The Solution for Cassandra at Scale

HGST has a long history of producing the largest and most reliable hard disk drives (HDDs) in the world. Couple those hard drives with the broad SanDisk-brand portfolio of flash storage technologies and you get a solution that enables Cassandra architects to achieve optimal Cassandra clusters. While modern multi-core servers allow parallel execution of multiple threads, Cassandra was not originally written to fully exploit them, which often leaves cores idling, waiting for data from storage. Ultimately, the low utilization causes server sprawl and wasteful spending of IT budgets.

Apache Cassandra™ database at scale can use both the cost-effective capacity of HGST-brand Ultrastar® helium hard drives and the density and performance capabilities of SanDisk®-brand solid-state drives (SSDs) to fully exploit flash and modern servers and to provide optimal performance and consolidation.

What SanDisk- and HGST-brand drives can do for Cassandra:

- Store more data in the same or smaller footprint

- Minimize query and data operation times on critical datasets

- Customer-optimized price and performance

- Reduce or eliminate JBODs and controllers

About Apache Cassandra

Cassandra is an open source NoSQL database written in Java and specifically optimized to be scalable, decentralized, fault tolerant, and, above all, performant. It is used at some of the web's largest properties and throughout financial services and other industries as a repository of record with multiple petabytes of data under its control.

As a NoSQL database, Cassandra was built from the ground up for scale-out architectures. Instead of investing in a large, centralized database server with massive amounts of storage and memory capacity, architects can deploy more modestly configured servers to perform the same types of operations and guarantee the same levels of uptime and data reliability.

Scale-out provides a powerful method for increasing database performance and capacity. Need more compute power? Add servers to distribute the workload. Need additional storage capacity? Add servers and rebalance. Yet all of these server additions, if not properly managed and minimized, can lead to a classic case of server sprawl with massive operational expenses from large and underutilized server farms.

Avoiding Cassandra Server Sprawl for Capacity

As described above, there are basically two reasons to add servers to a cluster: to expand capacity or to increase performance. Let's examine how HGST helium hard drives can help minimize the need for additional servers for petabyte-scale capacities.

Pain Point: Database Server Sprawl, Underutilized CPU

It is an axiom in the computer industry that data always grows to fill available space. This is a good problem to have because additional data enables Cassandra to perform deeper analytics and extract higher value insights from data. However, it can lead to adding servers simply for their storage, effectively wasting the initial cost of the rest of the server and its ongoing power, cooling, and maintenance.

Pain Point: Fixed Rack-Space, Increasing Database Size

In cases where your application is data-limited and not server-compute-limited, it can make sense to scale up your scale-out storage. HGST-brand Ultrastar® helium hard drives, in announced capacities of up to 12TB in an industry-standard 3.5 inch form factor, are offered with a choice of SAS or SATA interface. By loading 4 drives

SanDisk®

in a single rack unit server, nearly 50TB of raw storage and compute can exist in such a server, providing an optimal balance between capacity and compute for less-frequently accessed data.

Avoiding Cassandra Server Sprawl for Performance

SanDisk SSDs are built for performance. Depending on your performance needs, multiple SATA or SAS interfaced SSDs can reduce the I/O wait times dramatically when compared with traditional storage solutions, leading to higher CPU utilization and a decrease in server sprawl.

Pain Point: Database Overhead is Slowing Queries

When a Cassandra cluster is slow to return a response, the cause could be a bottleneck on the underlying storage. Cassandra has an on-disk data format, the SSTable, which is efficient for additions but needs occasional compaction (or garbage collection) as items are updated. When this

compaction takes place, one or more SSTables are consolidated and written into a new file. This process takes I/O performance away from the rest of the application, which is especially troublesome for high-write workloads. Even database reads can be stuck in the I/O queue behind these operations, which means that query performance can drop, sometimes dramatically, while actual server CPU usage will be minimal. A SAS SSD, such as the HGST-brand Ultrastar SS200 with its SAS interface and tuning for a mixed read/write workload, can help alleviate this bottleneck and maintain query performance during background operations.

Pain Point: Power Users Demand More Speed

For the absolute highest performance needs, the SanDisk brand also includes SSDs that completely skip the traditional storage stack by using NVM Express™ (NVMe), a direct-to-CPU attachment technology based on PCI Express that delivers dramatically lower I/O operation latencies than SATA or SAS.

Summary

Cassandra can be a powerful tool to store and extract value from massive amounts of data. However, like any scale-out tool, it needs to be applied carefully and thoughtfully, or it can result in a massive server sprawl and associated headaches.

For the largest Cassandra databases, adding HGST Ultrastar helium HDDs in industry-standard, fully serviceable chassis is ideal. This solution provides high capacity and good performance in a small footprint, and it enables the construction of cost-effective, massive clusters.

Ideal candidates for SanDisk SSDs are Cassandra databases in which queries take too long to return data or in which applications are not meeting their SLAs. In these cases, SanDisk SSDs, potentially in a direct-to-CPU connected NVMe interface form factor, may dramatically reduce query response times and allow you to maintain or reduce your server footprint at massively increased query performance.

Pain Point	SanDisk SSD CloudSpeed™ SATA	SanDisk SSD SkyHawk™ NVMe	HGST SSD Ultrastar® SS200 SAS	HGST SSD Ultrastar® Helium SATA/SAS
Database server sprawl, underutilized CPU	★★			★★★
Fixed rack space, increasing database size	★★★			★★★
Database overhead is slowing queries	★★★	★	★★	
Power users demand more speed	★	★★★	★★	
	Legend: ★ Good ★★ Better ★★★ Best			

Contact Information

SanDisk®

SB18-S-Cassandra-US-EN-0417-01