Advanced Performance with DataStax Enterprise

DATASTAX

CONTENTS

INTRODUCTION: FACTORS BETTER PERFORMANCE THAN OPEN SOURCE APACHE CASSANDRA™ AND APACHE SPARK™

Apache Cassandra is the de facto open source database technology at scale. DataStax Enterprise (DSE) is the world's most simple and powerful distributed cloud database, built on the best distribution of Apache Cassandra. DSE includes all the advantages of Apache Cassandra and Apache Spark with even more power and capabilities, including:

- Double the database performance:
  - **Thread per core and asynchronous architecture:** A coordination-free design, DSE's thread-per-core architecture provides up to 2-3x more throughput and reduced latency.
  - **Storage engine optimizations:** Improved latency and optimized compaction.
  - **DataStax Bulk Loader:** Up to 4x faster than current data loading utilities.
- More powerful and flexible data replication: DSE Advanced Replication delivers more flexibility in how data is replicated, both for operational and performance purposes.
- Faster analytics: Continuous paging improves analytics read performance by up to 3x over open source Apache Cassandra and Apache Spark.

These capabilities increase your teams' productivity and give your applications double the responsiveness and handle twice the throughput with the same hardware.

DOUBLE (OR MORE) THE PERFORMANCE

Thread-Per-Core and Asynchronous Architecture

Apache Cassandra uses a traditional staged event-driven architecture (SEDA). With the SEDA architecture, Apache Cassandra assigns thread pools to events or tasks and connects them via a message service. This architecture also uses multiple threads per task, meaning that threads need to be coordinated. Additionally, events in this architecture are synchronous, which can cause contention and slow downs. Because of this, adding CPU cores eventually sees diminishing returns.

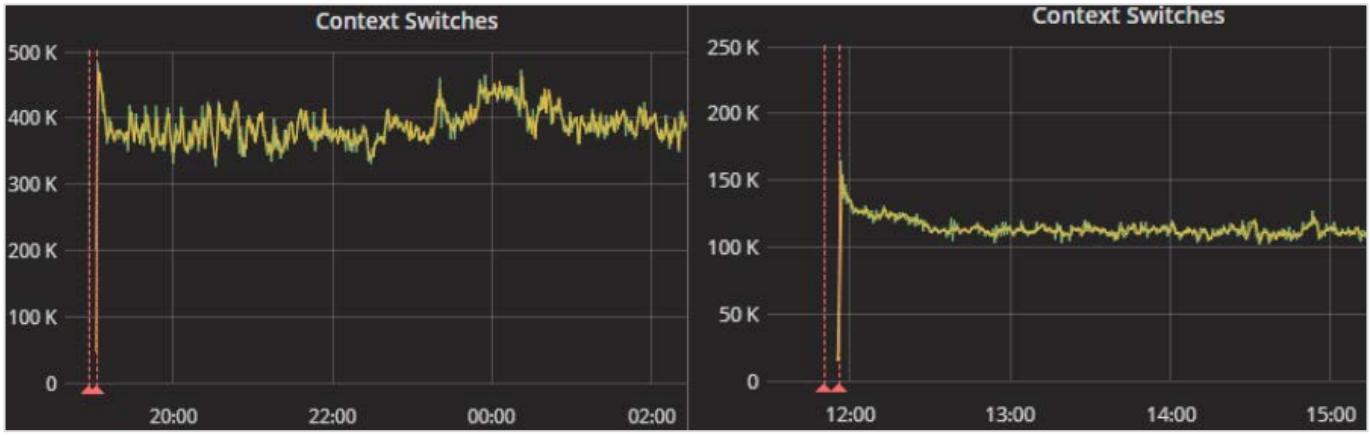

DSE 5.1 left, DSE 6 right - With the traditional SEDA architecture, we see much more context switching which is expensive and degrades performance.

DSE 6 utilizes a new coordination-free design — a thread-per-core architecture that yields incredible performance gains. In an independent benchmark conducted by zData, DSE 6 showed up to 3x throughput improvements and even greater latency reduction which you can see in the graphics below.

