

Workload XM 2.2.2

Workload XM Overview

Date published: 2020-12-04

Date modified: 2022-01-31

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

Understanding Workload XM.....	4
About the Workload XM User Interface.....	4
Using the Workload XM User Interface.....	5
Collecting Workload XM Diagnostic Metrics.....	10
Metric Sources Sent to Workload XM.....	10
Diagnostic Metrics Collection Details.....	11
Redaction Capabilities for Diagnostic Data.....	12

Understanding Workload XM

Workload XM is a Cloudera service that helps you interactively understand your workloads, clusters, and resources. Its wide range of metrics and health tests help you identify and troubleshoot existing and potential problems and the prescriptive guidance and recommendations help you quickly address and optimize those problems. When a workload completes, diagnostic information about the job or query and the cluster that processed them is collected by Telemetry Publisher, a role in the Cloudera Manager Management Service, and sent to Workload XM.

Workload XM enables you to interactively understand your workloads, clusters, and resources, and optimize your systems through:

- A wide range of metrics and health tests that help you identify and troubleshoot issues and potential issues.
- Prescriptive guidance and recommendations that help you quickly address and optimize those problems.
- Performance baselines and historical analysis that help you identify and address performance problems.

In addition, Workload XM also enables you to:

- Visually display your workload cluster's current and historical costs that help you plan and forecast budgets, future workload environments, and justify current user groups and resources.
- Trigger actions in real-time across jobs and queries that help you take steps to alleviate potential problems.
- Enable the daily delivery of your cluster statistics to your email address that help you to track, compare, and monitor without having to log in to the cluster.
- Break down your workload metrics into more meaningful views for your business requirements that help you analyze specific workload criteria. For example, you can analyze how queries that access a particular database or that use a specific resource pool are performing against your SLAs. Or you can examine how all the queries are performing on your cluster that are sent by a specific user.

About the Workload XM User Interface

Provides a brief introduction to the web interface and frequently used interface elements of Workload XM.

The Workload XM UI displays graphical reports of the health, performance, and status of your workloads, workload clusters, and resources. Its dashboard components include statistics, performance, health, and prescription guidance visually displayed in chart widgets or tabular views.

The navigation side-bar enables access to the following Workload XM features and engine services:

- Summary, which opens the Cluster Summary page that visually displays information about your cluster and engine health.
- Workloads, which opens the Workload Views page that enables you to create, view, monitor, and manage views specific to your business requirements.
- Auto Actions, which opens the Auto Actions page that enables you to create, view, monitor, and manage your Auto Actions.
- Chargeback, which opens the Chargeback page that enables you to create, view, monitor, and manage cost centers specifically for your workloads.
- Engines, which lists the Workload XM engine workload services:
 - Impala
 - Spark
 - Hive
 - Oozie
 - MapReduce

When an engine is selected, the engine's chart widgets display information about the workloads serviced by the selected engine, such as which jobs or queries have failed or are slow, their processing time, missed SLAs, user

and pool metrics, and outlier issues. The Impala and Workload View chart widgets also display information about the CPU and Memory usage, the tables used for your Impala queries, and the SQL statement types.

About the Workload XM Chart Widgets

The Workload XM chart widgets enable you to quickly observe real-time and historical patterns, trends, and outliers of your workload data. They provide quick insights into the health and performance of your workloads, clusters, and resources. Where:

- Hovering over an element with your mouse pointer, such as over a time-line or a data point, displays more information about the element underneath.
- Clicking a link within a chart widget or a bar within bar chart widget, such as the Suboptimal bar chart widget, opens the engine's Jobs or Queries page that contains more information in a tabular view for you to investigate further.

About the Jobs and Queries Pages

Depending on the engine chosen, the engine's Jobs or Queries page provides information about each of the jobs or queries that are serviced by the engine. You can filter further by selecting a filter option from one of the filter categories; Pool, User, Status, Health Check, or Duration.

