

Cloudera Observability On-Premises Release Notes

Date published: 2024-01-31

Date modified: 2024-08-14

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

Release Summary.....4

What's new in Cloudera Observability On-Premises 3.5.2..... 4

Known Issues.....5

Release Summary

Cloudera Observability On-Premises is a single pane of glass observability solution, continually discovering and collecting performance telemetry across data, applications, and infrastructure components running in your on-premises deployment. It enables you to interactively explore and understand your existing environments, workloads, clusters, and resources running in your on-premises environment.

With advanced intelligent analytics and correlations, it provides insights and easy to follow recommendations that reduce time to resolution of complex issues, help manage and optimize costs, and improve performance. Cloudera Observability On-Premises also supports better financial governance by tracking and reporting on the costs associated with your business' cost centers.

Cloudera Observability On-Premises helps Administrators and Developers to:

- Watch and protect against budget overruns with its financial governance capabilities, allowing you to define cost centers and chargeback reports.
- Keep workloads healthy with active system monitoring, so you not only know what's going on right now, you'll be comparing to previous trends and historical analysis, to predict issues before they happen, receive alerts to take actions, and get automatic mitigations when possible.
- Improve performance with automations that help things run as best as they can, helping you optimize resource utilization and improve performance. With recommendations, you'll get insights into how to tune, and with custom automatic actions, CDP can be configured to auto-tune, your way.
- Maintain end to end health by identifying and eliminating bottlenecks that impact performance, while also ensuring your entire system, from infrastructure to platform, and workload, is healthy and optimized.
- Get actionable insights through self-service analytics, putting easy to use visualizations into everyone's hands.

Cloudera Observability On-Premises collects and visualizes a wide range of metrics and health tests, enabling you to do the following:

- Gain insights on current and completed workload jobs and queries, resource consumption, and system performance from a wide range of metrics.
- Identify bottlenecks, performance, and resource health issues from a wide range of health tests.
- Address performance issues with performance tuning and prescriptive guidance and recommendations.
- Gain visibility into the workload resource costs of your environment's infrastructure with the Financial Governance Chargeback feature.
- Define workload thresholds and consumption rules, create actions and alerts, and securely control user access, with the Workload Views and Access Management features.

In this release, Cloudera delivers observability covering Hive, Impala, MapReduce, Oozie, and Spark.

Related Information

[Triggering actions across jobs and queries](#)

[Analyzing your environment costs with Cloudera Observability On-Premises cost centers](#)

[Classifying workloads for analysis with Workload Views](#)

[Assigning access roles in Cloudera Observability On-Premises](#)

[Hive, MapReduce, Oozie, and Spark Health Checks](#)

[Impala Health Checks](#)

What's new in Cloudera Observability On-Premises 3.5.2

Review the new features and functionality improvements in this release of Cloudera Observability On-Premises.

Migrate from Cloudera Workload XM On-Prem 2.3.0 to Cloudera Observability On-Premises 3.5.2

Cloudera now supports migration from the Cloudera Workload XM On-Prem 2.3.0 version to the Cloudera Observability On-Premises 3.5.2 version.

You can work closely with the Cloudera Professional Services team, and they will help you plan your migration.

Alternatively, you can contact your Cloudera account manager or create a support case on the Cloudera support portal.

Workload alerts

You can now enable or disable email alert notifications when the workload threshold or the number of failed jobs or queries is exceeded. When enabled, a maximum of one email for each calendar day is sent to the specified email address notifying them of the exceeded workload threshold.

For more information, see [Triggering email alerts for your workload views](#).

Job exports and downloading chargeback reports

A new Export option has been added to the Jobs or Queries page, enabling you to save job or query information as a CSV file on your system. You can use this report for further analysis using other tools or for printing and sharing with others.

For more information, see [Exporting a report about your workload jobs and queries](#).

New chart widgets on the Cluster Summary page

New chart widgets have been added to the top section of the Cluster summary page to display Total, Failed, and Slow queries. These charts also indicate the engines executing these queries. This data helps to gain insights into the health and performance of workload jobs and queries.

