



With small N2 standard VM instances:

Handle up to 1.22x the MariaDB database transactions per minute



With medium N2 standard VM instances:

Handle up to 1.22x the MariaDB database transactions per minute



With large N2 standard VM instances:

Get even greater performance of up to 1.25x the MariaDB database transactions per minute

Boost your MariaDB online transaction processing performance with N2 standard VM instances for Google Cloud Platform featuring 3rd Generation Intel Xeon Scalable processors

Compared to N2 standard VM instances with previous-generation Intel Xeon Scalable processors, these processed a higher rate of transactions per minute

Google Cloud Platform has recently introduced VM instances that use the latest generation of Intel Xeon Scalable processor. Because it may not be obvious which processors are at the heart of your VM instances, you may not realize you have the option to use the latest processors to improve the performance of your mission-critical workloads.

At Principled Technologies, we used an online transaction processing workload to assess the performance of two sets of Google Cloud Platform VM instances that ran MariaDB databases: N2 standard VM instances featuring 3rd Generation Intel Xeon Scalable processors and N2 standard VM instances featuring 2nd Generation Intel Xeon Scalable processors. In our tests, the VM instances with 3rd Gen Intel Xeon Scalable processors consistently processed a higher rate of transactions per minute—up to 1.25 times the rate that the VM instances with previous-generation processors accomplished.

By ensuring that you use N2 standard VM instances with this latest generation of Intel Xeon Scalable processor, your company could serve more online customers and potentially pave the way for future growth.

How we tested

We compared the online transaction processing performance of two sets of N2 standard VM instances that differed by processor. One set of VM instances contained 3rd Gen Intel Xeon Scalable processors (Ice Lake, or ICX), while the other set used 2nd Gen Intel Xeon Scalable processors (Cascade Lake, or CLX). We tested three sizes of VM instance for each set. Figure 1 illustrates key configuration information for the VM instances we tested.

Google Cloud Platform does not provide information on the specific model of processor powering the instances you spin up. However, N2 standard instances are available only with 3rd and 2nd Gen Intel Xeon Scalable processors. By matching processor speed data from the Google Cloud Platform interface to publicly available spec sheets for each generation of Intel Xeon Scalable processor, we were able to differentiate with reasonable certainty the generation of processor supporting each instance. We ensured that the processor speed for each VM instance remained consistent from test to test. Should you decide to take advantage of 3rd Gen Intel Xeon Scalable processors for your N2 instances, you would need to do the same.

To make certain that each VM instance had resources appropriate for the workload requirements and that each VM instance's storage would not negatively affect performance, we configured each test database to fit within the allocated memory footprint for each VM instance. While this may not always be practical in the real world, we generally find that sizing databases in this fashion results in strong and consistent performance while lowering cloud storage costs.

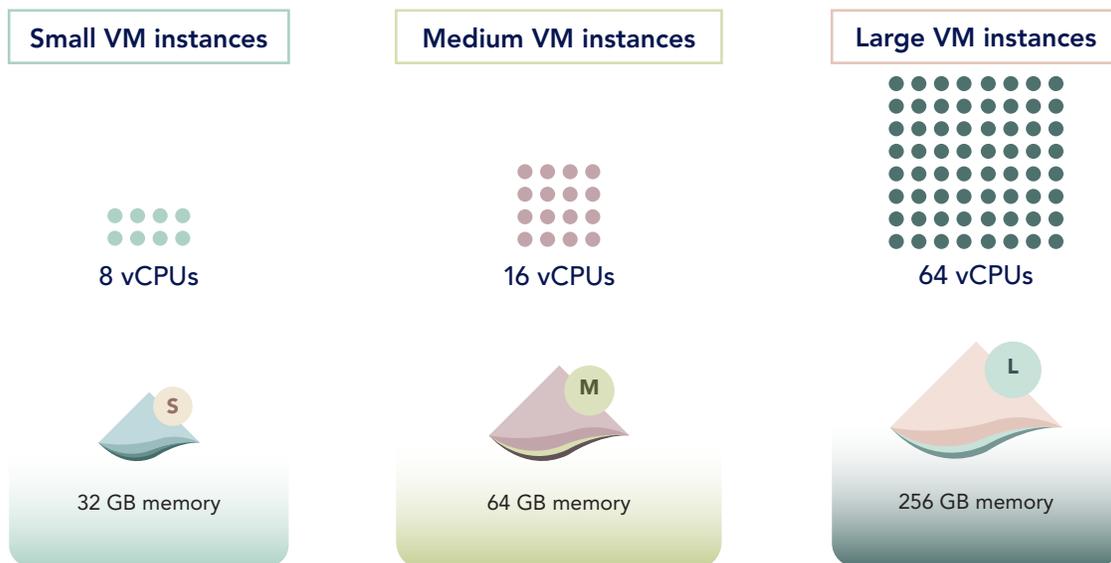


Figure 1: Virtual CPU and memory information for the VM instances we tested. All VM instances ran in the us-central-1 region. Source: Principled Technologies.

The HammerDB workload we ran on MariaDB

Ecommerce sites, food delivery services, financial institutions and more commonly use online transaction processing workloads during the normal course of business. To assess the OLTP performance of each VM instance, we used the TPROC-C workload from the HammerDB benchmark suite. While the HammerDB developers based TPROC-C on the TPC-C standard, it is not a full implementation of that standard. Therefore, the results we report in this study cannot be directly compared to official TPC-C results. For more information, see <https://www.hammerdb.com/docs/ch03s05.html>.

