



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

# Cassandra Installation and Configuration Guide

Performance Tuning

12/14/2025

# Performance Tuning

Besides configuring keyspaces and column families, it is possible to further tweak the performance of Cassandra by editing `cassandra.yaml` (Node and Cluster Configuration) or by editing `cassandra-env.sh` (JVM Configuration).

Descriptions of tunable properties can be found in both `cassandra.yaml` and `cassandra-env.sh`. A summary of these properties can be seen in the tables below:

## Performance Tuning Properties (cassandra.yaml)

| Option                                        | Default            | Description                                                                                                                                                                                                                                                                                        |
|-----------------------------------------------|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>column_index_size_in_kb</code>          | 64                 | The size at which column indexes are added to a row. Value should be kept small if only a select few columns are consistently read from each row as a higher value implies that more row data must be deserialized for each read (until index is added).                                           |
| <code>commitlog_sync</code>                   | periodic           | Allowed values are periodic or batch. In periodic mode, the value of <code>commitlog_sync_period_in_ms</code> determines how frequently the commitlog is synchronized to disk. Writes are acknowledged at every periodic sync. If set to batch, writes are not acknowledged until fsynced to disk. |
| <code>commitlog_sync_period_in_ms</code>      | 10000 (10 seconds) | Determines how often (in milliseconds) to sync commitlog to disk when <code>commitlog_sync</code> is set to periodic.                                                                                                                                                                              |
| <code>commitlog_total_space_in_mb</code>      | 4096               | When commitlog reaches specified size, Cassandra flushes memtables to disk for oldest commitlog segments. Reduces amount of data to replay on startup.                                                                                                                                             |
| <code>compaction_throughput_mb_per_sec</code> | 16                 | Throttles compaction to the given total throughput across entire system. Value should be proportional to rate of write throughput (16 to 32 times). Setting to 0 disables compaction throttling.                                                                                                   |
| <code>concurrent_compactors</code>            | 1 (per CPU core)   | Max number of concurrent compaction processes allowed on a node.                                                                                                                                                                                                                                   |

| Option                                      | Default              | Description                                                                                                                                                  |
|---------------------------------------------|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------|
| concurrent_reads                            | 16                   | Recommended setting is 16 * number_of_drives. This allows enough operations to queue such that the OS and drives can reorder them and minimize disk fetches. |
| concurrent_writes                           | 32                   | Number of concurrent writes should be proportional to number of CPU cores in system. Recommended setting is (8 * number_of_cpu_cores).                       |
| memtable_flush_writers                      | 1 per data directory | Number of memtable flush writer threads. Influences flush performance and can be increased if you have a large Java heap size and many data directories.     |
| memtable_heap_space_in_mb                   | 1/4 of heap          | Total memory used for all column family memtables on a node.                                                                                                 |
| stream_throughput_outbound_megabits_per_sec | 400                  | Max outbound throughput on a node for streaming file transfers.                                                                                              |

## JVM Configuration Settings

**Linux:** `conf/cassandra-env.sh`

**Windows:** `bin\cassandra.bat`

| Option                            | Default                           | Description                                                                                                                                                                |
|-----------------------------------|-----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| MAX_HEAP_SIZE                     | Half of available physical memory | Maximum heap size for the JVM. Same value is used for minimum heap size, allowing heap to be locked in memory. Should be set in conjunction with HEAP_NEWSIZE.             |
| HEAP_NEWSIZE                      | 100 MB per physical CPU core      | Size of young generation. Larger value leads to longer GC pause times while smaller value will typically lead to more expensive GC. Set in conjunction with MAX_HEAP_SIZE. |
| com.sun.management.jmxremote.port | 7199                              | Port on which Cassandra listens for JMX connections.                                                                                                                       |

| Option                                      | Default | Description                                   |
|---------------------------------------------|---------|-----------------------------------------------|
| com.sun.management.jmxremote.ssl            | false   | Enable/disable SSL for JMX.                   |
| com.sun.management.jmxremote.authentication | false   | Enable/disable remote authentication for JMX. |