Speeding time to insight: The Dell PowerEdge C6620 with Dell PERC 12 RAID controller for Apache Cassandra big data workloads

The new PowerEdge C6620 delivered better performance—both higher throughput and lower latency—than a previous-generation PowerEdge C6520 with PERC 11

Overview

Every day, individuals and organizations generate massive quantities of data, from text messages to location data to information from sensors on factory floors and beyond. This rapid proliferation of data offers enormous opportunities: If businesses can extract insights from that data, they can use it to improve their operations, grow their customer base, and provide a better experience to those customers.

That task is not simple, however. Much of this data is unstructured, meaning that it comes in many formats that traditional data models, such as SQL databases, cannot process. Processing and analyzing unstructured data may require different methods, such as utilizing a NoSQL database like Apache® Cassandra®. Organizations can use NoSQL databases to store, mine, and analyze unstructured data in its many forms and gain actionable information.

To efficiently analyze such large quantities of data, however, they need a powerful computing solution running the database system. Investing in newer server solutions with updated processing, storage, and networking components can offer greater performance and enable companies to get to those vital insights faster. To highlight the advantages of moving from an older server solution to a new one for big data workloads, we tested Apache Cassandra performance on a new Dell™ PowerEdge™ C6620 with a Broadcom®-based Dell PowerEdge RAID Controller (PERC) 12 and an older Dell PowerEdge C6520 with Dell PERC 11. On multiple performance metrics, the newer Dell PowerEdge C6620 with PERC 12 delivered stronger performance than its predecessor, offering businesses the chance to increase the value of their data and realize its benefits more quickly.

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."[1] 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

Testing the Dell PowerEdge C6620 with Broadcom-based PERC 12

If you're still relying on servers you purchased several years ago, it can be helpful to understand exactly how much you could gain by upgrading to a newer solution. We designed our testing to quantify the benefits of upgrading from older to latest-generation servers for organizations relying on Cassandra workloads for critical operations.

Our configurations

To set up our test environment, we installed VMware® vSphere® 8 on both servers. We then configured a separate infrastructure server with VMware ESXi™ and VMware vCenter® to manage the servers and to host client VMs that ran our test workload against our databases. The Dell PowerEdge C6620 server with Broadcom-based PERC 12 used two Dell U2 Gen4 NVMe® 3.84TB drives, while the Dell PowerEdge C6520 server with PERC 11 used six 960GB mixed-use SAS 12Gbps SFF drives. (See Table 1 for more details of our configuration.)

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

Server configuration information	Dell PowerEdge C6520	Dell PowerEdge C6620
Processors	2x Intel Xeon Gold 6330 28 cores, 2GHz	2x Intel Xeon Platinum 8452Y 36 cores, 2GHz
Storage controller	PERC H750 Adapter, 8GB 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."[2] 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."[3] 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.[4]

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).[5] 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.

See higher throughput and lower latency with the Dell PowerEdge C6620 with Broadcom-based PERC 12

Our testing with YCSB yielded three metrics: read latency, update (or write) latency, and throughput (measured in operations per second). The Dell PowerEdge C6620 with Broadcom-based PERC 12 offered stronger performance than the PowerEdge C6520 with PERC 11 on all three metrics, indicating that an upgrade can help speed your Cassandra workloads.

On the first and second metrics, read latency and update latency, the Dell PowerEdge C6620 was significantly faster than its previous-generation counterpart. Read latency measures the delay between the application requesting a piece of data and the database system delivering it; update latency measures the delay between the application changing or adding a piece of data and the database system completing the action. The shorter these delays, the faster a solution will be at completing user-facing requests, such as retrieving a customer's buying history when a store manager searches for it, and larger workloads, such as running analysis on a set of tens of thousands of data points.

On the surface, the differences in latency between the two solutions are very small: 0.49 milliseconds for read latency and 0.57 milliseconds for update latency. On a single operation, a delay of less than a millisecond would be impossible for a human to notice. But the database system isn't handling just one operation—it's handling thousands or millions of operations all at once. Our YCSB testing, for example, set the `maxexecutiontime` variable (or how long the benchmark should run) to 30 minutes. This means that at the Dell PowerEdge C6620 server's rate of 249,210 operations per second (which we show in Figure 3), it executed over 400 million operations during the 30-minute test. As tiny differences in latency scales up, they become very significant indeed. And the shorter these delays, the faster a solution will be at completing both user-facing requests, such as retrieving a customer's buying history when a store manager searches for it, and larger workloads, such as running analysis on a set of tens of thousands of data points.

