



## Complete more Apache Cassandra database work with Microsoft Azure Lsv3-series VMs enabled by 3<sup>rd</sup> Gen Intel® Xeon® Scalable processors

These VMs achieved as many or more operations per second (OPS) than Lasv3-series VMs enabled by 3<sup>rd</sup> Gen AMD EPYC processors

Organizations that support ecommerce, media streaming, social media, and other NoSQL applications need to provide scalable, high-access performance for massive amounts of data. That's where distributed databases such as Apache Cassandra® come in. Cassandra can provide reliable architecture for NoSQL applications, but if you're running database workloads in the cloud, you need to know that your VMs can handle user demands.

We compared Cassandra distributed database performance on two series of Microsoft Azure VMs: a three-node cluster of Lsv3-series VMs with 3<sup>rd</sup> Gen Intel® Xeon® Scalable processors and a three-node cluster of Lasv3-series VMs with 3<sup>rd</sup> Gen AMD EPYC™ processors. We tested each VM cluster at three sizes and found that at each size, the Lsv3-series VM clusters completed as many or more OPS than the Lasv3-series VM clusters.

Up to 11% more operations per second

compared to Lasv3-series VMs

## How we tested

We compared the performance of two series of storage-optimized Azure VM clusters at three sizes:

- A three-node cluster of Lsv3-series VMs featuring 3<sup>rd</sup> Gen Intel Xeon Platinum 8370C processors
- A three-node cluster of Lasv3-series VMs featuring 3<sup>rd</sup> Gen AMD EPYC 7763 processors

We tested each VM in the South Central US region of Azure. Figure 1 shows the size specifications of each VM cluster we tested.




Figure 1: Key specifications of the Microsoft Azure Lsv3- and Lasv3-series VM clusters we tested. Source: Principled Technologies.

## About Apache Cassandra

According to Apache, “Cassandra is 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>1</sup>

We used the `cassandra-stress` tool, a benchmark that comes built into Cassandra, to test the VM clusters. According to documentation, the tool “supports testing arbitrary CQL tables and queries, allowing users to benchmark their own data model.”<sup>2</sup>

To learn more about Cassandra, visit [https://cassandra.apache.org/\\_/index.html](https://cassandra.apache.org/_/index.html).

## Complete more operations per second with Azure Lsv3-series VMs

We tested each VM cluster with the cassandra-stress benchmark tool, which reports the number of operations per second each cluster completed. For each cluster, we ran the test three times, and we report the median result. At all three sizes, we saw that the Lsv3-series VM clusters with 3<sup>rd</sup> Gen Intel Xeon Scalable processors completed more OPS than the Lasv3-series VM clusters with 3<sup>rd</sup> Gen AMD EPYC processors. At the small size, the Lsv3-series VM clusters completed 11.73 percent more OPS; at the medium size, 8.50 percent more; and at the large size, 0.62 percent more.




Figure 2: Comparison of the number of operations per second that the cluster of small Lsv3 VMs completed, relative to the number of operations per second that the cluster of small Lasv3 VMs completed. Source: Principled Technologies.




Figure 3: Comparison of the number of operations per second that the cluster of medium Lsv3 VMs completed, relative to the number of operations per second that the cluster of medium Lasv3 VMs completed. Source: Principled Technologies.




Figure 4: Comparison of the number of operations per second that the cluster of large Lsv3 VMs completed, relative to the number of operations per second that the cluster of large Lasv3 VMs completed. Source: Principled Technologies.



### About Microsoft Azure Lsv3-series VMs

New storage-optimized Azure Lsv3-series VMs feature 3<sup>rd</sup> Gen Intel Xeon Scalable processors in hyperthreaded configurations. According to Microsoft, the VMs also offer the following specifications:<sup>3</sup>

- Up to 80 vCPUs and up to 800 GiB of RAM
- All-core turbo clock speed of up to 3.5GHz
- Intel Turbo Boost Technology
- Intel Advanced-Vector Extensions 512 (Intel AVX-512)
- Intel Deep Learning Boost

To learn more, visit <https://docs.microsoft.com/en-us/azure/virtual-machines/lsv3-series>.



## Conclusion

To run your NoSQL applications, you need flexibility and reliability from your distributed database. If you're using the cloud to support your database, you also need VMs that can support your users' needs. When we tested Apache Cassandra performance on three sizes of Lsv3-series VM clusters and three sizes of Lsav3-series VM clusters, we found that the Lsv3-series VM clusters offered more OPS.

1. "Apache Cassandra," accessed July 7, 2022, [https://cassandra.apache.org/\\_/index.html](https://cassandra.apache.org/_/index.html).
2. "Cassandra Stress," accessed July 8, 2022, [https://cassandra.apache.org/doc/latest/cassandra/tools/cassandra\\_stress.html](https://cassandra.apache.org/doc/latest/cassandra/tools/cassandra_stress.html).
3. "Lsv3-series," accessed July 7, 2022, <https://docs.microsoft.com/en-us/azure/virtual-machines/lsv3-series>.

Read the science behind this report at <https://facts.pt/3vjLnKq> ▶



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 Intel.