) and names: 4 Hosts, ATLAS-1, HBASE-1, HDFS-1, HUE-1, IMPALA-1, KAFKA-1, OOZIE-1, OZONE-1, RANGER-1, and SOLR-1. A bracket on the right side of the list groups ATLAS-1, HBASE-1, KAFKA-1, and SOLR-1, with the label 'Atlas and dependent services'.

Understanding Atlas status

The Atlas status bubbles up to a single health indicator, which appears in many places including the service name



To see more detail of the Atlas status:

- Click the links in the top-level status to drill into the recent health tests and events.
- Drill into specific health test to open the detail page for the test.

If the Atlas status is anything but a green check for "healthy", you can use the content of the page to identify the problem. Check the following:

- If Atlas is stopped, was it stopped deliberately or did it shutdown on its own? To restart the service, see [Restarting a Cloudera Runtime Service](#). To review what may have caused Atlas to shutdown, click the status icon to open a more detailed list of recent events.
- If Atlas shows a red exclamation point for "bad" health, check to see if one of the services it depends upon is stopped: HBase, Kafka, and Solr all run in the same cluster to support Atlas.

Related Information

[Cloudera Manager Health Tests](#)

[Viewing Service Status in Cloudera Manager](#)

Accessing Atlas logs

You can view and download Atlas logs through Cloudera Manager.

To access the Atlas Server log:

1. Open the Cloudera Manager Admin Console.
 - From CDP Public Cloud, [access Cloudera Manager](#) through your Data Lake.
 - From CDP Data Center, [access Cloudera Manager](#) directly in your browser.
2. Choose Clusters<*cluster name*>.
3. Open the Atlas service page by clicking the Atlas Server instance, typically "ATLAS-1".
4. In the Atlas Server detail page, click the Atlas Service role link.
5. In the Log Files menu, select Role Log File.

This option opens the Atlas application.log file, showing the latest contents of the log. You can filter the entries to a particular date range and save the log locally (Download button).

