

Cassandra Limitations

Limitations

CQL

- No join or subquery support, and limited support for aggregation. This is [by design](#), to force you to denormalize into partitions that can be efficiently queried from a single replica, instead of having to gather data from across the entire cluster.
- [Ordering is done per-partition](#), and is specified at table creation time. Again, this is to enforce good application design; sorting thousands or millions of rows can be fast in development, but sorting billions in production is a bad idea.

Storage engine

- All data for a single partition must fit (on disk) on a single machine in the cluster. Because partition keys alone are used to determine the nodes responsible for replicating their data, the amount of data associated with a single key has this upper bound.
- A single column value may not be larger than 2GB; in practice, "single digits of MB" is a more reasonable limit, since there is no streaming or random access of blob values.
- Collection values may not be larger than 64KB.
- The maximum number of cells (rows x columns) in a single partition is 2 billion.

Limitations that will be gone soon

- There is no cursor support, so large resultsets must be manually paged. Cursor support is [scheduled for 2.0](#).

<https://c.statcounter.com/9397521/0/fe557aad/1/> | stats