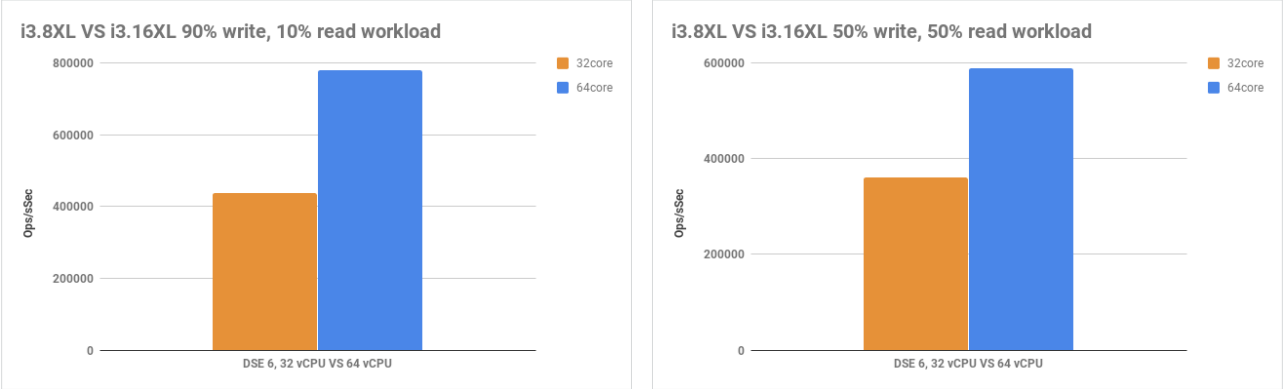

zData Independent Benchmark: 90% write, 10% read on the left. 50% write, 50% read on the right. 5 nodes, i3.8XL, 500GB density.

Each node in a cluster owns part of the token range; that's not new. What's new is that a respective node's token range is divided evenly among CPU threads: one per CPU core to be exact. A respective thread is now responsible for incoming writes for the part of the token range it owns, and any available thread can be used to handle read requests. This means that evenly distributed data in the cluster results in evenly distributed CPU utilization on a server. This architecture also means that very little coordination is needed between threads, which ensures that a CPU core can be used to its fullest capabilities.

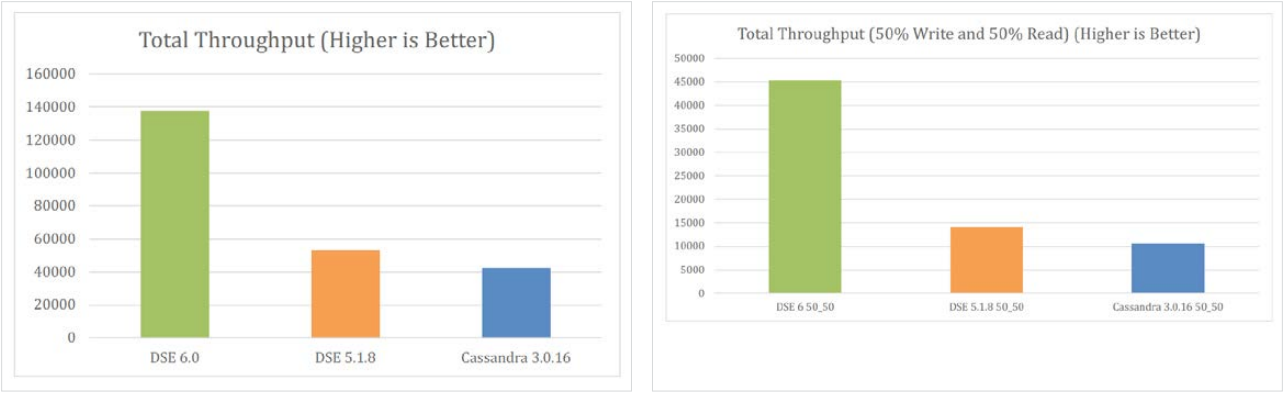

Cassandra-stress, 5 nodes, RF=3, CL=QUORUM, 500GB density

Since a single thread owns the writes for its respective token range, what about contention? In DSE 6, we've moved reads, writes, and other tasks from synchronous to asynchronous operations. This allows us to eliminate thread contention and always keep threads working. Combined with the thread-per-core architecture, this allows us to scale performance linearly as we scale the number of CPU cores. This is extremely important since multi-socket motherboards and high-core-count cloud instances have become the standard today.

