

Workload XM 2.1.3

Workload XM Cluster Optimization

Date published: 2020-12-04

Date modified: 2021-07-22

CLOUDERA

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

Legal Notice

© Cloudera Inc. 2022. 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

Working with Workload XM.....	4
Specifying a time range.....	4
Troubleshooting Abnormal Job Durations.....	4
Troubleshooting Failed Jobs.....	8
Determining the Cause of Slow and Failed Queries.....	10
Classifying Workloads for Analysis with Workload Views.....	13
Working with Auto Generated Workload Views.....	13
Defining Workload Views Manually.....	15
Troubleshooting with the Job Comparison Feature.....	21
Identifying File Size Storage Issues.....	26
Displaying File Size Metadata.....	26
Displaying the Metadata of a Table.....	27

Working with Workload XM

Tasks for identifying and troubleshooting job and query abnormalities and failures, optimizing workloads, and improving job performance with Workload XM.

The content is divided into sections and topics that provide all the tasks associated with identifying and troubleshooting job and query abnormalities and failures, optimizing workloads, and improving job performance with Workload XM.

Specifying a time range

Choose a time period in which your workload results are displayed in Workload XM for analysis and troubleshooting.

About this task

Describes how to specify a time period for displaying current or historical data about your cluster and the jobs performed in that cluster.

By default, Workload XM displays workload data for the last 24 hours. If there is no data available during that time, Workload XM displays the nearest date range that is available.

Procedure

1. Verify that you are logged in to the Cloudera Data Platform.
2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the **Clusters** page, select the cluster required for analysis.
5. From the time-range list in the Cluster Summary page, do one of the following:
 - For a predefined period, select one of the default periods of time that meets your requirements .
 - For an exact date and time range, select **Customize** and then either, enter the date and time range using the YYYY/MM/DD HH:MM:SS format for the beginning and the ending time period, or in the calendar element, select the beginning and ending time period.
6. Click **Ok**, which clears any existing workload data from the chart and table components for the existing period of time.

Results

All charts and tables in Workload XM are updated to reflect the workload data for the chosen time period.

Troubleshooting Abnormal Job Durations

Identify areas of risk from unusual job durations running on your cluster.

About this task

Describes how to locate and troubleshoot abnormal job duration time periods.

Steps with examples are included that explain how to further investigate and troubleshoot the cause of an abnormal job duration time period.

Procedure

1. Verify that you are logged in to the Cloudera Data Platform.

2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Clusters page do one of the following:
 - In the Search field, enter the name of the cluster whose workloads you want to analyze.
 - From the Cluster Name column, locate and click on the name of the cluster whose workloads you want to analyze.
5. From the time-range list in the Cluster Summary page, select a time period that meets your requirements.
6. Display the number of jobs with an abnormal duration that executed within the selected period by clicking the **Abnormal Duration** health check bar in the **Suboptimal Jobs** graph.



Tip: Any jobs that fall outside of the baseline duration are marked as slow. You can hover over the graph to see how many jobs triggered each health check.

Suboptimal Jobs



The Job page opens, listing all the jobs that have triggered the **Abnormal Duration** Health check.

- To specify a specific amount of time in which the job either ran less than or more than the Health check rule, from the **Duration** list, either select a predefined time duration or select **Customize** and enter the minimum or maximum time period.

Jobs

Type All Status All Health Check Duration Duration All Range Default (04/12/2018 - 04/20/20... Clear all

Type	Job	Status	Start Time	Duration	Health Issue	Execution ID
HV	--ins_from...	Succeeded	04/19/2018 1:08 AM PDT	0s - 1h 0m 0s	Abnormal Data Output Abnormal Data Input Abnormal Duration	
HV	insert over...	Succeeded	04/18/2018 10:48 PM P...	12m 10s	Task Output Data Skew Abnormal Data Output Task Input Data Skew Abnormal Data Input Abnormal Duration Data Processing Speed Skew	
HV	insert over...	Succeeded	04/18/2018 10:27 PM P...	16m 34s	Task Output Data Skew Abnormal Data Output Task Input Data Skew Abnormal Data Input	

Choose a duration range from the **Duration** list, or choose **Customize** to enter a custom maximum-minimum duration range.

- To view more details about a job, from the **Job** column, select a job's name and then click the **Health Checks** tab. The Baseline Health checks are displayed.

- To display more information about the job's duration, from the **Baseline** column, select **Duration**.

The following reveals that for this example the job finished much slower than the baseline:

Jobs

Log Analysis

Overview Health Checks Execution Details Baseline Trends 2/8/2019 1:08 PM 10m 44s

Hide Normal Stages

Baseline	Start Time	Job	Duration
Duration	13:08	SP Log Analysis	10m 44s
Input Size	13:09	Job 1	9m 20s
Output Size	13:09	Stage-2	9m 19s

Skew

- Task Duration
- Task Input Data
- Task Output Data
- Shuffle Input
- Data Processing Speed

Resources

- Task Wait Time

Log Analysis

Abnormal Duration

Finished in 10m 44s, slower than the median duration 2m 8s. View all metrics.

Start Time	Execution ID	Duration
2/8/2019 1:08 PM		11m
2/8/2019 1:08 PM		8m
2/8/2019 1:08 PM		6m

Click to further investigate this job.

10. To display more information about the length of time the processing tasks took within a job, from the **Baseline** column, select **Task Duration**.

The following reveals that for this example a particular task took an abnormal amount of time to finish:

The screenshot shows the 'Log Analysis' dashboard. In the 'Baseline' sidebar, 'Task Duration' is selected. The main table shows stages: 'SP Log Analysis' (10m 44s), 'Job 1' (9m 20s), and 'Stage-2' (9m 19s). The right-hand panel displays a 'Task Duration Skew' health check with a chart showing a median of 2s and a typical range of < 1s - 19s. A legend indicates that outliers are tasks that took more than 5m 46s. Below the chart, a table lists '1 Outlier Task':

Task Name	Host	Duration
! Task 160		5m 46s

11. To display more information about the abnormal task, click the abnormal task, which opens the **Task Details** panel.

The following reveals that for this example the garbage collection for Task 160 is taking significantly more time than the average task:

The screenshot shows the 'Task Details' panel for 'Task 000645'. The 'Task Details' table lists various metrics:

Metric	Task	Average
GC time elapsed	21m 53s	22s 206ms
CPU time spent	7h 34m	10m 16s
Successful Attempt Duration	26m 11s	6m 26s
Wait Duration	3s 985ms	10m
Data written (Local)	167.2 MiB	102 MiB
Map output materialized bytes	79.1 MiB	49.2 MiB
Duration	26m 15s	16m 26s
Data read (Local)	342.9 MiB	261.7 MiB
Map output bytes	781.9 MiB	612.2 MiB
Spilled Records	12.2M	9.7M

12. To display more information about the garbage collection for this job example, from the **Baseline** column, select **Task GC Time**.

13. In the **Task GC Time** page, click the **Execution Details** tab and then click one of the MapReduce stages:

The screenshot shows the Cloudera Workload Manager interface. At the top, there are tabs for Overview, Health Checks, Execution Details (selected), Baseline, and Trends. Below the tabs, there are buttons for 'Expand All' and 'Collapse All'. A table lists the stages of the job:

Time	Stage	Duration
20:51	HV insert over...	54m 37s
20:51	MR Stage-1	49m 9s
21:40	MR Stage-2	4m 47s
21:45	MR Stage-2	26s

To the right of the table is a Summary panel with fields for ID and Query, and a 'View Configurations' link.

14. In the **Summary** panel, click **View Configurations** and then locate the configuration for the garbage collection by entering part of the MapReduce memory configuration property name in the Search field:

The screenshot shows the Cloudera Workload Manager interface with the Configurations panel open for 'Stage-1'. The search field is populated with 'map.memory'. The configuration 'mapreduce.map.memory.mb' is displayed with a value of 1024.

The configuration for the garbage collection reveals that the setting is 1024, which might be causing the mapper JVM to have insufficient memory as well as triggering too many garbage collections. Increasing this number will improve cluster performance and remove this task as a potential risk.

