



Apache Cassandra performance advantages of the new Dell PowerEdge C6620 with Dell PERC 12 RAID controller

The PowerEdge C6620 with PERC 12 delivered lower latency and higher throughput than an HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller

Overview

Today's businesses both generate and take in enormous quantities of data as part of their daily operations. Smartphones, computers, servers, and the activities of people online are the sources of some of this data, but more and more of the data also come from a wide variety of other places, such as weather sensors, streaming video cameras, wearable devices, and onboard computers in vehicles, to name just a few examples. One estimate suggests that the number of connected Internet of Things (IoT) devices will reach over 29 billion by 2030.¹ With an ever-increasing mountain of data, much of it from non-traditional sources, organizations need a way to extract value from the noise.

NoSQL database systems such as Apache Cassandra can help organizations store, process, and analyze this data to glean useful insights. To be most effective, however, the database system should run on a high-performing computing platform that can complete big data workloads quickly and get insights into decision makers' hands fast. We assessed the ability of two platforms to handle Cassandra workloads. The first was the new Dell™ PowerEdge™ C6620 with Broadcom®-based Dell PowerEdge RAID Controller (PERC) 12, which companies might choose if they're upgrading to new servers to better handle their big data needs. The second was the older HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller, which represents a server that organizations might already have in their data centers. The Dell and Broadcom solution provided higher throughput and lower latencies in our testing, meaning that it completed more big data work in the same amount of time as the older HPE solution.

With a strong big data solution, businesses can put their data to work and use it to optimize processes, cut costs, improve customer experience, and grow their offerings. This report explores how and why running Apache Cassandra as a big data system on the Dell PowerEdge C6620 server with PERC 12 might be that solution for you.

About the Dell PowerEdge C6620 server

Part of the Dell modular infrastructure PowerEdge C-Series, Dell says the PowerEdge C6620 is “designed for compute-intensive workloads” but also “ideal for IOPS-heavy workloads.”² It features up to two 4th Generation Intel® Xeon® Scalable processors, with up to 56 cores per processor; offers memory speeds of up to 4,800 MT/s; and supports up to 16 NVMe® drives for workload acceleration. Optional liquid cooling is also available.

To learn more about the Dell PowerEdge C6620, visit <https://www.dell.com/en-us/shop/enterprise-products/c6620-two-socket-server-node-intel/spd/poweredge-c6620>.



Image provided by Dell

Assessing Cassandra performance on the Dell PowerEdge C6620 with Broadcom-based PERC 12

Upgrading to new servers is a big decision. You know that newer, more modern technology is likely to offer performance improvements, but exactly what will those benefits look like, and how much more will the new systems be able to handle? Our testing quantifies the performance boost you might see on your Cassandra workloads by moving from HPE ProLiant XL170r Gen9 servers with Smart Array P440ar controllers to new Dell PowerEdge C6620 servers with Broadcom-based PERC 12.

Our configurations

For our test environment, we installed VMware® vSphere® 8 on both servers before configuring a separate infrastructure server with VMware ESXi™ and VMware vCenter®. We used this infrastructure server to manage the servers and to host client VMs that ran our test workload against our databases. The Dell PowerEdge C6620 server with PERC 12 used two Dell U2 Gen4 NVMe® 3.84TB drives. The HPE ProLiant XL170r Gen9 server with HPE Smart Array P440ar controller used six 960GB mixed-use SAS 12Gbps drives. Table 1 highlights more details of our configuration.

Table 1: System configurations we used in our testing. Source: Principled Technologies.

Server configuration information	HPE ProLiant XL170r Gen9	Dell PowerEdge C6620
Processors	2x Intel Xeon processor E5-2699 v4 22 cores, 2.2GHz	2x Intel Xeon Platinum 8452Y 36 cores, 2GHz
Storage controller	Smart Array P440 Controller, 2GB cache	PERC H965i Adapter, 8GB cache
Disks	6x 960GB Toshiba PX05SVB096Y (12Gb SAS SSDs)	2x 3.84TB Dell Enterprise NVMe v2 AGN RI U.2 (NVMe SSDs)
Total memory in system (GB)	512	
OS and version number	VMware ESXi 8.0.0, 20513097	

On each server, we created a Cassandra gold VM and cloned it five times to create a total of six VMs, which we joined in a cluster configuration. We then used the Yahoo Cloud Serving Benchmark (YCSB) to create a 100GB database across the six VMs to take advantage of the distributed database functionality of Cassandra, ran YCSB workload B for 30 minutes, and recorded the results. In the results we highlight below, we provide two perspectives on the performance of each setup: the total throughput and the average read and write latency. Both results reflect the performance across all six VMs.