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.[6] 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.[7]

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)."[8]

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

Average read latency on YCSB workload B
Milliseconds | Lower is better

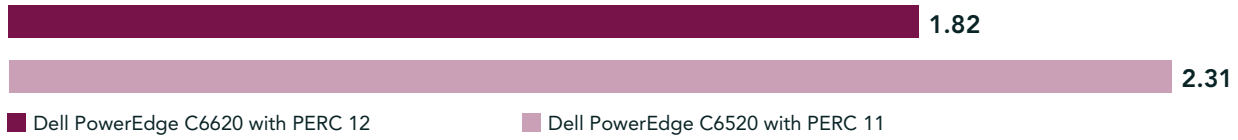

Figure 1: Average read latency the two solutions achieved on YCSB workload B. Lower is better. Source: Principled Technologies.

Average update latency on YCSB workload B
Milliseconds | Lower is better

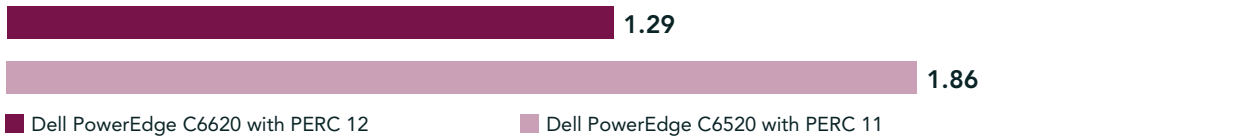

Figure 2: Average update latency the two solutions achieved on YCSB workload B. Lower is better. Source: Principled Technologies.

On the third metric, throughput, the Dell PowerEdge C6620 delivered 1.25 times as many operations per second as the previous-generation PowerEdge C6520. This increase in throughput is what we would expect to see based on the lower latencies: If a system is able to process operations faster (i.e., with lower latency), it will also boost how many operations the system can handle in a given time (i.e., better throughput). With greater throughput, depending on what read-intensive workloads your organization is running, you might see faster video streaming, quicker recommendations for customers, or an increase in the speed of users pulling up data.

Total operations per second on YCSB workload B
Higher is better

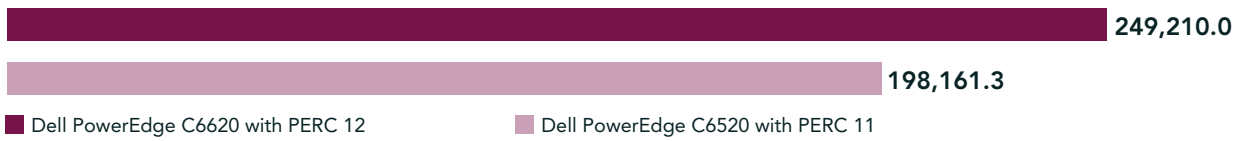

Figure 3: Total operations per second the two solutions achieved on YCSB workload B. Higher is better. Source: Principled Technologies.

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.[9,10,11] 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.[12] 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.[13] 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.[14]

Conclusion

The vast amounts of unstructured data that people and organizations generate daily have the potential to bring incredible value to companies that can utilize it quickly and correctly. Buried in the data are insights about consumer preferences, product performance, environmental trends, and more—but to access those insights at the speed of business, you need high-performing NoSQL databases. Aging servers may be holding you back from the full value of your data.

We found that the new Dell PowerEdge C6620 with Broadcom-based PERC 12 RAID controller can speed read-intensive Apache Cassandra database workloads compared to an older server solution. Faster read and update latencies and higher throughput, as we saw the PowerEdge C6620 deliver, can speed the retrieval, processing, and analysis of your unstructured data, enabling you to more effectively extract its value. To more fully utilize your data to inform your everyday business operations, consider the Dell PowerEdge C6620 with Broadcom-based PERC 12 RAID controller.

1. "PowerEdge C6620," accessed June 23, 2023, https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/poweredge-c6620-spec-sheet.pdf.
2. 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.
3. "The Ultimate YCSB Benchmark Guide (2021)," accessed June 23, 2023, https://benchant.com/blog/ycsb.
4. 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.
5. "brianfrankcooper/YCSB," accessed June 23, 2023, https://github.com/brianfrankcooper/YCSB/blob/master/doc/coreworkloads.html.
6. "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.
7. "SAS4116W 24G SAS Tri-Mode RAID-on-Chip (ROC)," accessed June 27, 2023, https://www.broadcom.com/products/storage/raid-on-chip/sas-4116w.
8. "SAS4116W 24G SAS Tri-Mode RAID-on-Chip (ROC)."
9. 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-latencyd-64f86b43589.
10. 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/.
11. 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/.
12. 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/.
13. "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.
14. "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.

Read the science behind this report at https://facts.pt/A5YAzMh ▶

Principled Technologies®

Facts matter.®

This project was commissioned by Dell Technologies.