Troubleshooting Failed Jobs

Steps for troubleshooting uncompleted workloads running on your cluster.

About this task

Describes how to locate and troubleshoot jobs that have failed to complete.

Steps with examples are included that explain how to further investigate and troubleshoot the root cause of an uncompleted job.

Procedure

1. Verify that you are logged in to the Cloudera Data Platform.
2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Clusters page do one of the following:
 - In the Search field, enter the name of the cluster whose workloads you want to analyze.
 - From the Cluster Name column, locate and click on the name of the cluster whose workloads you want to analyze.

- From the navigation panel, select **Jobs** under Data Engineering.
- From the **Health Check** list in the Jobs page, select **Failed to Finish**, which filters the list to display a list of jobs that did not complete.

Cloudera Workload XM

Jobs

Type All Status All Health Check All Duration All Range Default (04/12/2018 - 04/20/2018)

Type	Job	Status	Duration	User	Health Issue
HV	...	Failed to Finish	8 AM PDT 51s 406ms	...	Data Processing
HV	...	Failed to Finish	8 AM PDT 49s 608ms	...	
MR	...	Succeeded	8 AM PDT 1m 20s	...	
HV	...	Succeeded	7 AM PDT 1m 8s	...	Abnormal Data
MR	...	Succeeded	04/19/2018 2:37 AM PDT 1m 45s	...	Data Processing

Health Check dropdown menu:

- Failed - Any Healthchecks
- Passed All Health Checks
- All Jobs
- Failed to Finish**
- Baseline
- Duration
- Input Size
- Output Size

- To view more details about why the job failed to complete, from the **Job** column, select a job's name and then click the **Health Checks** tab.

The Baseline Health checks are displayed.

- From the **Health Checks** panel, select the **Failed to Finish** health check.

The following reveals that for this example the failure occurred in the **Map Stage** of the job process:

Jobs

Overview Health Checks Execution Details Baseline Trends

4/16/2018 3:39 AM 3m 4s

Health Check	Time	Job	Duration
Failed to Finish	03:39	HV	3m 4s
Baseline	03:39	MR Stage-1	3m 4s
Duration	03:39	Map Stage	2m 52s

Click Map Stage and then click Execution Details.

Map Stage
Failed to Finish [View Execution Details](#)

Operation failed to finish.

This health check determines whether a job succeeded or failed.

Check [Execution Details](#) for more information.

- To display more information about the Map Stage process, click **Map Stage** and then from the Map Stage panel, click **Execution Details**.

10. To see all the failed tasks, in the **Summary** panel, click on the number value in the **Failed** field:

The screenshot shows the 'Jobs' interface for a query named 'insert overwrite table ...'. The 'Execution Details' tab is active. The left sidebar shows a list of tasks: 'HV insert overw...' (3m 4s), 'MR Stage-1' (3m 4s), and 'Map Stage' (2m 52s). The 'Map Stage' is selected. The main panel shows the 'Summary' for 'Stage-1 / Map Stage'. It displays 'Completed' as 11 / 12 and 'Failed' as 1. An orange arrow points to the '1' in the 'Failed' field with the text 'Click the number of Failed tasks.' Below this, 'Average Map Time' is shown as 1m 30s.

11. For each failed attempt, display the error message by selecting each task.

For this example, the following Task KILL is received. Killing attempt! error message reveals that for this example more information is required to answer why a KILL task was received. To further troubleshoot the root cause, understanding what triggered the error is required. To investigate further, open the associated log file by clicking **Logs**.

The screenshot shows the 'Jobs' interface for the same query. The 'Execution Details' tab is active. The left sidebar shows the 'Map Stage' selected. The main panel shows the 'Failed Map Tasks' view for 'Stage-1 / Map Stage'. It displays a table with columns: Attempt, Host, Start Time, and Duration. The table has one row: 'Attempt 0' on host 'atl1c1r1data03.vldb-bo.secureworks.net' starting at 3:39 AM and lasting 2m 48s. Below the table, the 'Error Message' is expanded and highlighted with an orange box, showing 'Task KILL is received. Killing attempt!'. An orange arrow points to the 'Logs' link in the table with the text 'Click to view log file.'

Determining the Cause of Slow and Failed Queries

Identifying the cause of slow query run times and queries that fail to complete.

About this task

Describes how to determine the cause of slow and failed queries.

Steps with examples are included that explain how to further investigate and troubleshoot the cause of a slow and failed query.

Procedure

1. Verify that you are logged in to the Cloudera Data Platform.
2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Clusters page do one of the following:
 - In the Search field, enter the name of the cluster whose workloads you want to analyze.
 - From the Cluster Name column, locate and click on the name of the cluster whose workloads you want to analyze.
5. From the navigation panel under Data Engineering, select **Jobs**.
6. From the **Health Check** list in the Jobs page, select **Task Wait Time**, which filters the list to display a list of jobs with longer than average wait times to execute a process.

The screenshot shows the Cloudera Workload XM interface. On the left is a navigation sidebar with 'Jobs' selected. The main area displays a 'Jobs' table. Above the table, there are filters for 'Type', 'Status', 'Health Check', and 'Duration'. The 'Health Check' dropdown menu is open, showing options like 'Output Size', 'Task Retries', 'Task GC Time', 'Disk Spillage', 'Task Wait Time', and 'Skew'. 'Task Wait Time' is highlighted with an orange box and an orange arrow. The table below shows several job entries, some with 'Health Issue' labels like 'Data Processing Speed Skew' and 'Abnormal Data Input'.

Type	Job	Status	Start Time	Duration	User	Health Issue	Execution ID
HV	insert over...	Succeeded	04/19/2018 2:37 AM PDT	1m 45s			hive_...
HV	insert over...	Succeeded	04/19/2018 2:37 AM PDT	1m 45s		Data Processing Speed Skew	hive_...
MR	aggregate...	Succeeded	04/19/2018 2:37 AM PDT	1m 45s			job_...
HV	insert over...	Succeeded	04/19/2018 2:37 AM PDT	1m 45s		Abnormal Data Input	hive_...
MR	...	Succeeded	04/19/2018 2:37 AM PDT	1m 45s		Data Processing Speed Skew	job_...

7. To view more details, from the **Job** column, select a job's name and then click the **Health Checks** tab. The Baseline Health checks are displayed.

8. From the **Health Checks** panel, select the **Task Wait Time** health check.

The following reveals that for this example the long wait time occurred in the **Map Stage** of the job process due to insufficient resources:

The screenshot displays the Databricks Jobs interface for a job named 'select'. The 'Health Checks' tab is active, showing a list of health checks under 'Resources'. The 'Task Wait Time' health check is selected and highlighted in blue. An orange arrow points to this health check with the text 'Click to view details on tasks with abnormally long wait times.' Below the health check list, a table shows the job's progress: 'HV select' (18m 16s), 'MR Stage-8' (16m 54s), and 'Map Stage' (16m 40s). The 'Map Stage' is highlighted in blue. To the right, a 'Long Task Wait Time' health check details are shown, including a bar chart and a table of 'Top 20 Outlier Tasks'. The chart shows a median wait time of 7m 12s and a typical range of 3m 33s to 10m. The table lists three outlier tasks: 'Task 032985', 'Task 032975', and 'Task 032688', each with a warning icon and a host name.

9. To display more information about the Map Stage tasks that are experiencing longer than average wait times to execute, click one of the tasks listed under **Outlier Tasks**.

The following reveals that for this outlier task example, the **Wait Duration** time is above average, as confirmed by comparing this time with the time taken when the task successfully completes. Where, the successful value

is displayed in the **Successful Attempt Duration** field and is significantly better than the average time. This indicates that insufficient resources are allocated for this job.

The screenshot shows the Cloudera Workload Manager interface. The main view is 'Task Details' for a 'Map Stage' (Task 032953). The interface includes a navigation bar with tabs for Overview, Health Checks, Execution Details, Baseline, and Trends. The 'Task Details' table shows the following data:

Metric	Task	Average
Wait Duration	15m 20s	7m 9s
Duration	16m 5s	8m 8s
RECORDS_OUT_0	22K	12.8K
Input split bytes	551 B	797.73 B
Data written (HDFS)	6.5 MiB	5.2 MiB
Successful Attempt Duration	45s 10ms	59s 287ms
Map input records	6.1M	8M

Classifying Workloads for Analysis with Workload Views

The Workload View feature enables you to analyze workloads with much finer granularity. 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.