Why YCSB?

YCSB is an industry-standard benchmark for NoSQL databases. In 2010, a group from Yahoo! Research created it with “the goal of facilitating performance comparisons of the new generation of cloud data serving systems.”³ It is open source, meaning that anyone can access and modify the source code. In a recent interview, contributors to the YCSB open-source community note that it “is rather largely accepted by users” and “represents a series of scenarios that can be abstracted from the real world.”⁴ Apache Cassandra was one of the first four databases that the YCSB creators tested with the benchmark in 2010, and YCSB remains a good fit for testing Cassandra performance today.⁵

YCSB functions by letting users create a database populated with synthetic data on their database system of choice. Users can then run a pre-defined or customized workload against the database to gauge system performance. YCSB offers six core workloads, each of which represents a different type of database work. Our testing used the read-intensive workload B. This workload is 95 percent reads (pulling data from a database) and 5 percent writes (adding to or changing data in a database). YCSB gives one application example as photo tagging, where a user might occasionally add a tag to a photo (write) but will mostly search a library of tagged photos (read).⁶ A solution that offers higher performance on YCSB workload B is likely to improve performance on other read-intensive workloads, such as data analysis. We chose this workload to focus on reading and analyzing a database.

Upgrade to the Dell PowerEdge C6620 with Broadcom-based PERC 12 for lower latencies and better throughput

In our testing with YCSB, the Dell PowerEdge C6620 with PERC 12 offered better performance on all three metrics we measured: read latency, update (or write) latency, and throughput (measured in operations per second). The improvements were significant, meaning that trading out your HPE ProLiant XL170r Gen9 servers for new PowerEdge C6620 servers could enable your organization to handle substantially more Cassandra work.

The first metrics we examined were read latency, which measures the delay between the application requesting a piece of data and the database system delivering it, and update latency, which measures the delay between the application changing or adding a piece of data and the database system completing the action. The Dell PowerEdge C6620 with PERC 12 was much faster on both types of latency, with the largest advantage on update latency. There, it offered 60.2 percent lower—or 1.97 milliseconds less—latency than the HPE ProLiant XL170r Gen9 server with Smart Array P440ar controller.

It may seem like a sub-two-millisecond delay is inconsequential; if you were loading a webpage or pulling up a video, you wouldn't notice a two-millisecond difference. The significance of this advantage, however, is due to the enormous number of operations that the database system must perform before it can deliver usable results. For this testing in YCSB, we set the `maxexecutiontime` variable (or how long the benchmark should run) to 30 minutes. At a rate of 249,210 operations per second (see Figure 3), the Dell PowerEdge C6620 with Broadcom-based PERC 12 executed over 400 million operations during the 30-minute test. So, while a difference of one or two milliseconds might not mean much on a single operation, on 400 million operations, the benefit of the faster solution becomes clear.

About the Dell PERC 12 RAID controller

The Dell PowerEdge C6620 we tested features the PERC 12, which offers a single front controller with full RAID support for both NVMe and SAS.⁷ It brings 3,200MHz cache memory speed and a 16-lane host bus type and supports RAID levels 0, 1, 5, 6, 10, 50, and 60.⁸

The Dell PERC 12 is based on the Broadcom SAS4116W series chip. According to Broadcom, "this eighth-generation SAS RAID-on-Chip (ROC) is based on the industry-leading Fusion-MPT architecture and features Tri-Mode SerDes technology that enables a seamless operation of up to 16-wide direct-connect NVMe, SAS or SATA storage devices from any system design.... The Tri-Mode ROC device with 16-wide PCIe Gen 4.0 lanes provides SAS data transfer rates of 22.5, 12, 6Gb/s per lane and 6Gb/s SATA data transfer rates per lane. The high-port count ROC helps eliminate storage bottlenecks with support of x8, x4, x2, and x1 PCI Express® lanes and complies with the PCIe 4.0 specification, offering up to 6 million IOPS (random reads) and up to 900,000 IOPS in RAID (random writes)."⁹

To learn more about the Dell PERC 12, visit <https://infohub.delltechnologies.com/p/dell-poweredge-raid-controller-12/>.

With these lower latencies, a solution will be quicker to handle interactions with the Cassandra database, which might include anything from pulling up X-ray images in a hospital to analyzing a large set of data on an ecommerce business’s customer preferences.

Average read latency on YCSB workload B

Milliseconds | Lower is better




Figure 1: Average read latency the two solutions achieved on YCSB workload B. Lower is better. Source: Principled Technologies.