Storage Engine Optimizations

Besides ingesting and serving data faster with thread-per-core, we've also made improvements to the storage engine that improve latency and optimize compaction, which can also be a bottleneck for write-heavy workloads. In DSE 6, we see a compaction performance 22% faster than DSE 5.1 which is already 2x faster than open source Apache Cassandra. We're also seeing significant latency improvements on reads and writes.

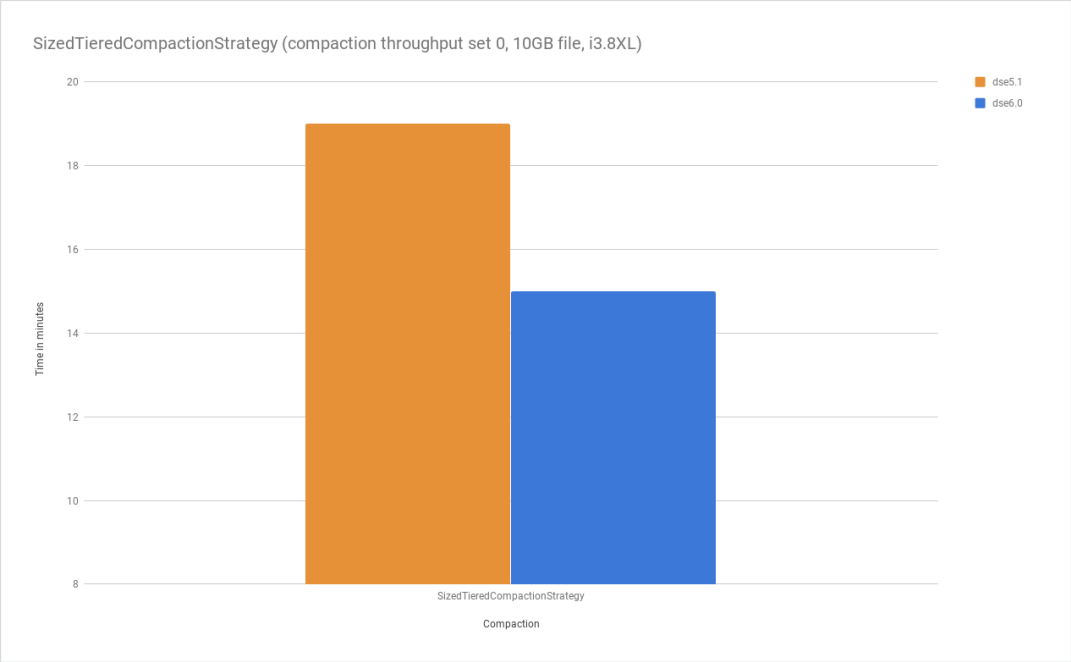

In DSE 5.1, DataStax introduced 2x compaction performance over Apache Cassandra. In DSE 6, compaction is even faster.