Working with Auto Generated Workload Views

Steps for using the Workload XM default workload views.

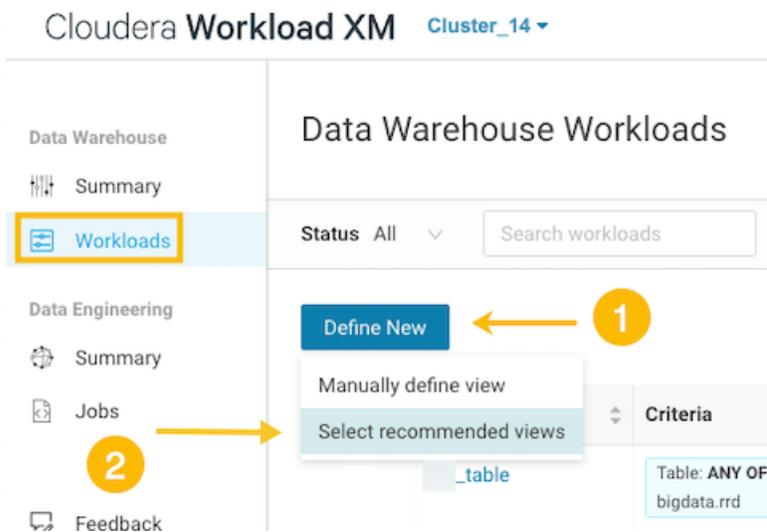
About this task

Describes how to use the workload views that Workload XM automatically generates.

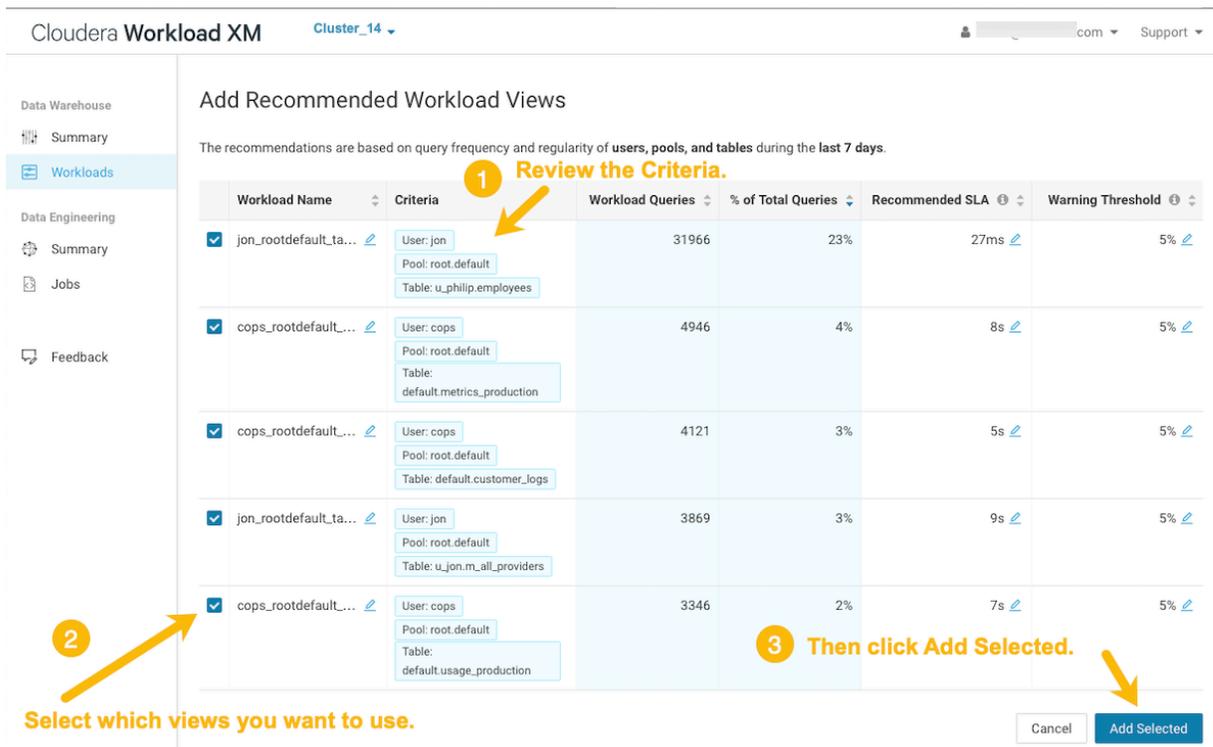
Procedure

1. Verify that you are logged in to the Cloudera Data Platform.
2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Clusters page do one of the following:
 - In the Search field, enter the name of the cluster whose workloads you want to analyze.
 - From the Cluster Name column, locate and click on the name of the cluster whose workloads you want to analyze.
5. From the navigation panel under Data Warehouse, select **Workloads**.

- In the Data Warehouse page, click the **Define New** list and then select **Select recommended views**:

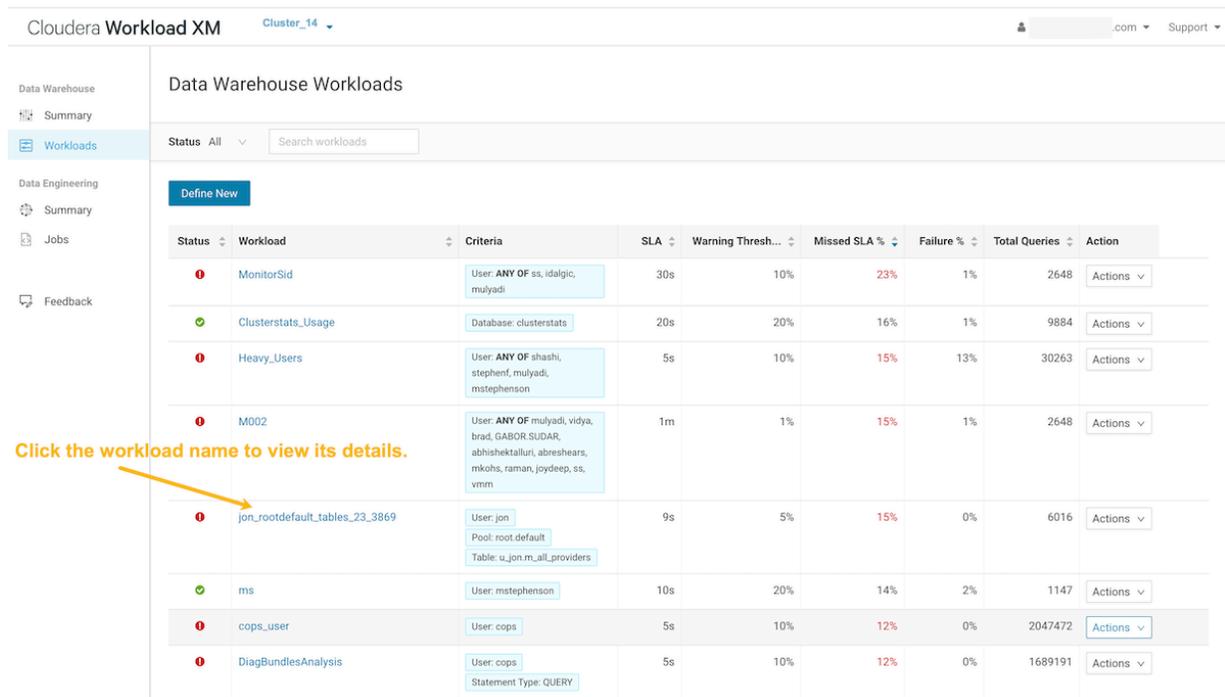


- From the **Criteria** column, examine the criteria that are used to create the workload views, select the workload views required, and then click **Add Selected**:



The workload views you selected are saved and displayed on the Data Warehouse Workloads page.