For more information about a specific job's or query's health, execution details, baseline, and trends, open their respective page by clicking its link in the Job or Query column and selecting the tab of interest. As well as providing more details for you to investigate further, these pages also provide prescriptive guidance and recommendations that enable you to address and optimize problems.

Using the Workload XM User Interface

Describes a few frequently used interface elements of Workload XM that help you identify and troubleshoot your workload issues.

The following examples describe a few interface elements of Workload XM that enable you to quickly identify workload problems, health issues, resource contentions, and abnormal or degraded performance problems.



Note: These examples assume that you have logged into the Workload XM UI.

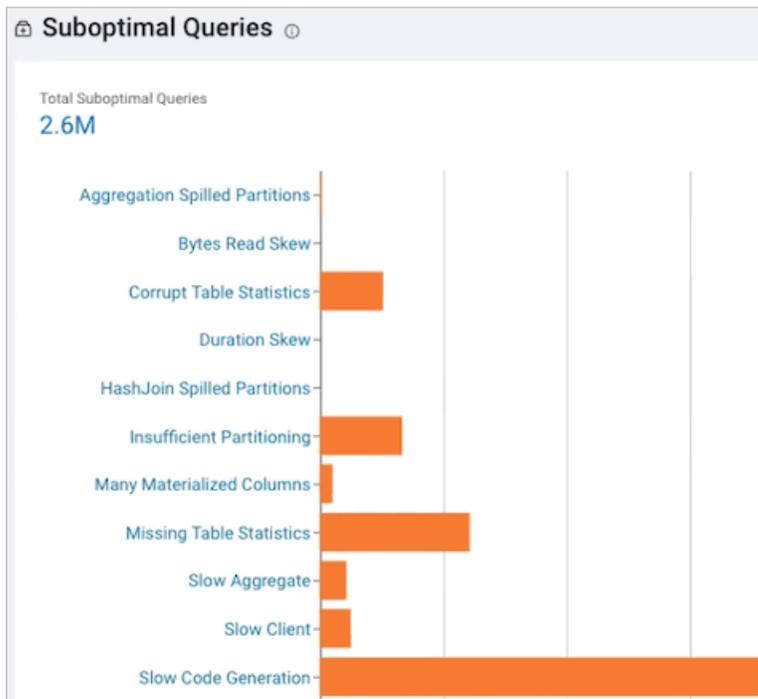
Identify Workload problems and Health Issues

Quickly locate what engines are running on your clusters and what jobs and queries are failing the health tests with the following chart widgets:

- The Usage Analysis chart widget, which displays what engines are running on the cluster, how many jobs or queries are processed by the engine, and how many jobs or queries have failed and missed their SLA.

Location: This chart widget is found on the Cluster Summary page, which is accessed by selecting the workload cluster for analysis and then from the navigation side-bar, selecting Summary.

- The Suboptimal chart widget, which displays the distribution of jobs and the jobs and queries that failed. This chart widget enables you to visually see at a glance what issues are currently impacting your jobs or queries and how they are executing on your cluster:

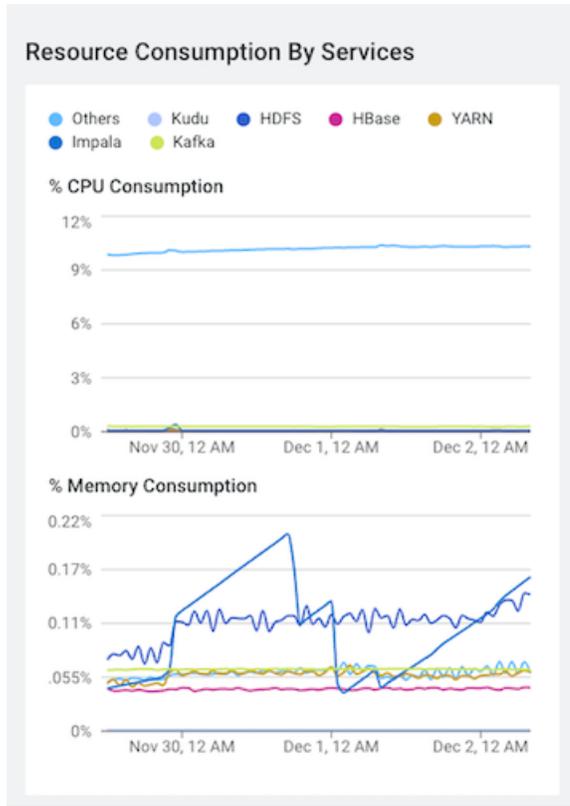


Location: This chart widget is found on the *engine's* page after selecting the workload cluster for analysis and then from the navigation side-bar, selecting an *engine*.

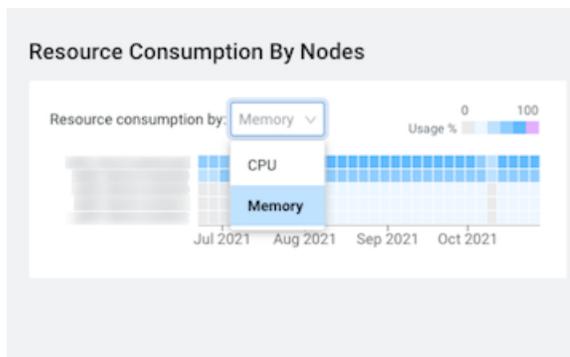
Identify and Address Resource Contentions

Workload XM provides the following chart widgets that help you analyze and identify resource consumption and contention problems:

- Location: The following chart widgets are found on the Cluster Summary page, which are accessed by selecting the workload cluster for analysis and then from the navigation side-bar, selecting Summary:
 - The Resource Consumption By Services chart widget displays the CPU and memory consumption for each service across the time range you selected. Hover your mouse over the time line, to display the amount of CPU or memory, as a percentage, that is consumed by each of the cluster's services.

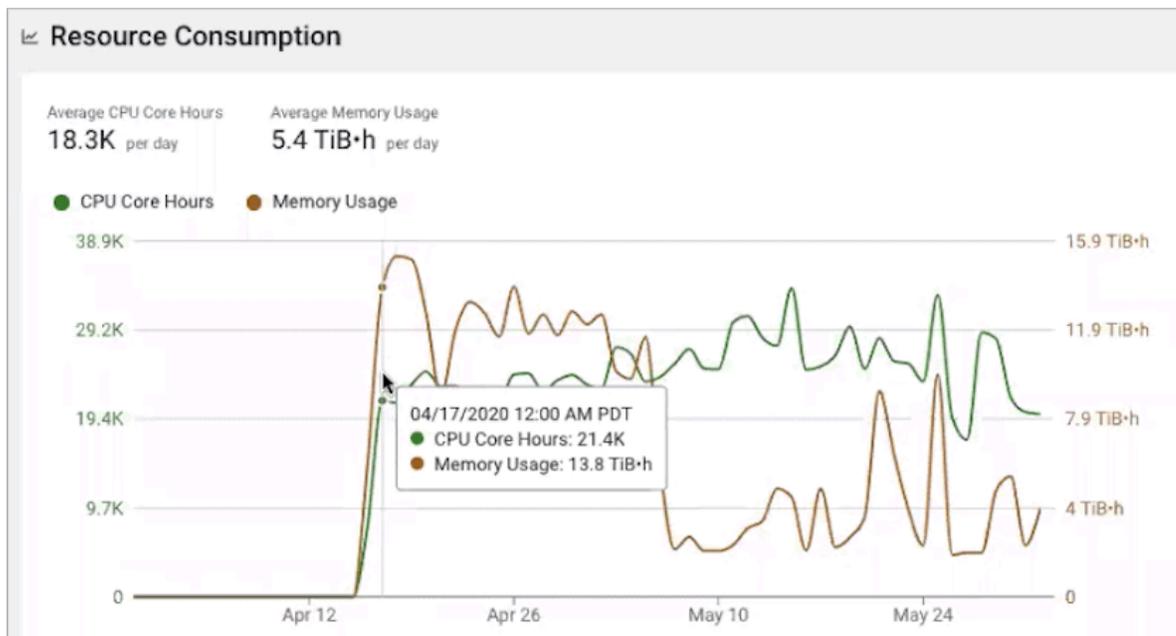


- The Resource Consumption By Nodes chart widget displays the CPU and memory consumption for each node in the cluster. Hover your mouse over the time line, to display the amount of CPU or memory, as a percentage, that is consumed by each node and its services.