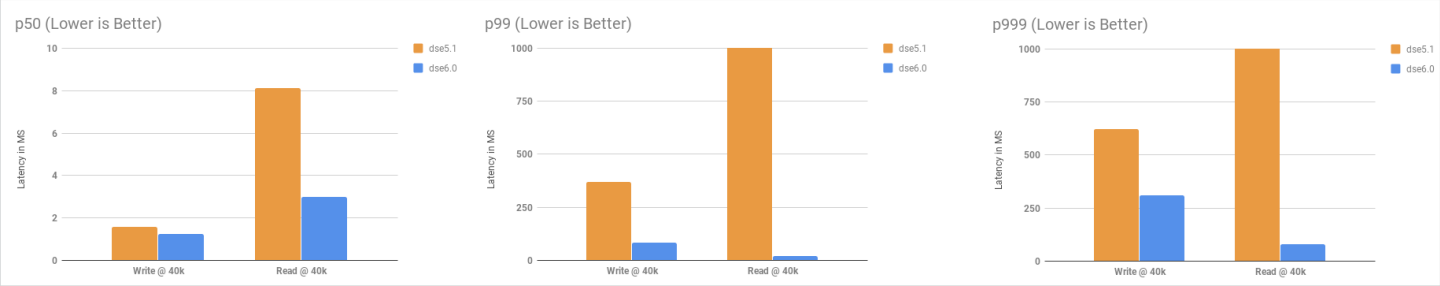

40k fixed throughput test, 3 nodes, RF=3, CL=QUORUM

DataStax Bulk Loader

Loading and unloading large volumes of data into and out of Apache Cassandra can be cumbersome, tedious, and slow. DataStax Bulk Loader simplifies this process and makes it up to 4x faster than current data loading utilities.

MORE POWERFUL AND FLEXIBLE DATA REPLICATION

DSE Advanced Replication

Apache Cassandra is unique in its abilities to support multiple data centers and clouds, and keep data synchronized between geographic locations. However, there are limits to Cassandra's built-in capabilities for both operational and application-specific use cases.

DSE Advanced Replication enables new application possibilities, new data flows, and new operational setups. It connects multiple clusters and keeps data in sync and up-to-date, whereas Cassandra operates within a single cluster. By replicating at the cluster level, DSE can configure workflows such as connecting points of presence at the edge of the enterprise to a central hub with a global view of the entire enterprise. Doing the same from within Cassandra results in a complex setup with subtle yet significant limitations, and in many situations proves prohibitive.

DSE Advanced Replication lets you have a single cluster that can define a primary hub with multiple spokes, allowing configurable distributed data replication from source clusters to destination clusters bi-directionally. It's designed to support microservice analytics commonly found in retail environments and tolerate sporadic connectivity that can occur in constrained environments, such as oil-and-gas remote sites and cruise ships.

The topological characteristics are not the only concern when considering connecting the edge of the enterprise to the central systems. You also need to consider the inherent physical constraints in such a deployment. For example, bandwidth is likely to be restricted, perhaps severely so, in many of these situations. Consider an industrial Internet of Things (IoT) example connecting field equipment to a central hub. The connection between these remote sites and headquarters is likely to go over a channel where bandwidth is either expensive or limited – or both. Moreover, the connection from these remote locations can frequently be unavailable, but applications at the edge continue to produce data. This data needs to be stored at the edge so that the data flow can resume when connectivity is restored, ensuring all messages are eventually delivered. This is also the case for mobile deployments, such as cruise ships, or rugged environments, such as oil drilling sites. In these environments, a system is required that allows data to be prioritized in order to make the best use of limited communications resources.

Data flows are another dimension which is constrained in this "Hub and Spoke" configuration. Consider the case of collecting data at these remote sites and wanting to query the data specific to the edge as well as a global view of data at the hub. Configuring keyspaces in Cassandra to accomplish the data movement is possible, though complex and error-prone, but getting a single view of the data for lookups or for analytic queries is actually prohibitively trickier than it appears.

DataStax Enterprise Advanced Replication takes a very straight-forward approach to solving these limitations in Cassandra. It allows many edge clusters – each acting independently and sized, configured, and deployed appropriately for each edge location – to send data to a central hub cluster. Data flows from an edge cluster can be prioritized to ensure the most important data is sent before the less important data. Inconsistent data connections are fully accommodated at the edge with a "store and forward" approach that will save mutations until connectivity is restored. Any type of workload is supported at both the edge and the hub, allowing for advanced Search and Analytic use cases at remote and central locations.