8. To verify your workload views, from the navigation panel under Data Warehouse, select **Workloads** and then on the Data Warehouse Workload page locate the workload view you just added. When verified, click that workload to view its details:



The screenshot shows the Cloudera Workload XM interface for Cluster_14. The 'Data Warehouse Workloads' page is displayed, featuring a table with columns for Status, Workload, Criteria, SLA, Warning Threshold, Missed SLA %, Failure %, Total Queries, and Action. An orange arrow points to the workload 'jon_rootdefault_tables_23_3869' with the text 'Click the workload name to view its details.'

Status	Workload	Criteria	SLA	Warning Thresh...	Missed SLA %	Failure %	Total Queries	Action
🔴	MonitorSid	User: ANY OF ss, idalgic, mulyadi	30s	10%	23%	1%	2648	Actions
🟢	Clusterstats_Usage	Database: clusterstats	20s	20%	16%	1%	9884	Actions
🔴	Heavy_Users	User: ANY OF shashi, stephent, mulyadi, mstephenson	5s	10%	15%	13%	30263	Actions
🔴	M002	User: ANY OF mulyadi, vidya, brad, GABOR.SUDAR, abhishektalluri, abreshears, mkohs, raman, joydeep, ss, vmm	1m	1%	15%	1%	2648	Actions
🔴	jon_rootdefault_tables_23_3869	User: jon Pool: root.default Table: u_jon_m_all_providers	9s	5%	15%	0%	6016	Actions
🟢	ms	User: mstephenson	10s	20%	14%	2%	1147	Actions
🔴	cops_user	User: cops	5s	10%	12%	0%	2047472	Actions
🔴	DiagBundlesAnalysis	User: cops Statement Type: QUERY	5s	10%	12%	0%	1689191	Actions

Defining Workload Views Manually

Steps for manually defining your workload views.

About this task

This task describes how to manually define your Workload Views.

To view this feature in action, watch the following video:

[Video: Classifying Workloads to Gain Insights](#)

Figure 1: Video: Classifying Workloads to Gain Insights

For better video quality, click **YouTube** in the lower right corner of the video player to watch this video on YouTube.com.

Procedure

1. Verify that you are logged in to the Cloudera Data Platform.
2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Search field of the Clusters page, enter the name of the cluster whose workloads you want to analyze.

5. From the time-range list in the Cluster Summary page, select a time period that meets your requirements.

The screenshot shows the Cloudera Workload XM interface. The navigation panel on the left has 'Summary' selected. The main content area shows a 'Trend' chart and an 'Outliers' section. A dropdown menu is open, showing a list of time ranges: 'Default (03/08/2018 - 05/03/2018)', 'Today', 'Yesterday', 'Last 7 Days', 'Last 30 Days', and 'Customize'. An orange arrow points to the 'Default' option.

6. From the navigation panel under Data Warehouse, select **Workloads**, and then click the **Define New** list and select **Manually define view**.

The screenshot shows the Cloudera Workload XM interface. The navigation panel on the left has 'Workloads' selected. The main content area shows a message: 'You have not defined any workloads yet.' and a 'Define Workload' button. An orange arrow points to the 'Workloads' option in the navigation panel.

The Define Data Warehouse Workload View opens, where you can define a set of criteria that enables you to analyze a specific set of queries.

For example, you can review all the failed queries using a specific database that are subject to a fifteen second SLA. For this example's view definition, a name is used that describes the view, as in **applog_db_under_15s**. Where the criteria will monitor queries that use the **applog** database, when **100** percent of these queries miss a

15s SLA, or when the total query execution time exceeds 15 seconds, then the workload is flagged with a failing status.:

Cloudera **Workload XM** Cluster_14 ▾

Define Data Warehouse Workload View

* Name ⓘ

* Criteria ⓘ Database ▾ = ▾ applog ▾ ⊞ ⊕

* SLA ⓘ
Example: 1h 2m 3s 5ms

* Warning Threshold ⓘ % queries missed SLA
Workload will be flagged in failing status when it reaches this threshold.

7. Click **Preview**, which displays a summary of the queries matching this criteria:

Cloudera **Workload XM** Cluster_14 ▾

Define Data Warehouse Workload View

* Name ⓘ

* Criteria ⓘ Database ▾ = ▾ applog ▾ ⊞ ⊕

* SLA ⓘ
Example: 1h 2m 3s 5ms

* Warning Threshold ⓘ % queries missed SLA
Workload will be flagged in failing status when it reaches this threshold.

ⓘ Cluster default date range is in the past, metrics reflect the status of the period.

03/08/2018 - 05/03/2018

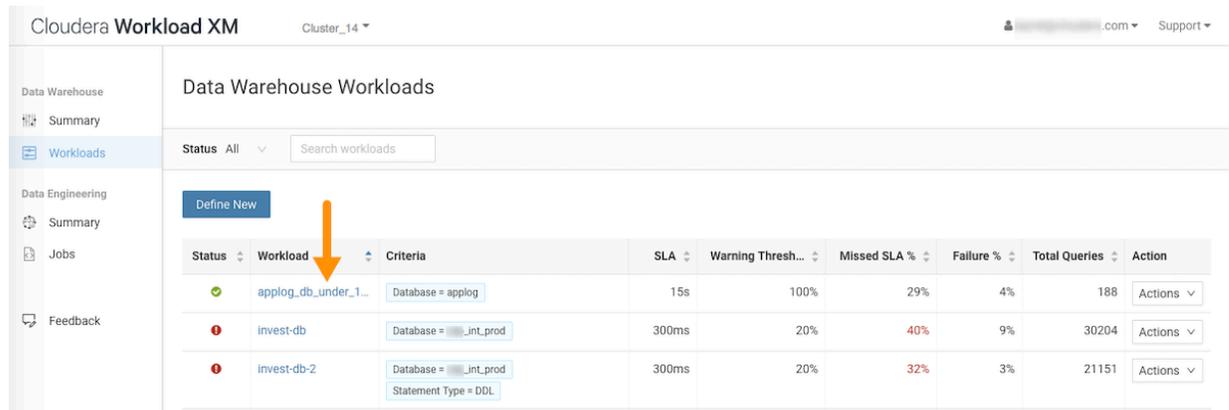
Total Queries	Missed SLA %
188	29%

Date range of queries in the workload.

Summary of queries that match the criteria.

8. When you are satisfied with the results, click **Save**.

The Data Warehouse Workloads page opens and your workload view appears in the **Workload** column.



Cloudera Workload XM Cluster_14

Data Warehouse Workloads

Status All Search workloads

Define New

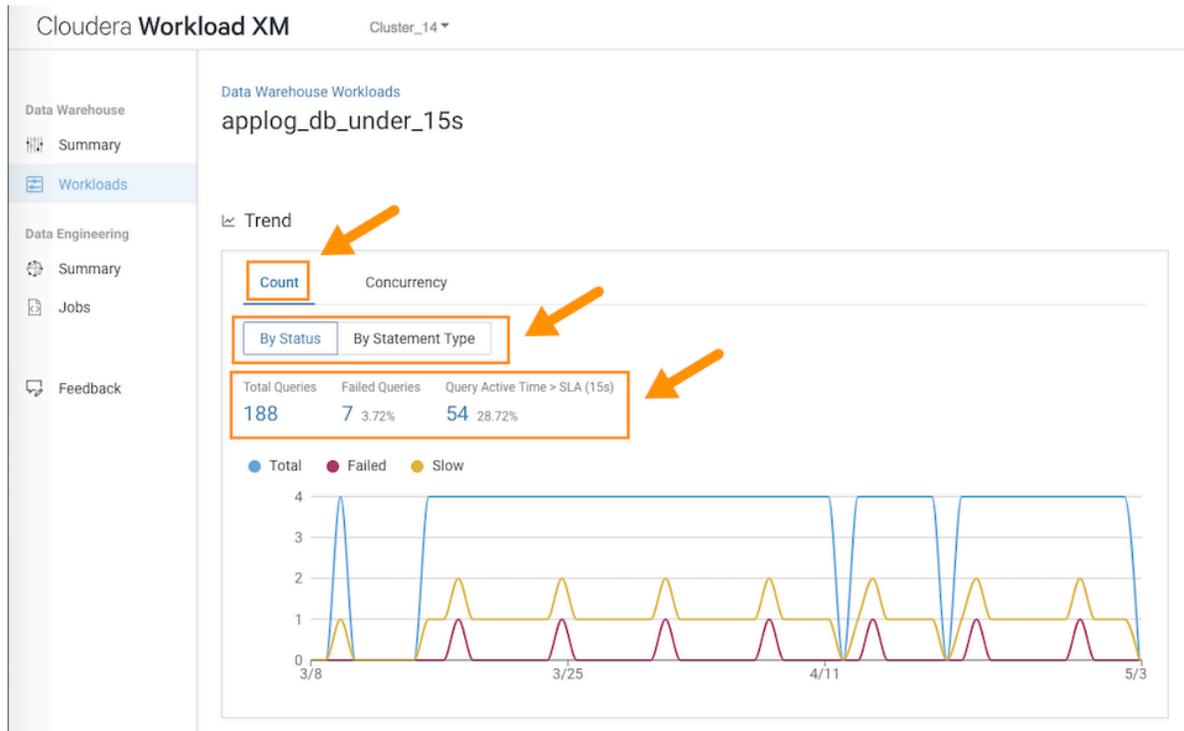
Status	Workload	Criteria	SLA	Warning Thresh...	Missed SLA %	Failure %	Total Queries	Action
🟢	applog_db_under_1...	Database = applog	15s	100%	29%	4%	188	Actions
🔴	invest-db	Database = _int_prod	300ms	20%	40%	9%	30204	Actions
🔴	invest-db-2	Database = _int_prod Statement Type = DDL	300ms	20%	32%	3%	21151	Actions