- Location: The following chart widgets are found on the *engine's* page, which is accessed by selecting the workload cluster for analysis and then from the navigation side-bar, selecting Impala:
 - The Memory Utilization chart widget, displays the aggregated maximum amount of memory that is used by the queries on any node performing the processing. It helps you identify inefficient queries that are consuming the most amount of memory during processing and if you need to allocate more memory to continue running your queries.

- The Resource Consumption chart widget, displays the concurrent use of CPU and memory consumption for a workload across the time line you selected.



Identify and Address Abnormal or Degraded Performance Problems

Workload XM enables you to identify and address abnormal or degraded performance problems by establishing baselines from health issues that also enable a performance comparison of your workloads. The Workload XM baseline metrics measure the current performance of a job against the average performance of previous runs. They use performance data from 30 of the most recent runs of a job and require a minimum of three runs. The baseline comparisons start with the fourth run of a job.

The following images show some of the baseline and comparison metrics that are provided:

- This image shows a few of the performance metrics listed in the Baseline tab:

Metric	Baseline
Active Tasks	0
Disk bytes Spilled	0 B
Duration	1h 36m
Executor Runtime	21h 36m

- This image shows the comparison between the baseline performance metrics and the current job run:

Metric 	Baseline 	Current Job	
Number of Task attempts	134	1K	+870
Number of Tasks	134	995	+861
Output bytes	728.5 MiB	0 B	-728.5 MiB
Output records	396.6K	0	-396.6K
Shuffle Read bytes	265.5 B	38.5 MiB	+38.5 MiB
Shuffle Read records	4.5	796K	+796K
Shuffle Write bytes	265.5 B	52.6 MiB	+52.6 MiB
Shuffle Write records	4.5	1.2M	+1.2M
Succeeded Task attempts	134	992	+858
Succeeded Tasks	134	992	+858
Total Task duration	24m 34s	2h 38m	+2h 13m

Location: The baseline for a job or a query is found on the Baseline page, which is accessed by selecting the workload cluster for analysis and then from the navigation side-bar, selecting Summary. In the Trend chart widget, select the engine tab for analysis and then, depending on the engine, click Total Jobs or Total Queries. From the Jobs or Queries page, select the job or query of interest and then select the Baseline tab.

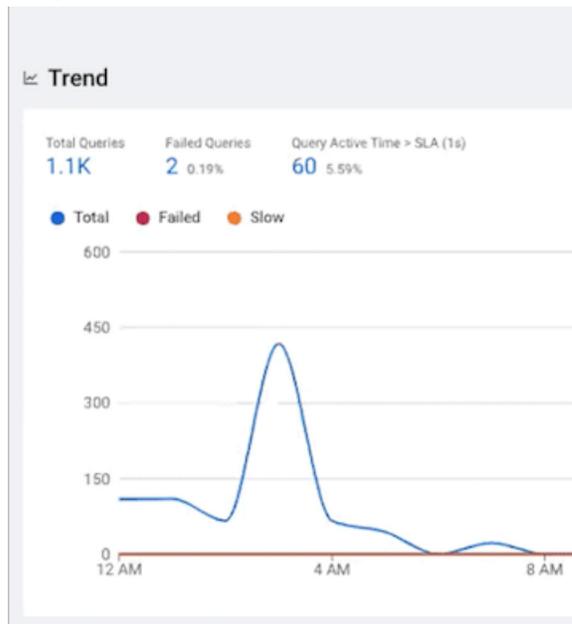


Note: Requires at least four job runs for Workload XM to create the job or query's baseline.