FASTER ANALYTICS THAN OPEN SOURCE

DSE Analytics improves analytics read performance by 3x over open source Apache Cassandra and Apache Spark. Continuous paging is an optimization specifically for DSE Analytics queries that scan a large portion of data. Tests in a number of scenarios show large performance gains.Whether selecting all columns or some columns, with or without a clustering-column predicate, we see a 2.5 to 3.5x performance improvement.

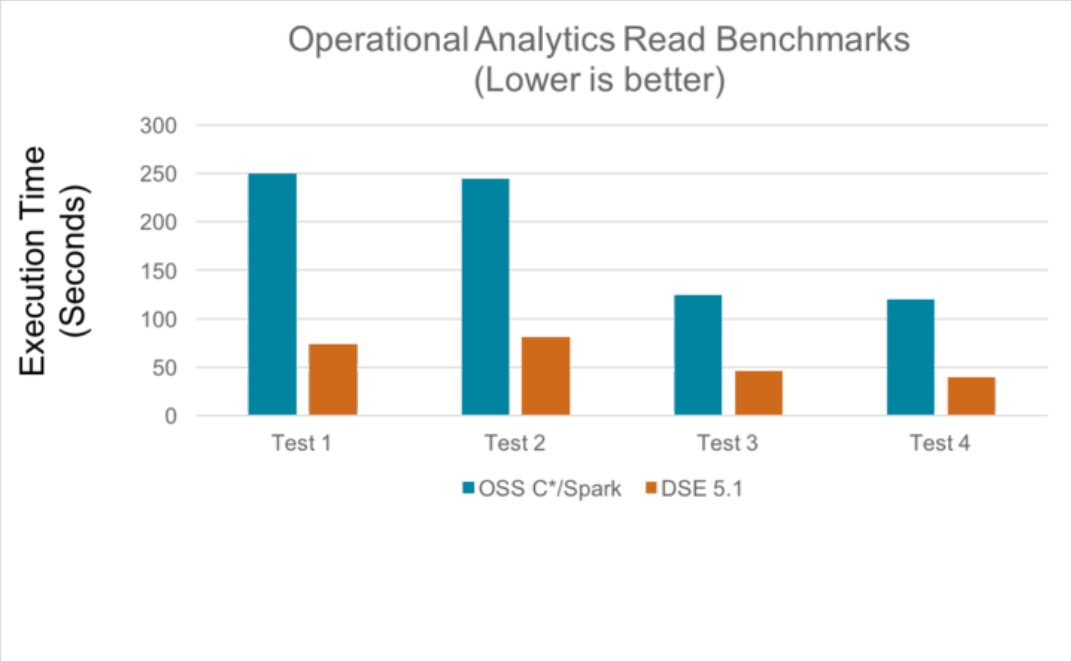

3x analytics read performance over open source Spark and Apache Cassandra.

CONCLUSION

The Advanced Performance capabilities offer significant improvements over open source Apache Cassandra and Apache Spark, all in a unified environment that includes search and a unified security model. With 2x throughput improvements, 2x latency improvements, 22% compaction throughput improvements, 4x loading and unloading of data, and 3x analytics improvements, you get significantly improved performance and productivity on the same infrastructure.

For information about DataStax Enterprise please visit www.datastax.com.

ABOUT DATASTAX

DataStax powers the Right-Now Enterprise with the always-on, distributed cloud database built on Apache Cassandra™ and designed for hybrid cloud. The foundation for real-time applications at massive scale, our flagship product, DataStax Enterprise, makes it possible for companies to exceed expectations through consumer and enterprise applications that provide responsive and meaningful engagement to each customer wherever they go. Our product also gives businesses full data autonomy, allowing them to retain control and strategic ownership of their most valuable asset in a hybrid cloud world. DataStax helps more than 400 of the world's leading brands like Capital One, Cisco, Comcast, eBay, McDonald's, Microsoft, Safeway, Sony, UBS, and Walmart transform their businesses through right-now applications focused on enterprise optimization and customer experience. For more information, visit DataStax.com and follow us on @DataStax.

Advanced Performance with DataStax Enterprise