Tip: When you have a long list of Workload views you can locate a specific view by clicking the up or the down arrow next to the **Workload** column heading, which sorts the column alphabetically in ascending or descending order.

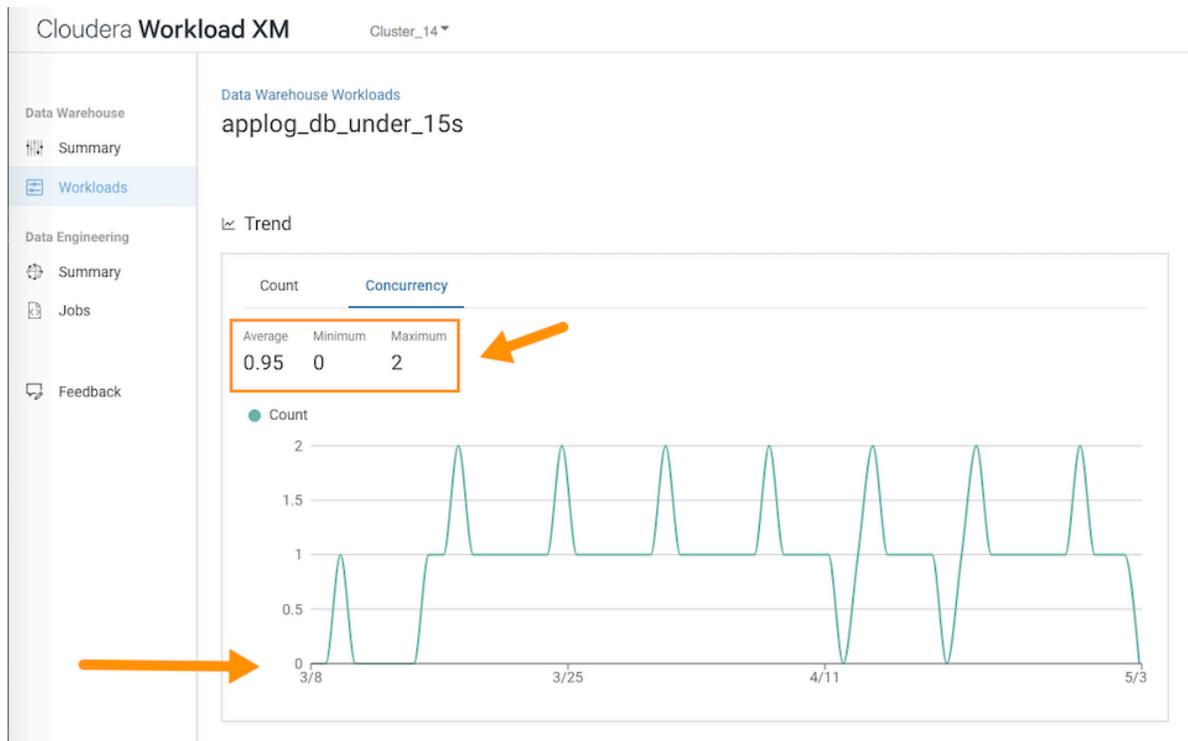
9. To view more information about the workloads using the view's formula, select the name of the workload view. The details page for your workload view opens visually displaying the views details as charts and tabs that you can use to analyze, as in this example, which visually displays how this group of queries is meeting its SLA:

- In the Trend section, you can display the counts of executing queries, either **By Status** or **By Statement Type**.



- To view further details, click the number under **Total Queries**, **Failed Queries**, and **Query Active Time**.
- To view the number of queries executing concurrently, click the **Concurrency** tab.

In this example, the maximum concurrency for this view is **2**. This indicates that for the queries flagged by this view, a maximum of only two queries access the same data at the same time during the specified time period. The graph, displays how the concurrency fluctuates over the date range specified for the workload view.



- In the Outliers section, you can display the **Slowest** queries and the **Most Complex** queries:

The screenshot shows the 'Outliers' section with a table of queries. The 'Slowest' and 'Most Complex' filter buttons are highlighted with an orange box and arrow. A text annotation says 'Click to view the slowest queries or the most complex.' Another arrow points to a query in the table with the annotation 'Click a query to view more details, such as the query code.'

Status	Query	Duration	User	Pool
❗	compute stats snaptor	44m 35s		Unknown
❗	compute stats snaptor	44m 14s		Unknown
❗	compute stats snaptor	43m 21s		Unknown
❗	compute stats snaptor	42m 58s		Unknown
❗	compute stats snaptor	42m 43s		Unknown
❗	compute stats snaptor	42m 25s		Unknown

- You can also display the users who are executing the queries in your workload view and the pools and databases being used. To view further details about the specific queries being executed, click a user, pool, or the database name:

Outliers

Queries **Top User/Pool/Database**

Users Pools Databases

User	# Queries	%
user1	188	100%

Click to view details.

- You can also display the different statement types contained in the workload view, the active time of the queries, and you can further drill down to view more granular details on each query.

Troubleshooting with the Job Comparison Feature

Steps for comparing two different runs of the same job, which is especially useful when you notice unexpected changes. For example, when you have a job that consistently completes within a specific amount of time and then it starts taking longer, comparing two runs of the same job enables you to analyze the differences so that you can troubleshoot the cause.

About this task

Describes how to compare any two runs of a job using the Job Comparison tool.

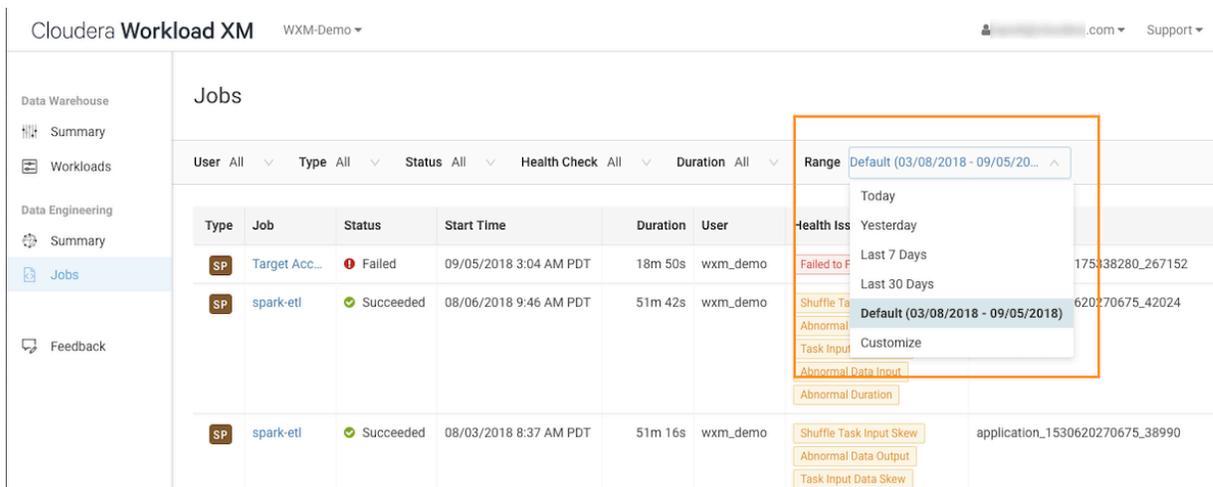


Note: When a job is flagged as slow, there is a **Compare with Previous Run** link on the Job page that opens the Job Comparison tool and compares the current run of the job with its last run.

