



Get insight from document-based distributed MongoDB databases sooner and have CPU headroom for additional data analysis workloads

With additional drive bays and 2nd Generation Intel Xeon Scalable processors, Dell EMC PowerEdge R640 servers handled more Yahoo Cloud Serving Benchmark (YCSB) operations per second than previous-generation servers and handled them more efficiently

Data analysis from distributed MongoDB databases can help improve operational efficiency and services by unifying information from a range of sources. The sooner you can get analyses into the hands of decision makers, the better informed their choices will be. Older servers running data analysis workloads can slow the decision-making process and might not be able to support necessary big data growth.

You can help decision makers get useful insight sooner by moving MongoDB data analysis workloads to current-generation servers. The bump in performance can allow more business units to generate in-depth analyses in less time.

Running read-intensive big data workloads in our data center, a cluster of three current-generation Dell EMC™ PowerEdge™ R640 servers powered by 2nd Generation Intel® Xeon® Scalable processors outperformed a cluster of three previous-generation Dell EMC PowerEdge R630 servers. The current-generation solution had lower processor utilization as well, which could leave headroom to run more data analysis workloads as needs grow.



Analyze data faster...

36% more YCSB operations per second



...while leaving headroom for additional workloads on the system

Improve performance and increase processor efficiency



How additional storage and newer processors help your organization

New generations of servers typically introduce improved or new technology, such as greater storage capacity, increased network speeds, and faster processors. These upgrades in server technology contribute significantly to performance gains in the data center. For example, more processing power allows more users to access and use data.

Compared to older servers, current-generation servers sometimes feature design improvements that promote better performance. Current-generation Dell EMC PowerEdge R640 servers have two more drive bays than previous-generation Dell EMC PowerEdge R630 servers. The added drive bays can expand the server's storage capacity with SATA, SAS, and NVMe SSDs. As both the previous-generation and current-generation servers used in our testing fit in 1U of rack space, replacing the older servers with newer ones would increase storage capacity in the same amount of rack space, which increases the overall storage density of your data center and helps limit data center sprawl.

About 2nd Generation Intel Xeon Scalable processors

The latest from Intel, the 2nd Generation Intel Xeon Scalable processor platform features a wide range of processors to support the workloads you run, including Bronze, Silver, Gold, and Platinum. According to Intel, the 2nd Generation Intel Xeon Scalable platform can handle a variety of workloads, including enterprise, cloud, HPC, storage, and communications.¹ This new processor line also supports a new memory and storage technology to further accelerate workloads, Intel Optane DC persistent memory.


To learn more about the 2nd Generation Intel Xeon Scalable processor family, visit <https://www.intel.com/content/www/us/en/products/docs/processors/xeon/2nd-gen-xeon-scalable-processors-brief.html>.

About the Dell EMC PowerEdge R640

The Dell EMC PowerEdge R640 is a dense 1U, two-socket server. It features 24 DDR4 DIMM slots and up to 76.8 TB of storage.

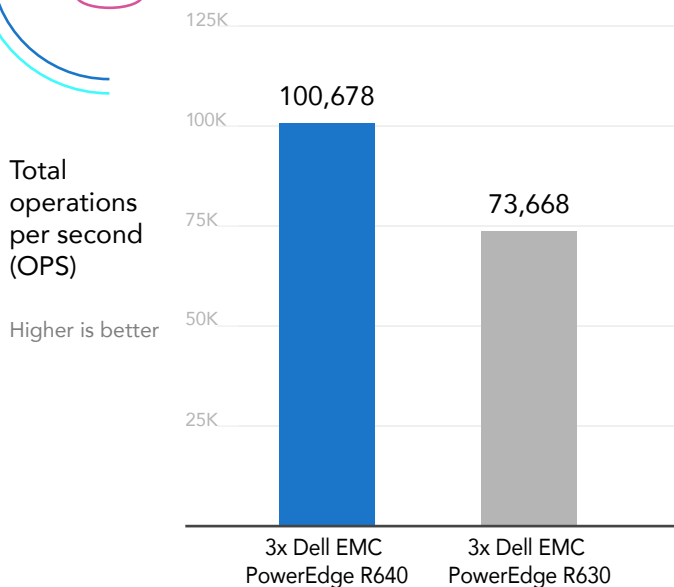
To learn more about the Dell EMC PowerEdge R640, visit <https://www.dell.com/en-us/work/shop/povw/poweredge-r640>.





Analyze data faster

36% more YCSB operations per second



Gain insight from document-based distributed data sooner

We ran YCSB workload B on each three-node solution with three YCSB driver VMs to generate the test load. The current-generation Dell EMC PowerEdge R640 servers, powered by 2nd Generation Intel Xeon Scalable processors and using two Dell EMC Express Flash NVMe P4610 SSDs in the additional drive bays, delivered 36 percent more YCSB operations per second (OPS) than the previous-generation solution. The performance boost could allow more of your business units to run in-depth analysis without experiencing slowdowns or bottlenecks while also helping your organization accumulate data.

Have more CPU headroom to add workloads

The added drive bays that the Dell EMC PowerEdge R640 offers could allow your organization to boost data analysis workload performance with additional SSDs. For our tests, we configured each server in both solutions with eight SATA SSDs and then increased the current-generation solution's storage capacity by adding two Dell EMC Express Flash NVMe P4610 SSDs in the available drive bays of each server.


NVMe SSDs typically increase processor utilization, but the NVMe drives we added to the current-generation servers didn't affect processor utilization when we ran the data analysis workload in the MongoDB database. The combination of 2nd Generation Intel Xeon Scalable processors and additional Dell EMC NVMe SSDs in each current-generation server delivered better performance for the data analytics workload than the previous-generation solution while utilizing a lower overall percentage of the processors. By replacing older servers with the current-generation solution running Dell EMC NVMe SSDs, your organization could improve the performance of its data analysis workloads now while simultaneously creating headroom to run additional workloads. Adding workloads could allow more users to run data analysis or allow your organization to save power, cooling, and licensing expenses by consolidating to fewer current-generation servers.



Create headroom for additional workloads on the system.

Improve performance and increase processor efficiency

Average CPU utilization



Conclusion

Data-driven decisions require speedy data analysis. You could be slowing the decision-making process if you continue to use older servers to run your data analysis workloads. We found that current-generation Dell EMC PowerEdge R640 servers with 2nd Generation Intel Xeon Scalable processors handled more YCSB OPS than previous-generation servers. Replacing previous-generation Dell EMC PowerEdge R630 servers with current-generation PowerEdge R640 servers can deliver a performance boost that allows more of your business units to run in-depth analysis more quickly. In addition, the current-generation servers had more storage drive bays and adding Dell EMC Express Flash NVMe P4610 SSDs to increase the solution's capacity did not affect the level of server processor utilization. With more processor headroom, you could run additional workloads in the cluster, allowing more business units to run data analysis.



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."² 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>.

- 1 Intel, "2nd Gen Intel Xeon Scalable Processors Brief," accessed October 8, 2019, <https://www.intel.com/content/www/us/en/products/docs/processors/xeon/2nd-gen-xeon-scalable-processors-brief.html>.
- 2 "Yahoo Cloud Serving Benchmark," accessed October 8, 2019, <https://research.yahoo.com/news/yahoo-cloud-serving-benchmark>.

Read the science behind this report at <http://facts.pt/4m6apen> ►



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 EMC.