



Crunch data in up to 40% less time



The PowerEdge R7525 completed a k-means analytics workload in 2h 8m while the HPE server\* took 3h 34m

Process up to 67% more data per second



The PowerEdge R7525 processed 121,241 KB/s while the HPE server\* handled 72,518

74% better performance/cost ratio



The PowerEdge R7525 had a performance/cost ratio of 3.15 KB/s per dollar while the HPE server's\* ratio was 1.81 KB/s per dollar

## Get better k-means analytics workload performance for your money with the Dell EMC PowerEdge R7525

The solution beat an HPE ProLiant DL380 Gen10 server by achieving faster analysis at a similar cost

At Principled Technologies, we wanted to gauge the level of big data analytics performance we could get out of servers that cost roughly \$40,000 each. To that end, we tested two current-generation, dual-socket servers: a Dell EMC™ PowerEdge™ R7525 powered by AMD EPYC™ 7502 processors, and an HPE ProLiant DL380 Gen10 powered by Intel® Xeon® Gold 6240 processors. The Dell EMC server cost \$38,482.50, while the HPE server cost \$39,846.00 (a difference of just three percent).

To test each server, we used the k-means clustering workload from the Spark-Bench benchmark. The PowerEdge R7525 server required less time to process the 888GB dataset, finishing in 2 hours, 8 minutes compared to 3 hours, 34 minutes for the ProLiant DL380 Gen10 server. On average, the Dell EMC server processed 67 percent more data each second than the HPE server—121,241 kilobytes per second (KB/s) versus just 72,518 KB/s.

Because the Dell EMC server completed its work faster than the HPE solution and at a similar price, we determined that it presented a 74 percent better value in terms of hardware cost vs. data processing performance.


\*HPE ProLiant DL380 Gen10 server

## About Dell EMC PowerEdge R7525 servers

The Dell EMC PowerEdge R7525 we tested was powered by two AMD EPYC 7502 processors that contained 32 cores each. According to Dell EMC, this server also boasts the following:

- Up to 24 NVMe direct connections
- 32 DDR4 RDIMM or LRDIMM memory module slots
- Automated server life cycle management
- Security features such as AMD Secure Memory Encryption (SME) and Secure Encrypted Virtualization (SEV)

To learn more, visit <https://www.dell.com/en-us/work/shop/povw/poweredge-r7525>.



## K-means clustering: Crunch data in less time

K-means clustering is a popular type of analytics algorithm that organizations can use to mine data, find similarities between data points, and uncover patterns therein. For example, you might use k-means clustering on a population map to determine the most advantageous locations for you to offer services. The sooner your servers can complete a k-means workload, the sooner your company can use the data it collects for business-critical endeavors.

Note that most big data analytics workloads will use many servers to process many terabytes of data. We tested with just one server for each solution as a proof of concept.

Figure 1 shows the results of our Spark-Bench tests. The Dell EMC PowerEdge R7525 server required 40 percent less time to process 888 GB of data in our k-means workload, finishing in 2 hours, 8 minutes compared to the HPE ProLiant DL380 Gen10 server's 3 hours, 34 minutes—a difference of 1 hour, 26 minutes.




Figure 1: Time to complete the Spark-Bench k-means workload on an 888GB file. Lower is better.  
Source: Principled Technologies.



## K-means processing rate

The Dell EMC PowerEdge R7525 server completed the k-means clustering workload more quickly than the HPE ProLiant DL380 Gen10 server. To calculate average rate of processing, we divided the size of the 888GB data set by the time it took each solution to complete the k-means test. Figure 2 shows the rate of processing for each solution in terms of kilobytes per second (KB/s). On average, the Dell EMC server processed 121,241 KB/s, which is 67 percent more than the HPE server's 72,518 KB/s rate. Note this is not disk or networking throughput, but rather a rate of processing the dataset in our k-means clustering workload.




Figure 2: Data processing rate during the Spark-Bench k-means workload (average). Higher is better. Source: Principled Technologies.

Process up to 67% more data per second

## Cost and value: Get more from your hardware investment

We configured the servers in this study to have a similar price point. The two systems had comparable drives, RAM size, and network cards, but differed in power supplies, RAM speed, and CPU cores. We were able to give the Dell EMC server higher wattage power supplies (2,400W vs 800W), slightly faster RAM (3,200MHz vs 2,933MHz), and processors with a higher core count (32 vs. 18) for around the same total hardware price as the HPE server.

Figure 3 shows the hardware cost of each server.<sup>1,2</sup> Hardware costs fluctuate often, but at the time of this writing, the prices for the two servers we tested were similar, differing by just three percent.

Even though the prices are similar, the magnitude of the Dell EMC PowerEdge R7525 server's performance win means it offered a better value in terms of hardware cost and data processing performance. Comparing these values, as shown in Figure 4, reveals that the Dell EMC solution had a performance/cost ratio of 3.15 KB/s per dollar, whereas the HPE server's ratio was only 1.81 KB/s per dollar. In this configuration and workload, the Dell EMC PowerEdge R7525 presented a 74 percent better performance/cost ratio for our specific test case.




Figure 3: Hardware cost. Lower is better.  
Source: Principled Technologies.




Figure 4: Performance/cost ratio. Higher is better.  
Source: Principled Technologies.

**74% better  
performance/  
cost ratio**



## Conclusion

When it comes to data analytics, time is money. Powerful servers earn their keep by analyzing data quickly—the more time your servers spend crunching numbers, the less time your business has to make critical decisions to stay ahead of your competition and extend the reach of your sales.

In our tests comparing a Dell EMC PowerEdge R7525 server to an HPE ProLiant DL380 Gen10 server, we found that the Dell EMC solution:

- Completed a k-means clustering workload in **40 percent less time**
- Processed **67 percent more data per second**
- Carried a **74 percent better performance/cost ratio** in terms of data processing performance vs. hardware price

The performance boost will come in handy for businesses who need to analyze and sort large datasets. Though businesses will typically use many server nodes for their data analytics work, these results should encourage IT decision-makers to explore the benefits of the Dell EMC PowerEdge R7525.

- 1 On February 25, 2020, Dell EMC sent us the list price for our configuration of the Dell EMC PowerEdge R7525.
- 2 On March 2, 2020, a certified HPE reseller sent us the list price for our configuration of the HPE ProLiant DL380 Gen10.

Read the science behind this report at <http://facts.pt/s8ite7f> ▶



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.