Procedure

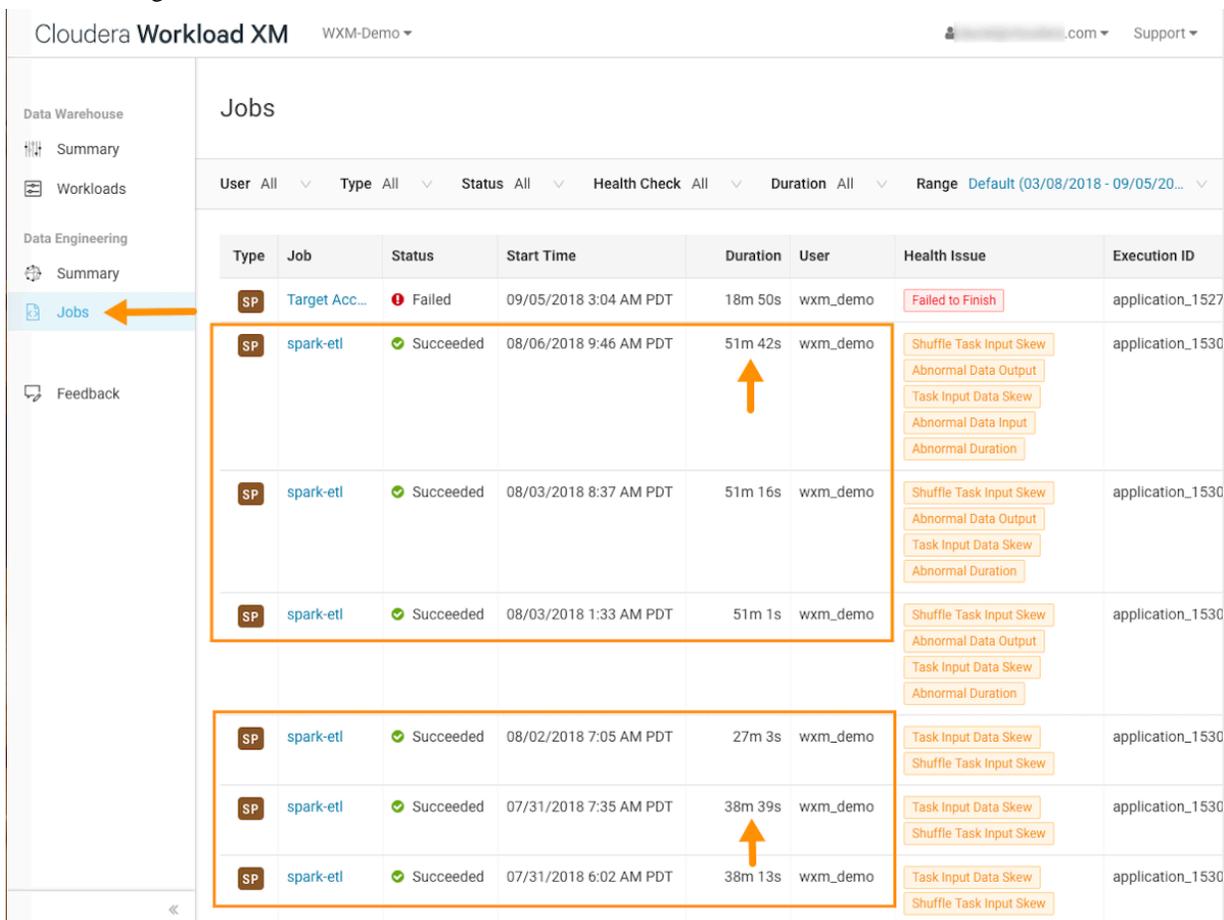
1. Verify that you are logged in to the Cloudera Data Platform.
2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Search field of the Clusters page, enter the name of the cluster whose workloads you want to analyze.

- From the navigation panel under Data Engineering, select **Jobs** and then from the time-range list in the Cluster Summary page, select a time period that meets your requirements.



- View the list of jobs that have executed during the selected time period:

The following example reveals that though the `spark-etl` job runs often, the last three runs have taken significantly longer. Where, on August 2, the duration was 27 minutes, but on August 3, the duration almost doubled to 51 minutes. The Job Comparison tool will enable you to examine both runs to determine why the duration changed:



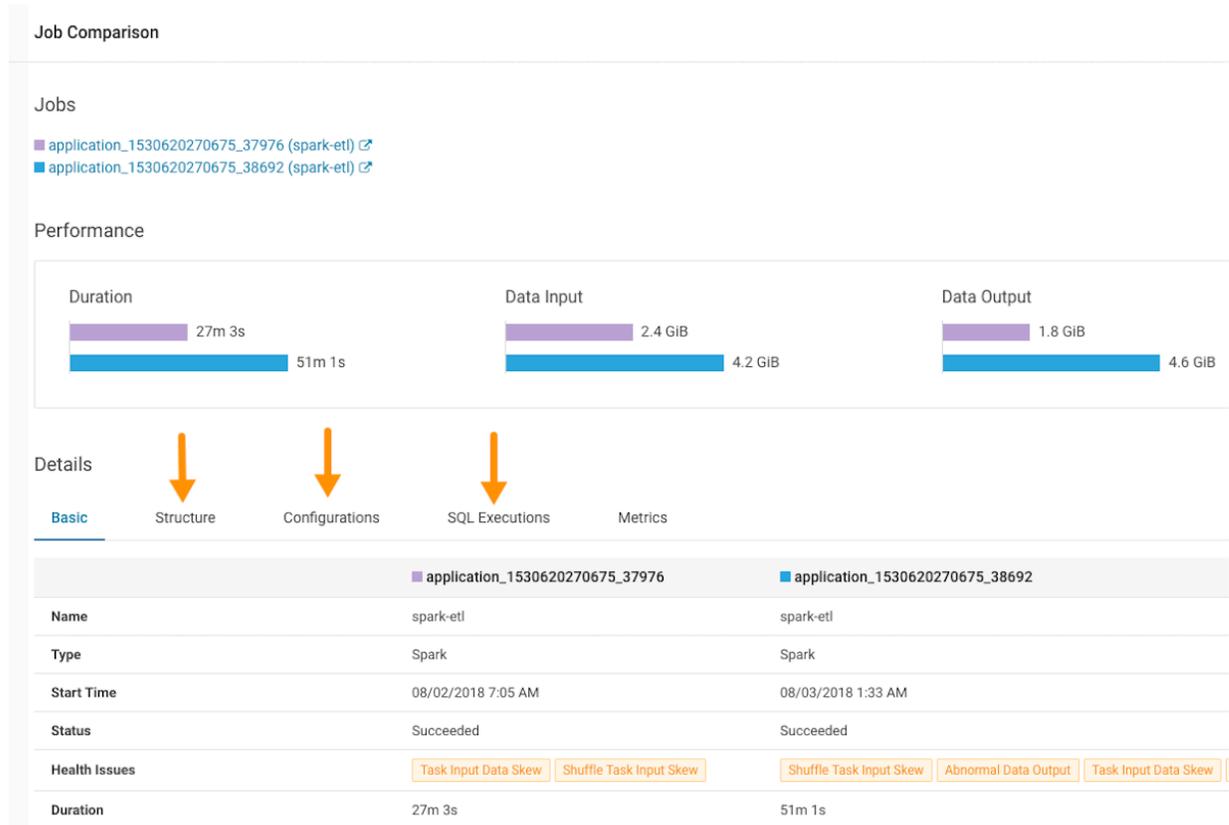
7. Select one of the runs of the `spark-etl` job, and then in the Jobs detail page, click the **Trends** tab:

