

Tune Application Performance

Contents:

- *Application Performance*
 - *Other applicable parameters*
- *Limit Application Memory Utilization*
- *Photon Running Environment Performance*

This section provides general guidelines for tuning the configuration parameters of the Trifacta® node for varying loads.

NOTE: These guidelines are estimates of what should provide satisfactory performance. You should review particulars of the variables listed below in detail prior to making recommendations or purchasing decisions.

For more information on tuning the performance of the connected cluster, see *Tune Cluster Performance*.

Application Performance

Some Trifacta platform services running on the Trifacta node can use multiple processes to serve more requests in parallel (e.g., `webapp` and `vfs-service`). By increasing the number of processes, these services are able to serve more requests in parallel and improve the application's response time under load.

Other services, such as `batch-job-runner`, `data-service`, and `scheduling-service` use multiple threads within a single process as needed to serve concurrent requests. Each service may have tuning parameters that can be adjusted according to load.

Tip: Unless specific recommendations apply below, you should use the default configuration.

For an expected number of peak concurrent (active) users, P , set the following parameters.

You can apply this change through the *Admin Settings Page* (recommended) or `trifacta-conf.json`. For more information, see *Platform Configuration Methods*.

Parameter	Default Value	Recommended Value	Example
<code>webapp.numProcesses</code>	2	$P / 15$, rounded up to the nearest integer, minimum of 2	for $P = 40$, set to 3
<code>webapp.database.pool.maxConnections</code>	10	$5 * \text{webapp.numProcesses}$	for <code>webapp.numProcesses = 3</code> , set to 15
<code>vfs-service.numProcesses</code>	2	$P / 50$, rounded up to the nearest integer, minimum of 2	for $P = 225$, set to 5

Other applicable parameters

The following configuration parameters also affect application performance.

NOTE: Avoid modifying these parameters unless instructed by *Trifacta Support*.

Parameter	Default Value
<code>batch-job-runner.systemProperties.httpMaxConnectionsPerDestination</code>	50

Limit Application Memory Utilization

Several Trifacta node services allow limitations on memory utilization by varying their JVM configuration's `-Xmx` value. These can be limited by modifying the following parameters:

Parameter	Default Configuration
<code>batch-job-runner.jvmOptions</code>	<code>-Xmx1024m</code>
<code>diagnostic-server.jvmOptions</code>	<code>-Xmx512m</code>
<code>data-service.jvmOptions</code>	<code>-Xmx128m</code>
<code>spark-job-service.jvmOptions</code>	<code>-Xmx128m</code>

Other services have low memory requirements.

Photon Running Environment Performance

Jobs run on the Trifacta node and "Quick Scan" samples are executed by the Photon running environment embedded on the Trifacta node, running alongside the application itself. Two main parameters can be used to tune concurrency of job execution and throughput of individual jobs:

Parameter	Description	Default
<code>batchserver.workers.photon.max</code>	Maximum number of simultaneous Photon processes; once exceeded, jobs are queued	2
<code>photon.numThreads</code>	Number of threads used by each Photon process	4

Increasing the number of concurrent processes allows more users' jobs to execute concurrently. However, it also leads to resource contention among the jobs and the application services.

Photon's execution is purely in memory. It does not spill to disk when the total data size exceeds available memory. As a result, you should configure limits on Photon's memory utilization. If a job exceeds the configured memory threshold, it is killed by a parent process tasked with monitoring the job.

Parameter	Description	Default
<code>photon.memoryMonitorEnabled</code>	Set <code>true</code> to enable the monitor, and set to <code>false</code> (limited only by operating system) otherwise	<code>true</code>
<code>photon.memoryPercentageThreshold</code>	Percentage of available system memory that each Photon process can use (if <code>photon.memoryMonitorEnabled</code> is <code>true</code>)	50

A reasonable rule of thumb: the input data size should not exceed one tenth of the job's memory limit. This rule of thumb accounts for joins and pivots and other operations that can increase memory usage over the data size. However, this parameter is intended as a safeguard; it is unlikely that all running jobs would approach the memory limit simultaneously. So you should "oversubscribe" and use slightly more than $(100 / \text{batchserver.workers.photon.max})$ for this threshold.

In addition to a memory threshold, execution time of any Photon job can also be limited via the following parameters:

Parameter	Description	Default
<code>photon.timeoutEnabled</code>	Set <code>true</code> to enable the timeout, and set to <code>false</code> (unlimited) otherwise	<code>true</code>
<code>photon.timeoutMinutes</code>	Time in minutes after which to kill the Photon process (if <code>photon.timeoutEnabled</code> is <code>true</code>)	180

For more information, see *Configure Photon Running Environment*.