Our results

Across the small, medium, and large VM instance sizes we tested, the N2 standard VM instances with 3rd Generation Intel Xeon Scalable processors handled a higher rate of MariaDB transactions per minute compared to N2 standard VM instances with 2nd Gen Intel Xeon Scalable processors. This consistent advantage across multiple VM instance sizes could translate to benefits for a variety of business types. Figures 2 through 4 showcase our test results.

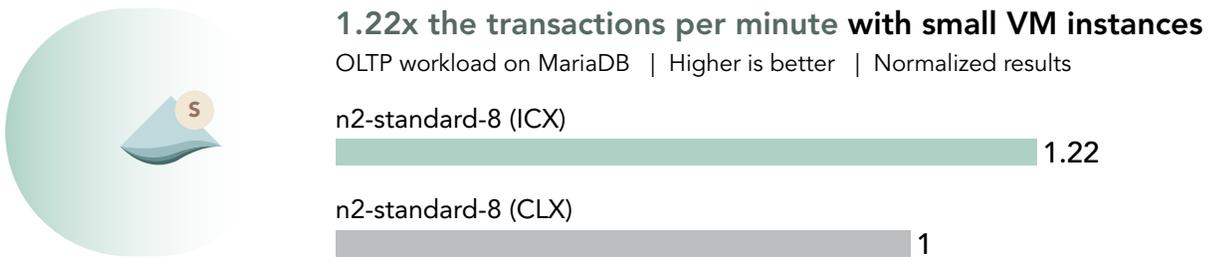


Figure 2: Comparison of the average rate of transactions per minute each n2-standard-8 VM instance achieved, relative to the rate of the N2 standard VM instance with 2nd Gen Intel Xeon Scalable processors. Source: Principled Technologies.

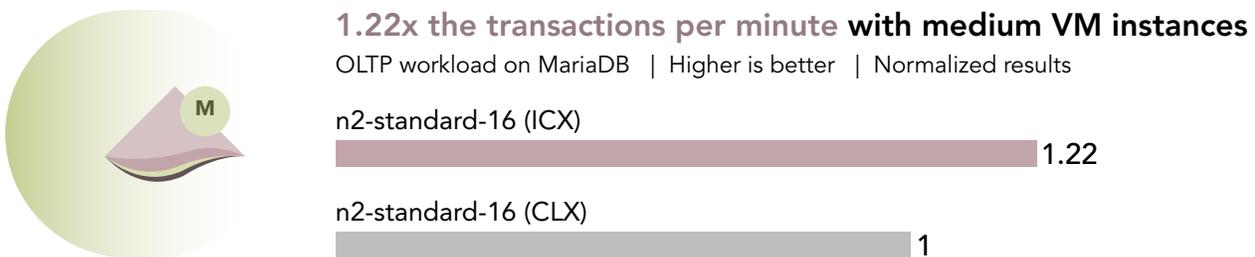


Figure 3: Comparison of the average rate of transactions per minute each n2-standard-16 VM instance achieved, relative to the rate of the N2 standard VM instance with 2nd Gen Intel Xeon Scalable processors. Source: Principled Technologies.

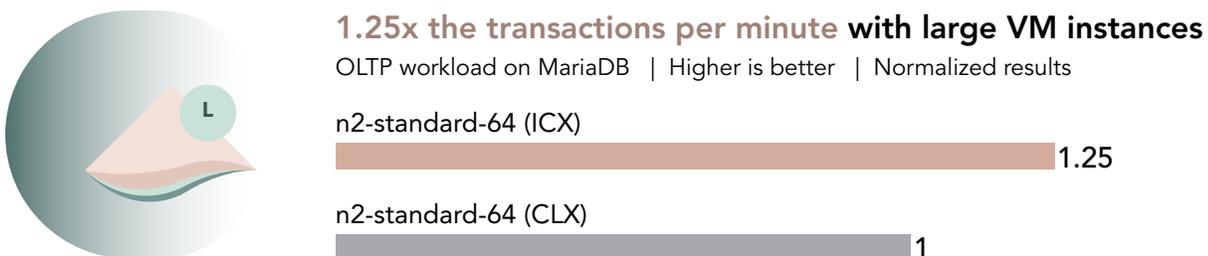


Figure 4: Comparison of the average rate of transactions per minute each n2-standard-64 VM instance achieved, relative to the rate of the N2 standard VM instance with 2nd Gen Intel Xeon Scalable processors. Source: Principled Technologies.



Conclusion

If your company uses N2 standard VM instances from Google Cloud Platform, stronger performance for your online transaction processing workloads is already within reach.

In our OLTP tests with MariaDB databases, N2 standard VM instances featuring 3rd Generation Intel Xeon Scalable processors handled a higher rate of transactions per minute compared to N2 standard VM instances featuring 2nd Generation Intel Xeon Scalable processors. This remained true across three different sizes of VM instances for an up to 1.25x advantage.

By ensuring your VM instances use 3rd Generation Intel Xeon Scalable processors, you could accelerate your OLTP work at no additional cost.

Read the science behind this report at <https://facts.pt/JuCr08n> ►



Facts matter.®

This project was commissioned by Intel.

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.