

DataStax Enterprise – In Memory Option

Nearly all IT analyst groups recommend the use of in-memory technology as a way to increase performance for today’s modern applications that require very fast response times. While in-memory computing does not lend itself for use in all application use cases, in those where its use is appropriate, the performance difference between in-memory and disk-based technologies can be many orders of magnitude with some differentials reaching anywhere from 30-100X.

The in-memory option of DSE provides in-memory computing abilities to Cassandra, and lets developers, architects and administrators easily choose what parts (some or all) of a database reside fully in RAM. The in-memory option of DSE is designed to service use cases that lend themselves to in-memory computing while allowing disk-based workloads to be quickly serviced by Cassandra’s traditional storage model. This design allows applications with both in-memory and disk-based needs to be supported by one database platform.

Simple to Use

DSE’s in-memory option is very easy to use and understand. A Cassandra table can be created as an in-memory or traditional disk-based table. Existing disk-based tables can also be easily altered to be in-memory tables with the reverse also being true. Developers and DBAs can utilize options on in-memory tables to specify characteristics such as how large an in-memory table can grow in size, how often it flushes changed data to disk for backup, and more.

Application and System Transparency

Once created, an in-memory table acts and presents as any other Cassandra table, making it completely transparent to an application and the overall database. There is nothing special a developer has to do to interact with an in-memory table vs. a disk-based table. An in-memory table provides all the primary benefits of traditional disk-based Cassandra tables – automatic distribution of table data across the RAM in nodes that comprise a cluster, powerful replication abilities, multi-data center and cloud availability zone support, indexing, flexible data type support, etc.

Writes to in-memory tables function in the same way as with disk-based tables. Data is first written to a commit log to ensure data durability and then to the in-memory table. Unlike disk-based tables, all read operations are satisfied completely in RAM. DSE’s in-memory option utilizes memory in both the Java heap and non-Java memory structures to hold data and is therefore able to address large amounts of RAM on each node in a cluster.

Fast Performance for Key Use Cases

DSE’s in-memory option delivers lightning-fast read performance for use cases that lend themselves to in-memory operations. Typical use cases that benefit from DSE’s in-memory computing option include those with primarily read-only workloads with slowly changing data and/or semi-static datasets. An example use case might be product catalogs that are refreshed nightly, but read constantly during the day. Workloads that are not suitable for DSE’s in-memory option include those that heavily changing data or monotonically growing datasets that are predicted to exceed the RAM capacity on the nodes/cluster.



DataStax Enterprise delivers constant uptime and linear scale performance for online applications needing transactional, analytical, search, and in-memory workload support in a single platform.

No Data Loss

There is no possibility for data loss with in-memory tables as all incoming writes go first to a disk-based commit log and are then written to the in-memory table. Because Cassandra is so fast at writing data, the commit log portion of the write operation completes in microseconds. In-memory tables can be backed up and restored as any other Cassandra table.

In Memory Management and Monitoring

In-memory table usage can easily be monitored with DataStax OpsCenter. Information for all in-memory tables can be easily viewed, with historical usage information for in-memory tables being automatically collected and stored by OpsCenter.

In-memory performance and usage can be viewed over time in historical trend analysis fashion with OpsCenter. OpsCenter also supplies the ability to forecast the future state of in-memory tables (e.g. how large will an in-memory table be in 6 months?) Lastly, alerts on in-memory tables can be set so administrators can be proactively notified of events such as an in-memory table nearly its specified size limit.

Easy Tiered Storage Management

DSE's in-memory option provides architects and DBAs with full tiered storage management for their NoSQL database platform. Data may be assigned to in-memory objects, SSD's, traditional spinning disks, and cloud-based storage platforms all in the same database cluster.

Further Reading

DSE's in-memory option for Cassandra combines the power of in-memory computing with the fastest and most scalable NoSQL database in the market for applications that need both in-memory and big data handling capabilities. For more resources and [downloads](#) of DataStax Enterprise, visit www.datastax.com today.