

Dataproc

Clusters

CREATE CLUSTER

REFRESH

START

STOP

DELETE

REGIONS

+ 5 RECOMMENDED ALERTS

LEARN


SHOW INFO PANEL

Filter

Search cluster by properties, press Enter

Sorry, the server was not able to fulfill your request.

☐	Name ↑	Status	Region	Zone	Total worker nodes	Flexible VMs?	Scheduled deletion	Cloud Storage staging bucket	Created	Labels
☐	example-cluster	Running	us-central1	us-central1-f	4	No	Off	dataproc-staging-us-central1-590053044234-qnywjida	Jul 22, 2025, 1:02:14 PM	goog-dataproc-: example-cl... goog-dataproc-: eaab6b01-8... goog-dataproc-: cluster-ea... goog-dataproc-: us-central1



Job details

CLONE

DELETE

STOP

REFRESH

Job ID

Job UUID

Type

Status

job-d965796e

ae3017b1-806a-45cc-ef5c-6b6e08a8fcc2

Dataproc Job

Succeeded

MONITORING

CONFIGURATION

The charts below represent the metrics from the cluster this job ran on, scoped to the time that this job was running. It is possible for more than one job to run on a cluster at a time, so these metrics may not reflect this job's resource usage accurately. Metrics for a job may lag behind the job run by several minutes.

SAVE AS DASHBOARD

RESET ZOOM

YARN memory

150GB

YARN pending memory

10GB

YARN NodeManagers

2

HDFS capacity

300GB

CPU utilization

50%

1 hour

6 hours

12 hours

1 day

2 days

4 days

7 days

14 days

30 days

Custom 1:03 PM - 1:08 PM

Output

LINE WRAP: OFF

Dismiss

Spark jobs take ~60 seconds to initialize resources.

25/07/22 11:07:27 INFO SparkEnv: Registering MapOutputTracker

25/07/22 11:07:27 INFO SparkEnv: Registering BlockManagerMaster

25/07/22 11:07:27 INFO SparkEnv: Registering BlockManagerMasterHeartbeat

25/07/22 11:07:27 INFO SparkEnv: Registering OutputCommitCoordinator

25/07/22 11:07:29 INFO MetricsConfig: Loaded properties from hadoop-metrics2.properties

25/07/22 11:07:29 INFO MetricsSystemImpl: Scheduled Metric snapshot period at 10 second(s).

25/07/22 11:07:29 INFO MetricsSystemImpl: google-hadoop-file-system metrics system started

25/07/22 11:07:30 INFO DataprocSparkPlugin: Registered 188 driver metrics

25/07/22 11:07:31 INFO DefaultHadoopFilesystemProxyProvider: Connecting to ResourceManager at example-cluster-m.us-central-1-f.c.qwiklabs-gcp-01-b6ed1e98aac.internal./10.128.0.3:8032

25/07/22 11:07:32 INFO HSFProxy: Connecting to Application History server at example-cluster-m.us-central-1-f.c.qwiklabs-gcp-01-b6ed1e98aac.internal./10.128.0.3:10200

25/07/22 11:07:34 INFO Configuration: resource-types.xml not found

25/07/22 11:07:34 INFO ResourceUtils: Unable to find 'resource-types.xml'.

25/07/22 11:07:34 INFO ResourceUtils: Submitted application application_1753182328888_0001

25/07/22 11:07:44 INFO RequestTracker: Detected high latency for [url=https://storage.googleapis.com/storage/v1/b/dataproc-temp-us-central1-598853844234-hwvrp2h/s/eaabb01-8e21-489e-9984-f988688d5cb9?spark-job-history?fields=bucket_name,timeCreated,updated,generation,metageneration,size,o

25/07/22 11:07:44 INFO GoogleCloudStorageImpl: Ignoring exception of type GoogleJsonResponseException: verified object already exists with desired state.

25/07/22 11:07:45 INFO GoogleHadoopOutputStream: hFlush(): No-op due to rate limit (RateLimiter[stableRate=0.2ops/s]): readers will "not" yet see flushed data for gs://dataproc-temp-us-central1-598853844234-hwvrp2h/eaabb01-8e21-489e-9984-f988688d5cb9/spark-job-history/application_17531823288

25/07/22 11:08:02 INFO DhfsGlobalStorageStatistics: periodic connector metrics: [action_http_post_request=4, action_http_post_request_duration=374, action_http_post_request_max=148, action_http_post_request_mean=93, action_http_post_request_min=46, action_http_put_request=2, action_http_put_

CONTEXT rateLimit_period=5 MINUTES]

P2 is roughly 9.1416765914167657

25/07/22 11:08:21 INFO DhfsGlobalStorageStatistics: Detected potential high latency for operation op_hFlush. latencyMs=612; previousMaxLatencyMs=443; operationCount=15; context-gs://dataproc-temp-us-central1-598853844234-hwvrp2h/eaabb01-8e21-489e-9984-f988688d5cb9/spark-job-history/applica

25/07/22 11:08:21 INFO DataprocSparkPlugin: Shutting down driver plugin. metrics=[action_http_patch_request=0, files_created=1, gcs_api_server_timeout_count=0, op_get_list_status_result_size=0, op_open=0, action_http_delete_request=2, gcs_api_time=3563, gcs_backoff_count=0, gcs_api_client_un

EQUIVALENT COMMAND LINE

Request to update cluster example-cluster submitted

Revision #3

Created 11 November 2022 21:24:46 by Nehemie

Updated 22 July 2025 11:13:34 by Nehemie