



Dremio Software

Monitoring Dremio

To maximize your investment in Dremio and to proactively identify and resolve issues related to Dremio before they have a broader impact on workload, it's essential to deploy a good monitoring solution. The solution should ensure overall cluster health and performance. Dremio provides a large set of metrics that can be utilized for a monitoring solution along with infrastructure-related metrics such as CPU, memory, etc.

Metrics to Monitor

Dremio can be deployed in various ways. Hence, the collection and reporting of infrastructure-related metrics depend on your deployment model (physical or virtual host, Kubernetes, or cloud). You can collect important metrics for monitoring via JMX, JDBC, ODBC, or the REST API. You can monitor Dremio directly with open-source tools such as Prometheus/Grafana or with commercially available tools such as AppDynamics, Datadog, etc. It is essential to align the monitoring solution with the existing monitoring infrastructure.

JMX Metrics

Java Management Extensions (JMX) is a specification for monitoring and managing Java applications. Since Dremio is a Java application, it uses JMX to expose several important metrics utilizing this framework.

NOTE

For information about modifying `dremio.conf` to enable JMX metrics, see [Enabling Node Metrics](#).

This configuration exposes JMX metrics on the URL `http://<dremio host>:<port>/metrics`. The following table lists important metrics, which Dremio node role supports them, and alert threshold details.

Metric Name	Description	Master Coordinator	Secondary Coordinator	Executor	Alert Threshold (if any)
<code>jobs.active</code>	Currently active jobs	Yes	Yes	No	
<code>jobs.active_15m</code>	Number of jobs in the previous 15-minute period	Yes	No	No	
<code>jobs.failed_15m</code>	Number of failed jobs in the last 15-minute period	Yes	No	No	WARN: 5% of total jobs CRITICAL: 10% of total jobs
<code>jobs.queue.<queue_name>.waiting</code>	Number of current waiting jobs	Yes	No	No	WARN: > 0 CRITICAL: > 50% of allowed concurrency for the queue
<code>dremio.memory.direct_current</code>	Direct memory used by the execution engine	Yes	Yes	Yes	WARN: 90% of allocated value CRITICAL: 95% of allocated value
<code>fragments.active</code>	Number of active query fragments (threads): Indicator on how starved Dremio is for CPU (monitor on executors)	Yes	Yes	Yes	WARN: 0.9 x number of CPU cores on executors for 5 min CRITICAL: 1 x number of CPU cores on

Metric Name	Description	Master Coordinator	Secondary Coordinator	Executor	Alert Threshold (if any)
					executors for 5 min
gc.G1-Young-Generation.time/ gc.G1-Young-Generation.count gc.G1-Old-Generation.time/ gc.G1-Old-Generation.count	Time per GC event (for young and old generations), assuming default G1GC is used for Garbage collection	Yes	Yes	Yes	WARN: > 10s CRITICAL: > 1m
memory.heap.usage	Ratio of memory.heap.used to memory.heap.max	Yes	Yes	Yes	WARN: > 75% CRITICAL: > 80% Comment: Monitor on Coordinator. Coordinator's JVM monitor automatically kills queries when heap utilization is > 85%.
reflections.failed	Currently failed data reflections	Yes	No	No	WARN: > 0 CRITICAL: > 10% of all reflections
reflections.active	Currently active data reflections	Yes	No	No	
reflections.refreshing	Data reflections currently refreshing or pending a refresh	Yes	No	No	
rpc.failure_15m	RPC connection failures in 15 minute period	Yes	Yes	Yes	WARN: seen on ~10% of available executors CRITICAL: seen on >25% of available executors

API Metrics

The Dremio coordinator exposes the REST API on the web UI port, which is 9047 by default. For more information about how to connect, authenticate, and submit requests, see [API Reference](#).

API endpoint	Description	Expected Values	Alert Threshold (if any)
GET /apiv2/server_status	Coordinator Status	"OK"	Anything other than "OK" or a slow response indicates that coordinator is either down or unhealthy

API endpoint	Description	Expected Values	Alert Threshold (if any)
GET /api/v3/source	Source Status	"good"	Any status other than "good"

SQL Metrics

SQL commands can be executed using ODBC, JDBC, or REST interfaces.

i NOTE

The API call does not return query results when executing SQL with the REST API. The API returns the ID of the submitted query.

SQL	Description	Alert Threshold (if any)
SELECT COUNT(*) FROM sys.memory	Canary query. Can be executed against a user dataset (small) or a Dremio internal table. This query should return results in ms. It is an indicator to overall cluster health.	WARN: > 30s CRITICAL: > 1m

POSIX Metrics

These metrics monitor metrics outside of Dremio that have an impact on stability and performance of Dremio.

Command	Description	Alert Threshold (if any)
df -h <Directory where Rocks DB is mounted>	Catalog DB free space. Catalog DB must have free space to create reflections, update profiles, and run jobs.	WARN: > 80% used CRITICAL: > 90% used