Average update latency on YCSB workload B

Milliseconds | Lower is better



Figure 2: Average update latency the two solutions achieved on YCSB workload B. Lower is better. Source: Principled Technologies.

The Dell PowerEdge C6620 also offered an enormous advantage on throughput, delivering over twice the operations per second of the HPE ProLiant XL170r Gen9. Given the lower latencies we saw, this is unsurprising—because the PowerEdge C6620 could process operations faster (with lower latency), we would expect it to also be able to handle more operations in a given time. Depending on the read-intensive workloads you’re running, this increase in throughput could translate to quicker load times for your customers or faster data analysis, among other possibilities.

Total operations per second on YCSB workload B

Higher is better




Figure 3: Total operations per second the two solutions achieved on YCSB workload B. Higher is better. Source: Principled Technologies.

NoSQL databases and Cassandra in today's business landscape

For this study, we tested with Apache Cassandra, a widely used NoSQL database system. NoSQL, or non-relational, databases are a category of database system that store and query data that do not have a traditional data structure. Traditional SQL databases organize data in a column-row format for finding or creating relationships across the data. To store data in a SQL database, all data in each table must have the same structure and fit a pre-defined schema, with every row in each table including the same columns and formats every time. NoSQL databases, however, can organize data more dynamically. They can deal with data from documents, graphs, key-values, and more. This flexibility lets people use them to analyze documents or data that don't follow identical structuring formats. For organizations that need to store and analyze unstructured data—which may include data from Internet of Things (IoT) applications, audio, video, text files, social media posts, and more—a NoSQL database is a great option.

There are many types of NoSQL database systems; Apache Cassandra is a type of key-value and wide-column store. These databases have essentially two fields: One is the key, and the other is the value. The value can be any type of data (text, numbers, etc.). Taking our previous example, a key-value database could have some keys that correspond to a date, others that are numbers, and so on. A wide-column database, which Cassandra uses, is a two-dimensional key-value database, where instead of mapping to just one value, the keys can map to several columns of values.

Apache Cassandra is a distributed database, meaning that it can run on multiple nodes while acting as a single entity. This makes it resilient and highly scalable. Its scalability, combined with the flexibility afforded by its hybrid key-value/tabular model, allows it to handle many types of big data work very well. Cassandra is also open-source and free, a compelling benefit for organizations seeking to save on licensing fees.

The flexibility of Cassandra makes it suitable for a very large range of use cases. For example, Instagram uses Cassandra to support its content feed, Spotify uses it to store playlist metadata, and Intuit uses it as part of their largest production clusters supporting TurboTax.^{10,11,12} Common uses of Cassandra include:

- Analysis of customer data for personalization and recommendation, such as in ecommerce environments and content sharing or streaming websites
- Storage and analysis of IoT data, such as data gathered from mobile and wearable devices, environmental sensors, and edge devices
- Fraud detection, especially for financial organizations
- Messaging, such as for organizations' internal messaging platforms

We chose to test with Cassandra in part because so many organizations rely on it for everyday operations. Approximately 90 percent of Fortune 100 companies use Apache Cassandra in some capacity.¹³ If your organization uses Cassandra or is considering doing so, to get the most value from it, you will want to ensure that the solution backing your implementation offers high performance. As our testing highlights, the Dell PowerEdge C6620 with Broadcom-based PERC 12 can deliver just that.

Dell PowerEdge servers: A proven history of strong Apache Cassandra performance

In this study, we tested the Apache Cassandra performance of a new Dell PowerEdge C6620 server compared to an HPE ProLiant XL170r Gen9 server, but this isn't the first time we've seen strong Cassandra performance on a latest-generation Dell server.

In 2019, we tested Apache Cassandra performance on a 14th generation PowerEdge C-series server, the Dell EMC PowerEdge C6420. Pitted against an older modular solution of HPE ProLiant XL170r Gen9 server nodes, the PowerEdge C6420 accomplished double the amount of work in the same amount of rack space.¹⁴ Two years prior, in 2017, we assessed a different product line from the 14th generation of PowerEdge servers—the Dell EMC PowerEdge FC640 server—and found that it delivered dramatically more throughput and consistently lower latency than a legacy solution of PowerEdge R710 servers.¹⁵

Conclusion

Data proliferation today is rapid, and its growth shows no signs of stopping. For businesses that can take advantage of that data, there is tremendous potential value. One recent McKinsey study notes that “companies that are using data-driven B2B sales-growth engines report above-market growth and EBITDA increases in the range of 15 to 25 percent.”¹⁶ With data flooding in so quickly and in so many different forms, however, companies need high-performing big data solutions to have a chance at utilizing that data effectively.

