



Image provided by Dell



## Process more data

2.1 times the YCSB operations per second\*

## Dell PowerEdge C6620 server with Dell PowerEdge RAID Controller (PERC 12) analyzed Cassandra database data more quickly

than an HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller

As the speed of business accelerates, catching anomalies in data requires quick access to and analysis of collected data. Organizations using Apache® Cassandra® distributed database workloads can improve performance—and put detailed insights into the hands of decision makers earlier—by selecting the right hardware platform.

At Principled Technologies, we used Yahoo Cloud Serving Benchmark (YCSB) workload B to measure read-intensive Apache Cassandra distributed database performance of two server solutions in virtualized environments:

- A Dell™ PowerEdge™ C6620 server with a Dell PERC 12 RAID controller
- An HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller

The PowerEdge C6620 server featuring the Dell PERC 12 RAID controller delivered 2.1 times as many YCSB operations per second as the HPE solution with 60.2 percent lower application update latency. A hardware solution that can process operations at a speedier pace lets you catch anomalies sooner, which can lead to better business decisions.

\*Dell PowerEdge C6620 server with a Dell PERC 12 vs. HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller

## Our test approach

The Apache organization describes Cassandra as “an open source NoSQL distributed database trusted by thousands of companies for scalability and high availability without compromising performance. Linear scalability and proven fault-tolerance on commodity hardware or cloud infrastructure make it the perfect platform for mission-critical data.”<sup>2</sup>

We investigated the Cassandra distributed database capabilities of two environments, one with the Dell PowerEdge C6620 server with a Dell PERC 12 RAID controller and one with the HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller. The PowerEdge C6620 used two Dell U2 Gen4 3.84TB drives, while the HPE solution used six mixed-use SAS 12Gbps drives. Both environments used six Apache Cassandra database VMs joined in a cluster configuration.

To generate the test load and measure performance, we used YCSB workload B on each solution from a group of YCSB driver VMs. The Dell PowerEdge C6620 solution with a PERC 12 RAID controller delivered 2.1 times as many operations per second (OPS) as the HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller and achieved 60.2 percent lower application update latency. For an organization that chose the Dell solution, this improvement could translate to finding problems in data sooner, which could, in turn, lead to more effective business decisions. This could also result in savings from needing less hardware to do a given amount of work.



Image provided by Dell<sup>3</sup>

## About the Dell PowerEdge C6620 server

Part of the Dell modular infrastructure PowerEdge C-Series, the Dell PowerEdge C6620 features up to two 4<sup>th</sup> Generation Intel® Xeon® Scalable processors, with up to 56 cores per processor.

The PowerEdge C6620 offers memory speeds of up to 4,800 MT/s and supports up to 16 NVMe drives for workload acceleration. According to initial specification sheets we received from Dell, the PowerEdge C6620 provides flexible I/O options, including several PCIe slots and SNAP I/O support.

To learn more about the Dell PowerEdge C6620, check out information about the PowerEdge C-Series at <https://www.dell.com/en-us/shop/modular-infrastructure/sf/modular-infrastructure>.

## About Yahoo Cloud Serving Benchmark

According to Yahoo!, "the goal of the Yahoo Cloud Serving Benchmark (YCSB) project is to develop a framework and common set of workloads for evaluating the performance of different 'key-value' and 'cloud' serving stores."<sup>4</sup> The benchmark serves many databases including Apache HBase and Cassandra, two NoSQL databases that can handle large datasets.

For more information on YCSB, visit <https://research.yahoo.com/news/yahoo-cloud-serving-benchmark>.

## Put more detailed insights into the hands of decision-makers earlier

In our testing with YCSB workload B, the Dell PowerEdge C6620 server with a Dell PERC 12 RAID controller delivered 2.1 times as many operations per second as the HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller and achieved 60.2 percent lower application update latency. The increased speed on this type of Apache Cassandra workload could mean retrieving data insights earlier.

### Total operations per second on YCSB workload B

Higher is better



Figure 1: Total operations per second the two solutions achieved on YCSB workload B. Higher is better. Source: Principled Technologies.

### Average read latency on YCSB workload B

Milliseconds | Lower is better



Figure 2: Average read latency on YCSB workload B. Lower is better. Source: Principled Technologies.

### Average update latency on YCSB workload B

Milliseconds | Lower is better



Figure 3: Average update latency on YCSB workload B. Lower is better. Source: Principled Technologies.

# About the Dell PERC 12 RAID controller

The latest series of Dell RAID controllers, the PERC 12 Series, offers expanded support and capabilities compared to previous versions. According to documentation we received from Dell, PERC 12 RAID controllers offer support for 24Gbps SAS drives, have 3,200 MHz cache memory speed, support the 16-lane host bus type, and support both NVMe and SAS on the front controller.

To learn more about the PERC 12, visit <https://www.dell.com/support/kbdoc/en-us/000131648/list-of-poweredge-raid-controller-perc-types-for-dell-emc-systems>.



## Conclusion

Detailed, up-to-the-moment data analysis is becoming increasingly important for those who make business decisions. In our hands-on tests, a Dell PowerEdge C6620 server with a Dell PERC 12 RAID controller handled more operations per second for data analysis Cassandra workloads than an HPE ProLiant XL170r Gen9 server with an HPE Smart Array P440ar controller. By selecting the Dell solution, you could give the decision-makers in your organization earlier access to in-depth, actionable intelligence, by finding anomalies quickly in unstructured data.

1. Dell provided the image showing a fully populated Dell PowerEdge C6600 chassis. Our Dell PowerEdge C6600 chassis included four Dell PowerEdge C6620 blades and eight disks. We conducted our testing on one blade and two disks.
2. Apache, "Apache Cassandra Open Source NoSQL Database," accessed November 16, 2022, [https://cassandra.apache.org/\\_/index.html](https://cassandra.apache.org/_/index.html).
3. Dell provided the image showing a fully populated Dell PowerEdge C6600 chassis. Our Dell PowerEdge C6600 chassis included four Dell PowerEdge C6620 blades and eight disks. We conducted our testing on one blade and two disks.
4. "Yahoo Cloud Serving Benchmark," accessed November 14, 2022, <https://research.yahoo.com/news/yahoo-cloud-serving-benchmark>.

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.