# Cloudera Observability On-Premises Release Notes

**Date published: 2024-01-31**
**Date modified: 2024-08-14**

## CLOUDERA

# Legal Notice

# Contents

# 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 3.5.3

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

### Resource efficiency analysis

You can now identify inefficient jobs or queries that over-allocate resources such as viewing the top queries or jobs by CPU wastage and memory wastage or viewing the CPU and memory consumption for the selected query. Using the resource efficiency analysis feature, you can ensure resources are allocated more efficiently, reducing waste and improving overall system performance.

You can use the query cost analysis to gain insights into costs associated with different resources used by the query.

**Note:** Currently, the Resource Efficiency Analysis feature is available for Hive on Tez, while the Query Cost Analysis is supported for both Hive and Impala.

For more information, see Query and job resource optimization using resource efficiency analysis.

### Support for purging data from /cloudera-sigma-olap directory

The weekly execution of the purger now includes periodic clean-up of the /cloudera-sigma-olap directory using the purge event parameter.

For information on purge event parameters, see Cloudera Observability On-Premises purge event parameters.

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

**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.

Upgrade Cloudera Manager to 7.11.3 cumulative hotfix 9 (CHF9) version.

**Telemetry publisher test altus connection fails for Cloudera Manager 7.11.3 hotfix (CHF6, 7, and 8) versions**

Test connection fails with the following error:

```
Exception in thread "main" java.lang.NoSuchMethodError: 'com.goo
gle.common.collect.ImmutableSet com.google.common.collect.Immuta
bleSet.copyOf(java.util.Collection)'
  at com.cloudera.cdp.http.HttpCodesRetryChecker.<init>(HttpCode
sRetryChecker.java:57)
```

```
   at com.cloudera.cdp.client.CdpClientConfigurationBuilder.<init>(
CdpClientConfigurationBuilder.java:53)
   at com.cloudera.cdp.client.CdpClientConfigurationBuilder.default
Builder(CdpClientConfigurationBuilder.java:400)
   at com.cloudera.cdx.client.TestDatabusConnection.main(TestDatab
usConnection.java:55)
```

This issue only affects the test connection method.

Upgrade Cloudera Manager to 7.11.3 cumulative hotfix 9 (CHF9) version, and then start Telemetry Publisher.

**Related Information**
Adding a user
Impala Authorization

# Fixed Issues

Review the list of issues that are resolved in the Cloudera Observability On-Premises 3.5.3 release.

**OBS-2173: cloudera-sigma-olap directory consumes significant storage space in HDFS**

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.

This issue is fixed now by providing a support for purging data from the /cloudera-sigma-olap directory.

**OBS-4600: 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.

This behavior is resolved now. You can now view the Log details page.

**OBS-5246: Workload alerts missing or Kafka connection error in Admin API logs with TLS enabled**

Admin API server starts, however, the following warning is displayed in the role logs:

```
WARN org.apache.kafka.clients.NetworkClient: [kafka-producer-net
work-thread | producer-1] [Producer clientId=producer-1] Connect
ion to node -1 (example-1.esc12345.root.comops.site/192.0.2.0:24
) terminated during authentication. This may happen due to any o
f the following reasons: (1) Authentication failed due to invali
d credentials with brokers older than 1.0.0, (2) Firewall blocki
ng Kafka TLS traffic (eg it may only allow HTTPS traffic), (3) T
ransient network issue.
WARN org.apache.kafka.clients.NetworkClient: [kafka-producer-net
work-thread | producer-1] [Producer clientId=producer-1] Bootstr
ap broker example-1.esc12345.root.comops.site:24 (id: -1 rack: n
ull) disconnected
```

This issue occurs when SSL Truststore path (JKS) and password are not configured for the Admin API Server component.

This issue is fixed now.