Up to 30 runs prior to the selected job are displayed. Notice that in the **Input** and **Output** columns there are different amounts of data processed by the job. For example, on August 2, the job processed 2.4 GB of data and outputted 1.8 GB. However, on August 3, the job processed 4.2 GB, almost twice as much data, and it outputted 4.6 GB. The Job Comparison tool will enable you to examine both runs to determine why the amount of data changed:

Start Time	Execution ID	Status	Duration	Input	Output	Health Issues
08/03/2018 1:33 AM	application_1530620270675_38692	Succeeded	51m 1s	4.2 GiB	4.6 GiB	Shuffle Task Input Skew, Abnormal Data Output, Task Input Data Skew, Abnormal Duration
08/02/2018 7:05 AM	application_1530620270675_37976	Succeeded	27m 3s	2.4 GiB	1.8 GiB	Task Input Data Skew, Shuffle Task Input Skew
07/31/2018 7:35 AM	application_1530620270675_35552	Succeeded	38m 39s	4.2 GiB	3.6 GiB	Task Input Data Skew, Shuffle Task Input Skew
07/31/2018 6:02 AM	application_1530620270675_35463	Succeeded	38m 13s	4.2 GiB	3.6 GiB	Task Input Data Skew, Shuffle Task Input Skew
05/24/2018 11:45 AM	application_1521297608030_101095	Failed	17m 3s	1.5 GiB	0 B	Shuffle Task Input Skew, Failed to Finish
03/21/2018 8:23 AM	application_1521297608030_4085	Failed	1h 16m 2s	64.2 GiB	0 B	Shuffle Task Input Skew, Failed to Finish
03/08/2018 4:49 AM	application_1520383227161_2794	Failed	15m 16s	17.9 GiB	0 B	Shuffle Task Input Skew, Failed to Finish

8. To compare two job runs, select the check boxes adjacent to the job runs you require, in this case the runs for August 2 and August 3 are selected, and then click **Compare**.

The Job Comparison page opens displaying more details about each job. For this example's comparison, the tabs that will contain more information are the **Structure**, **Configurations**, and the **SQL Executions** tabs:

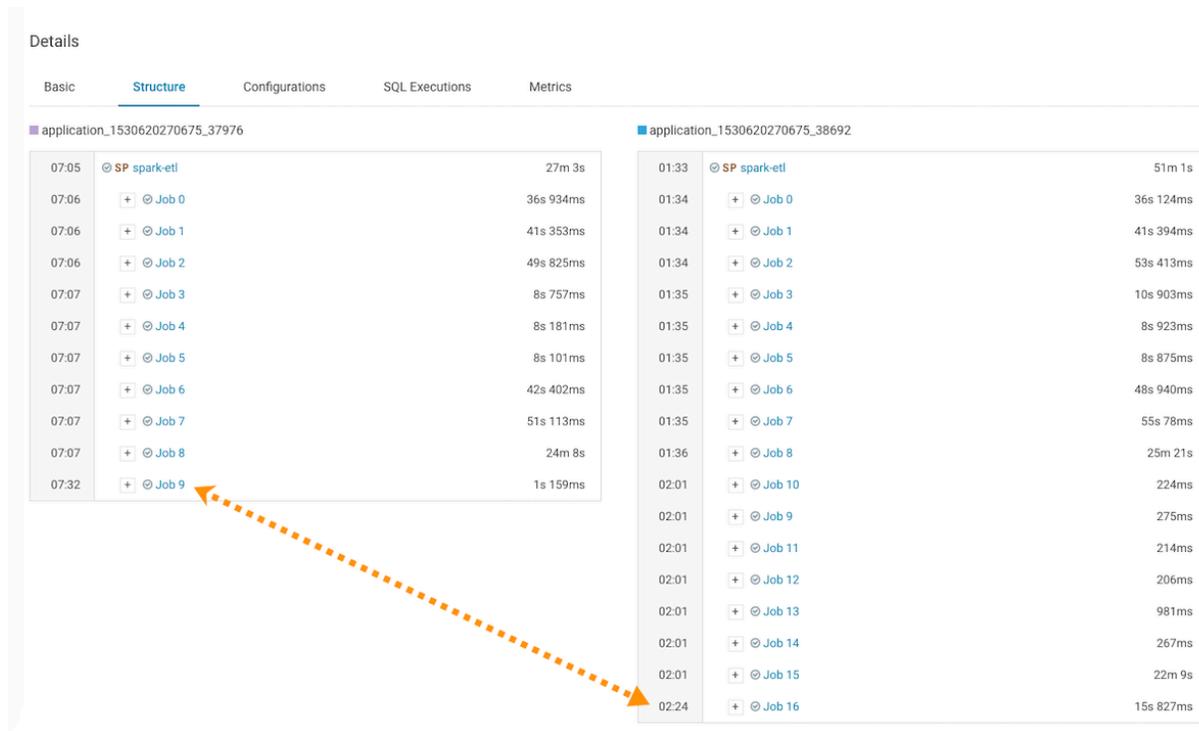


Note: The **SQL Executions** tab is only available for Spark jobs.

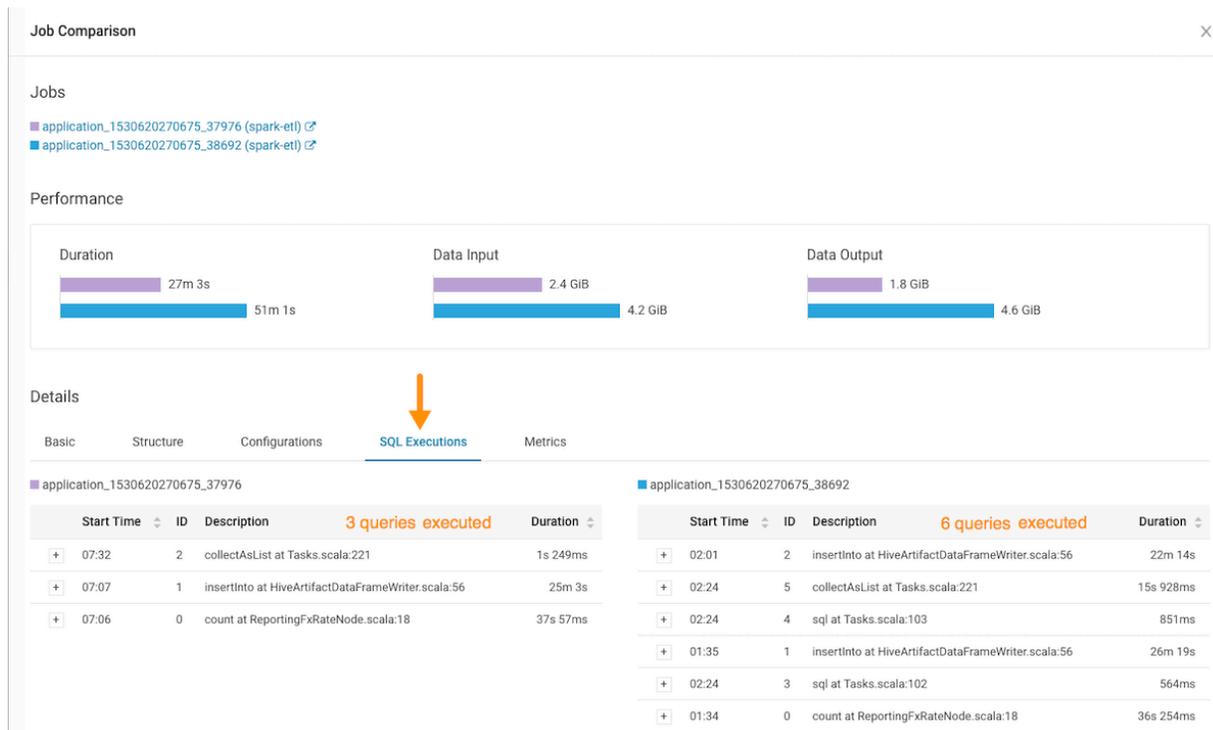
9. The **Structure** tab page, displays the sub-jobs executed for both runs of the `spark-etl` job:

In this example, the job that took 27 minutes only executed 9 sub-jobs and the job that took 51 minutes, almost twice as much time, executed 16 sub-jobs, almost twice as many. Selecting any of the listed sub-jobs displays more details.