We tested the performance of two platforms with a read-intensive Apache Cassandra database system big-data workload to assess which might be better suited to speedily deliver the insights decision makers need. Compared to an older HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller, the new Dell PowerEdge C6620 with Broadcom-based PERC 12 RAID controller delivered faster read and update latencies and more than twice the throughput. This improvement in performance can help you glean more value from your unstructured data more quickly. If you're watching your stores of unstructured data grow but are still leaning on older servers for your critical Cassandra workloads, it may be time for an upgrade.

1. Lionel Sujay Vailshery, "Number of Internet of Things (IoT) connected devices worldwide from 2019 to 2021, in forecasts from 2022 to 2030," accessed July 13, 2023, <https://www.statista.com/statistics/1183457/iot-connected-devices-worldwide/>.
2. "PowerEdge C6620," accessed June 23, 2023, <https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-c6620-spec-sheet.pdf>.
3. Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, Russell Sears, "Benchmarking Cloud Serving Systems with YCSB," accessed June 23, 2023, <https://courses.cs.duke.edu/fall13/compsci590.4/838-CloudPapers/ycsb.pdf>.
4. "The Ultimate YCSB Benchmark Guide (2021)," accessed June 23, 2023, <https://benchant.com/blog/ycsb>.
5. Brian F. Cooper, Adam Silberstein, Erwin Tam, Raghu Ramakrishnan, Russell Sears, "Benchmarking Cloud Serving Systems with YCSB," accessed June 23, 2023, <https://courses.cs.duke.edu/fall13/compsci590.4/838-CloudPapers/ycsb.pdf>.
6. "brianfrankcooper/YCSB," accessed June 23, 2023, <https://github.com/brianfrankcooper/YCSB/blob/master/doc/coreworkloads.html>.
7. "Dell PowerEdge RAID Controller 12 User's Guide PERC H965i Adapter, PERC H965i Front, and PERC H965i MX," accessed June 27, 2023, <https://www.dell.com/support/manuals/en-us/perc-h965i-front/perc12/dell-technologies-poweredge-raid-controller-12?guid=guid-5889415d-b297-43a0-9197-113a56c33c79&lang=en-us>.
8. "SAS4116W 24G SAS Tri-Mode RAID-on-Chip (ROC)," accessed June 27, 2023, <https://www.broadcom.com/products/storage/raid-on-chip/sas-4116w>.
9. "SAS4116W 24G SAS Tri-Mode RAID-on-Chip (ROC)."
10. Instagram Engineering, "Open-sourcing a 10x reduction in Apache Cassandra tail latency," accessed June 27, 2023, <https://instagram-engineering.com/open-sourcing-a-10x-reduction-in-apache-cassandra-tail-latency-d-64f86b43589>.
11. Kinshuk Mishra and Matt Brown, "Personalization at Spotify using Cassandra," accessed June 27, 2023, <https://engineering.atspotify.com/2015/01/personalization-at-spotify-using-cassandra/>.
12. Denson Pokta, "Pronto! Intuit Releases First Open Source Cassandra Cluster Manager," accessed June 27, 2023, <https://thenewstack.io/pronto-intuit-releases-first-open-source-cassandra-cluster-manager/>.
13. Jeff Carpenter, "How the world caught up with Apache Cassandra," accessed June 27, 2023, <https://techcrunch.com/sponsor/datastax/how-the-world-caught-up-with-apache-cassandra/>.
14. "Move your private cloud to Dell EMC PowerEdge C6420 server nodes and boost Apache Cassandra database analysis," accessed June 23, 2023, <https://www.principledtechnologies.com/Dell/Power-Edge-C6420-Apache-Cassandra-1019-v2.pdf>.
15. "Update your private cloud with 14th generation Dell EMC PowerEdge FC640 servers and do more work in less space," accessed June 23, 2023, https://www.principledtechnologies.com/Dell/PowerEdge_FX2s_FC640_Apache_Cassandra_1117.pdf.
16. Jochen Böringer, Alexander Dierks, Isabel Huber, and Dennis Spillecke, "Insights to impact: Creating and sustaining data-driven commercial growth," accessed July 13, 2023, <https://www.mckinsey.com/capabilities/growthmarketing-and-sales/our-insights/insights-to-impact-creating-and-sustaining-data-driven-commercial-growth>.

Read the science behind this report at <https://facts.pt/cGgc9vo> ▶



Facts matter.®

Principled Technologies is a registered trademark of Principled Technologies, Inc. All other product names are the trademarks of their respective owners. For additional information, review the science behind this report.

This project was commissioned by Dell Technologies.