Identify Performance Trends

You can identify trends as well as baselines by analyzing your engine's or cluster's performance trends from the Trends chart widget and the Trend tab. Where:

- The Trends time-series chart widget, displays more detailed metrics about the processed jobs and queries and enables you to view historical trends for analysis when you select a predefined or custom time period from the Range filter list.



Location: This chart widget is found on the *engine's* page, which is accessed by selecting the workload cluster for analysis and then from the navigation side-bar, selecting the *engine* of interest.

- The Trends tab, displays the job or query's instances executed during the selected time range. Depending on the engine, the Trends page displays a job's historical trend from Duration, Data Input, and Data Output histogram charts or lists the runs of the query to show how its performance changes overtime.

Location: This chart widget is found on the Jobs or Query's page, which is accessed by selecting the workload cluster for analysis and then from the navigation side-bar, selecting the *engine* of interest. Depending on the engine chosen, from the Trends chart widget, click Total Jobs or Total Queries and then from the Jobs or Queries page, select the job or query of interest and then select the Trends tab.

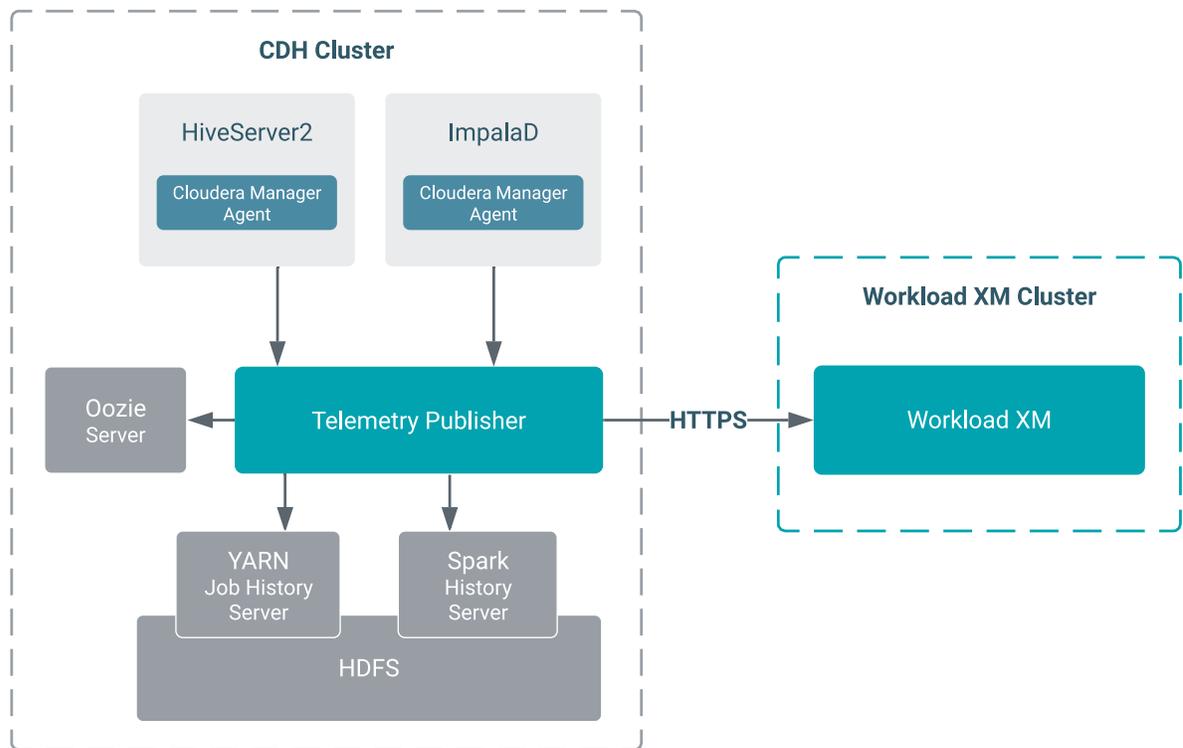
Collecting Workload XM Diagnostic Metrics

When you enable Workload XM, the Cloudera Management Service starts the Telemetry Publisher role. Telemetry Publisher collects and transmits metrics as well as configuration and log files from Impala, Oozie, Hive, YARN, and Spark services for jobs running on your clusters to Workload XM. Telemetry Publisher collects metrics for all the clusters that use Workload XM-enabled environments.