Disable TLS/SSL for Phoenix, Kafka, or Impala services

Cloudera Observability On-Premises now supports enhanced security features for clusters by default. Specifying the `ssl.truststore.path` properties ensures that TLS/SSL is enabled for supporting services such as Kafka, Impala, and Phoenix. However, in certain cases, you might need to disable TLS. You can now selectively disable TLS for supporting services.

For more information, see [Disabling TLS/SSL for Phoenix, Kafka, or Impala services](#).

Support Cloudera Observability On-premises with Apache Ranger

If you have installed Apache Ranger for Cloudera Observability On-Premises, Cloudera suggests you set the Apache Ranger properties to specify super users for the supported services.

For more information, see [Supporting Cloudera Observability On-Premises with Apache Ranger](#).

Known Issues

Current known issues and limitations in Cloudera Observability On-Premises.

Impala does not support super user configuration for the observability user for Apache Ranger-enabled cluster

The observability user requires full privileges on the Observability cluster. Required services such as Kafka, HDFS, HBase, and Hive support super user setup by specifying the `ranger.plugin.[service].super.users` property to observability. However, this super user setup is not supported for Impala.

Manually add a new user named observability in Apache Ranger and assign full privileges.

- For information, see *Adding a user* in CDP Private Cloud Base documentation.
- For information on granting user access using Apache Ranger, see *Impala Authorization* in CDP Private Cloud Data Warehouse Runtime documentation.

/cloudera-sigma-olap directory consumes significant storage space in HDFS



Note: The /cloudera-sigma-olap directory stores all the aggregated data. Deleting anything from this directory leads to data loss in dashboards and affects chargeback costs. Consider cleaning data older than six months to avoid unnecessary loss.

The Cloudera Observability On-Premises 3.5.2 version enables weekly execution of the purger by default. However, periodic clean-up of /cloudera-sigma-olap is not currently included in this purger.

Run the command to identify which specific directory consumes the most space.

```
hdfs dfs -du -h /cloudera-sigma-olap/
0          0          /cloudera-sigma-olap/auto_action_audit
2.0 G     6.0 G     /cloudera-sigma-olap/cluster_events
82.1 G    246.4 G   /cloudera-sigma-olap/cluster_metrics
0         0         /cloudera-sigma-olap/hive_on_mr_table
37.7 M    113.2 M   /cloudera-sigma-olap/hms_partition
13.2 M    39.6 M    /cloudera-sigma-olap/hms_partition_json_schem
a_v1
93.4 M    280.3 M   /cloudera-sigma-olap/hms_table
7.8 M     23.4 M    /cloudera-sigma-olap/hms_table_json_schema_v1
798.1 K   2.3 M     /cloudera-sigma-olap/impala_query
52.7 M    158.1 M   /cloudera-sigma-olap/pse_root
8.1 K     24.3 K    /cloudera-sigma-olap/schema
0         0         /cloudera-sigma-olap/yarn_app_metrics
```

Consider a scenario where you want to clean up the cluster_metrics directory.

1. Navigate into the folder structure to review the directories for each date.

```
/cloudera-sigma-olap/cluster_metrics/accountid=accountid/clu
sterid=/dt=2024-07-07
```

2. Run the following commands to manually clean up the directory:

Clean up for a year:

```
hdfs dfs -rmr /cloudera-sigma-olap/cluster_metrics/accountid
=accountid/clusterid=/dt=2023-*
```

Clean up for a month

```
hdfs dfs -rmr /cloudera-sigma-olap/cluster_metrics/accountid
=accountid/clusterid=/dt=2023-07-*
```

Exporting of Impala queries fail for Telemetry Publisher with Cloudera Manager 7.11.3

Telemetry Publisher for Impala queries does not work with Cloudera Manager 7.11.3

Upgrade Cloudera Manager from 7.11.3 to 7.11.3 cumulative hotfix 6 (CHF6) version to successfully export Impala queries.

Auto Action trigger for Impala Engine

Impala Auto Action triggers do not work for the Kerberos-enabled Private Cloud base cluster running on Cloudera Manager 7.9.5 and 7.11.3.

None

Full log link fails to open the log details page on Mozilla Firefox

On the Spark Job Details page, clicking the Full Log link does not open the Log details page. This issue occurs on Mozilla Firefox.

Use Google Chrome or Internet Explorer.

Related Information

[Adding a user](#)

[Impala Authorization](#)