When the **Configurations** tab was examined, it revealed that the configurations between the two runs of this job were identical, so a configuration change probably did not cause this anomaly.



10. When the **SQL Executions** tab was selected, it revealed that twice as many Spark queries executed for the job that took the longest duration:



Results

The analysis from the Job Comparison tells us that either the Spark SQL code was changed by the Job Developer or that the data which the code ran against triggered more of the Spark queries in the job. The Workload XM Job Comparison tool helped narrow the number of causes that produced this anomaly. In this example, the change in job duration appears to be expected so no further troubleshooting is required.

Identifying File Size Storage Issues

Data stored in small files or partitions may create performance issues. The File size reporting feature helps you identify data that is stored inefficiently in small files or partitions.

A table's data maybe stored in a large number of files, perhaps millions of files. For example, the first time you run an Impala query, Impala also loads the metadata for each file, which can cause processing delays. In addition, every time you change a query, refresh the metadata, or add a new file or partition, Impala reloads the metadata. This puts pressure on the NameNode, which stores each file's metadata. For more information about the problems caused by small files and what you can do to fix those problems, see [Handling Small Files on Hadoop with Hive and Impala](#) on the Cloudera Engineering Blog.

The Workload XM file size reporting enables you to identify tables that have a large number of files or partitions. For example, if a query runs slowly or your Impala cluster crashed, you can view the table's metadata that the query is using to determine if a large number of files or partitions is causing the problem.



Important: At this time the Workload XM File Size Report feature is only supported on CDH Workload clusters, version 6.3 to version 7.0, with Cloudera Navigator enabled. CDP Workload clusters are not supported.



Note: Before you can view file size metadata in Workload XM, you must enable file size reporting in Cloudera Manager. Once enabled, the file size metadata is saved in HDFS. Telemetry Publisher gets the metadata from HDFS and sends it to Workload XM.

Displaying File Size Metadata

Steps for displaying a table's File Size report and the metadata about the table's file's size distribution.

About this task

Describes how to open a table's File Size report and the metadata of a file.



Important: At this time the Workload XM File Size Report feature is only supported on CDH Workload clusters, version 6.3 to version 7.0, with Cloudera Navigator enabled. CDP Workload clusters are not supported.

Procedure

1. Verify that you are logged in to the Cloudera Data Platform.
2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Search field of the Clusters page, enter the name of the cluster whose workloads you want to analyze.
5. From the navigation panel under Data Warehouse, select **File Size Report**.

6. In the **File Size Report** page, either search for a specific table, or locate the table by sorting the tables by the number of files, the number of partitions, or the table size.

For example, the File Size Reports shows that the Animantarx table has 7 million files and 913 partitions.

Cloudera Workload XM MyCluster ▼ triceratops@cloudera.com Support ▼

Data Warehouse

- Summary
- Workloads
- File Size Report**

Data Engineering

- Summary
- Jobs
- Feedback

Table File Size Report

As of Mon, Apr 15, 2019 5:13 PM

Search

Table	Files	Median File Size	File Size Distribution	Partitions	Table Size	Database
Animantarx	7M	36.7 KiB		913	229.6 GiB	Carnotaurus
Bonapartenykus	3.1M	1 MiB		397.3K	3.3 TiB	Bruhathkayosaurus
Balaur	1.7M	469 KiB		1K	1.7 TiB	Chasmosaurus
Alwalkeria	595.3K	2.5 MiB		1.7K	1.4 TiB	Cetiosaurus
Atlasaurus	401.8K	1.2 KiB		4	477.6 MiB	Chilantaisaurus
Angolatitan	358.9K	168 KiB		7.1K	455.9 GiB	Cerasinops
Anatosaurus	346.9K	1.9 KiB		5.1K	27.3 GiB	Byronosaurus

7. To display details about the table's file size distribution, select a table name.

For example, the following table's details window shows that the Aerosteon table uses 42 data files that range from 10 to 24.5 GiB and the graph displays the Q1 and Q3 file size distribution.

Cloudera Workload XM MyCluster ▼ triceratops@mycompany.com Support ▼

Data Warehouse

- Summary
- Workloads
- File Size Report**

Data Engineering

- Summary
- Jobs
- Feedback

Atlascopcosaurus	42	41.1 KiB		0	1.7 MiB	Chilantaisaurus
Agustinia	42	2.9 KiB		0	168.3 KiB	Chialingosaurus
Apatosaurus	42	1.2 MiB		0	24.1 MiB	Cerasinops
Bistahieversor	42	9.8 MiB		0	28.9 MiB	Camelotia
Amurosauros	42	28.9 MiB		0	28.9 MiB	Brachyophosaurus
Acrotholus	42	2 GiB		0	2 GiB	Chungkingosaurus
Aorun	42	12.4 MiB		0	12.4 MiB	Carnotaurus
Aerosteon	42	11.9 GiB		0	572.5 GiB	Bruhathkayosaurus
Anatosaurus	42	1.2 MiB		0	48.7 MiB	Chasmosaurus
Afrovenator	42	338.9 KiB		0	14.5 MiB	Byronosaurus
Barosaurus	42	367.3 KiB		0	14.9 MiB	Carcharodontosaurus
Aucasaurus	42	789 B		0	32.9 KiB	Brachyrachelopan

File size distribution for Aerosteon

Median 11.9 GiB

10 GiB —●— 24.5 GiB

Q1 10.8 GiB Q3 16.5 GiB

There are 42 files backing this table. They range from 10 GiB to 24.5 GiB with a median of 11.9 GiB.

Displaying the Metadata of a Table

Steps for displaying a table's metadata that could be causing a query to run slowly.

About this task

Describes how to display the metadata of table used in your query, such as the table's file size distribution that could be causing your query statement to run slowly.

Procedure

1. Verify that you are logged in to the Cloudera Data Platform.

2. In a supported browser, log in to Workload XM.
3. From the **Your Enterprise Data Cloud** landing page, select the **Workload Manager** tile.
4. In the Search field of the Clusters page, enter the name of the cluster whose workloads you want to analyze.
5. From the navigation panel under Data Warehouse, select **Summary**.
6. In the Queries page, select the query of interest and then select the **HDFS Tables Scanned** tab.

For example, the **Duration** column shows that the query took over six hours to run and the **HDFS Tables Scanned** section displays the metadata for the tables that were scanned.



Note: This is not the number of files accessed, but the total number of files that were in the table the last time a HDFS snapshot was taken before the query was run.

Cloudera Workload XM MyCluster ▾ triceratops@cloudera.com ▾ Support ▾

Data Warehouse

- [Summary](#)
- [Workloads](#)
- [File Size Report](#)

Data Engineering

- [Summary](#)
- [Jobs](#)

[Feedback](#)

Queries

1c4d1a6ca67e94f4:7977bd4d00000000

Summary Trend ✔ Succeeded 📅 04/24/2019 2:05 AM CDT 👤 stegosaurus 🔗 DML 📧 root.default 👤 Profile

Joins	Duration	Rows Produced	Aggregate CPU Usage	Aggregate Memory Usage	Peak Memory Usage
2	6h 34m 3s	571	41h 29m 23s	30.6 TiB × s	1.5 GiB

Basic Operators Hosts **HDFS Tables Scanned**

As of Mon, Apr 15, 2019 5:13 PM

🔔 The table metadata below is taken from the last HDFS snapshot before the query was run.

Table	Files	Median File Size	File Size Distribution	Partitions	Table Size	Database
Aragosaurus	739	39.5 KiB	↑-----	0	78.8 MiB	Brachytrachelopan

7. To display the file size distribution details for a table, click the Table name .