Understanding the sources of information sent to Workload XM and how that data is redacted is described in the following topics.

Metric Sources Sent to Workload XM

Describes the resources from which you can configure Telemetry Publisher to collect diagnostic metrics.



Telemetry Publisher collects and transmits metrics as well as configuration and log files from Impala, Oozie, Hive, YARN, and Spark services for jobs running on your clusters to Workload XM, as shown in the above diagram. The metrics are collected as follows:

- Pull — Telemetry Publisher pulls diagnostic metrics from Oozie, YARN, and Spark periodically (by default, once per minute).
- Push — A Cloudera Manager Agent pushes diagnostic data from Hive and Impala to Telemetry Publisher within 5 seconds after a job finishes.

After the diagnostic data reaches Telemetry Publisher, it is stored temporarily in its data directory and periodically (once per minute) exported to Workload XM.

Diagnostic Metrics Collection Details

Describes the type of data collected by Telemetry Publisher and the Cloudera services that provide the data.

Telemetry Publisher collects and sends the following diagnostic metrics to Workload XM:

- MapReduce Jobs — Telemetry Publisher polls the YARN Job History Server for recently completed MapReduce jobs. For each of these jobs, Telemetry Publisher collects the configuration and jhist file, which is the job history file that contains job and task counters, from HDFS. Telemetry Publisher can be configured to collect MapReduce task logs from HDFS and send them to Workload XM. By default, this log collection is turned off.
- Spark Applications — Telemetry Publisher polls the Spark History Server for recently completed Spark applications. For each of these applications, Telemetry Publisher collects their event log from HDFS. Telemetry

Publisher can be configured to collect the executor logs of Spark applications from HDFS and send them to Workload XM. By default, this data collection is turned off.



Important: Telemetry Publisher only collects Spark application data from Apache Spark 2.x versions. At this time, Apache Spark 3.x versions are not supported.

As CDH version 5.x is packaged with Spark version 1.6, you cannot configure Telemetry Publisher data collection for CDH 5.x clusters unless you are using CDS 2.2 Powered by Apache Spark or later versions with those clusters.

- Oozie Workflows — Telemetry Publisher polls Oozie servers for recently completed Oozie workflows and sends the details to Workload XM.
- Hive Queries — The Cloudera Manager agent periodically searches for query detail files that are generated by HiveServer2 after a query completes and then sends the details from those files to Telemetry Publisher.



Important: Hive query audits must be enabled.

- Impala Queries — A Cloudera Manager agent periodically looks for query profiles of recently completed queries and sends them to Telemetry Publisher.

Redaction Capabilities for Diagnostic Data

Describes the resources that you can configure for redaction. Cloudera recommends enabling redaction even if you are not sending diagnostic data to Telemetry Publisher.

The diagnostic data collected by Telemetry Publisher may contain sensitive information in job configuration or log files. The following lists the data and resources that you can configure for redacting sensitive data before it is sent to Telemetry Publisher:

- Log and query redaction — You can redact information in logs and queries collected by Telemetry Publisher based on filters created with regular expressions.
- MapReduce job properties redaction — You can redact job configuration properties before they are stored in HDFS. Since Telemetry Publisher reads the job configuration files from HDFS, it only fetches redacted configuration information.
- Spark event and executor log redaction — The Spark2 on YARN service on CDH clusters has the `spark.redaction.regex` configuration property that can be used to redact sensitive data from event and executor logs. When this configuration property is enabled, Telemetry Publisher sends only redaction data to Workload XM. By default, this configuration property is enabled, but it can be overridden by using safety valves in Cloudera Manager or in the Spark application itself.