Achieve faster online analytics processing work with newer VM instances for Google Cloud Platform powered by 2nd Generation Intel Xeon Scalable Processors – Cascade Lake

The newer VM instances, powered by 2nd Generation Intel Xeon Scalable Cascade Lake processors, completed a data analytics workload faster than older VM instances powered by processors from a previous generation

Great business decisions don't just spring out of the ether. The choices that control day-to-day operations often require input from reports and other insights generated by robust data analytics applications. If your business is going to use the cloud to power these applications, you'll want VM instances that can complete analysis jobs quickly, as getting a head start on the decision-making process can be key to your organization's success.

At Principled Technologies, we used an online analytics processing (OLAP) workload to test general-purpose VM instances from two Google Cloud Platform series: newer N2 standard series VM instances powered by 2nd Generation Intel® Xeon® Scalable Cascade Lake processors, and older N1 standard series VM instances powered by processors from a previous generation. In our tests, the N2 standard series VM instances finished the OLAP workload faster than the N1 standard series VM instances. This held true for both medium and large-sized VM instances, resulting in up to a 1.68x advantage. Faster data analytics could help your company reach key decisions sooner, and even help to save on uptime costs by shrinking your daily analytics window.

Medium VM instances: Complete data analysis up to 1.27 times as fast*

Large VM instances: Complete data analysis up to 1.68 times as fast**

*Medium VM instances contained 16 vCPUs and ran a 30GB database

**Large VM instances contained 64 vCPUs and ran a 100GB database

How we tested

VM instances in Google Cloud Platform

We tested two types of VM instances for Google Cloud Platform: N2 standard series VM instances powered by 2nd Generation Intel Xeon Scalable Cascade Lake processors, and N1 standard series VM instances powered by older processors. Note that Google Cloud Platform does not make available the specific model number for the processor powering each VM instance. We made sure that the processor speed information they provided was consistent from test to test, but we cannot say for certain which specific processors we used for each test. We provide additional configuration information in the Science behind this report.

VM instance sizes

We tested two sizes of VM instance for each series: Medium (containing 16 vCPUs and a 30GB database) and Large (containing 64 vCPUs and a 100GB database). We chose these database sizes for two reasons: First, we wanted to size the databases to match the same medium and large real-world workloads represented by the two different VM instance sizes. We also wanted to ensure the OLAP workload was putting more stress on each VM instance's processor rather than storage; limiting the database size to fit within the limits of each VM instance's allocated RAM helped us achieve this goal. Figure 1 is a visual representation of the VM instances we used for testing. We tested all VM instances in the us-east1-b region.

Figure 1: Breakdown of the VM instances we used for testing. Source: Principled Technologies.

HammerDB workload

We tested each VM instance with an OLAP workload from the HammerDB suite of benchmarks. The HammerDB developers based this workload on the TPC-H specification. However, because the workload is not a full implementation of the TCP-H benchmark, the results in this paper are not directly comparable to official, published TPC-H results. The HammerDB OLAP test measures the time required for VM instances to analyze a "stream" of database queries. (HammerDB defines one query stream as 22 serialized database queries.)

We began our tests by having each VM instance complete the workload with just a single stream. Though VM instances can analyze multiple query streams simultaneously, the Transaction Processing Council (TPC) lays out recommendations for the maximum number of query streams to use for databases of varying sizes.[1] For each test, we completed a single query stream test to cache the database, then immediately ran a concurrent stream test. We repeated these steps at increasing concurrent stream counts until we reached the maximum number TPC recommends for each VM size (four streams for the Medium VM instances and five streams for the Large VM instances). To ensure we were fully representing the performance of each processor, we confirmed that each VM instance's processors achieved full saturation during the tests that used the maximum number of streams from TPC-H recommendations.

How N2 series VM instances help your business

The N2 series offers a few advantages that contribute to the improved performance compared to the N1 series, including:[2]

- Higher processor frequency (2.8GHz vs. 2.2GHz)
- Increased memory at each size
- Increased egress network bandwidth limits

Our results

Figure 2 shows the full results of our tests with Medium VM instances, while Figure 3 shows the Large VM instance test results. Both figures show that the N2 standard series VM instances had markedly better data analysis times than the older N1 standard series VM instances. For Medium VM instances, the n2-standard-16 VM instances completed analysis from 1.22 to 1.27 times as fast as the n1-standard-16 VM instances. For Large VM instances, the n2-standard-64 VM instances completed analysis from 1.48 to 1.68 times as fast as the n1-standard-64 VM instances.

It can be difficult to translate results like these to your own business, so we've devised a hypothetical scenario that can help to illuminate the benefits that N2 standard series VM instances powered by 2nd Generation Intel Xeon Scalable processors can bring.

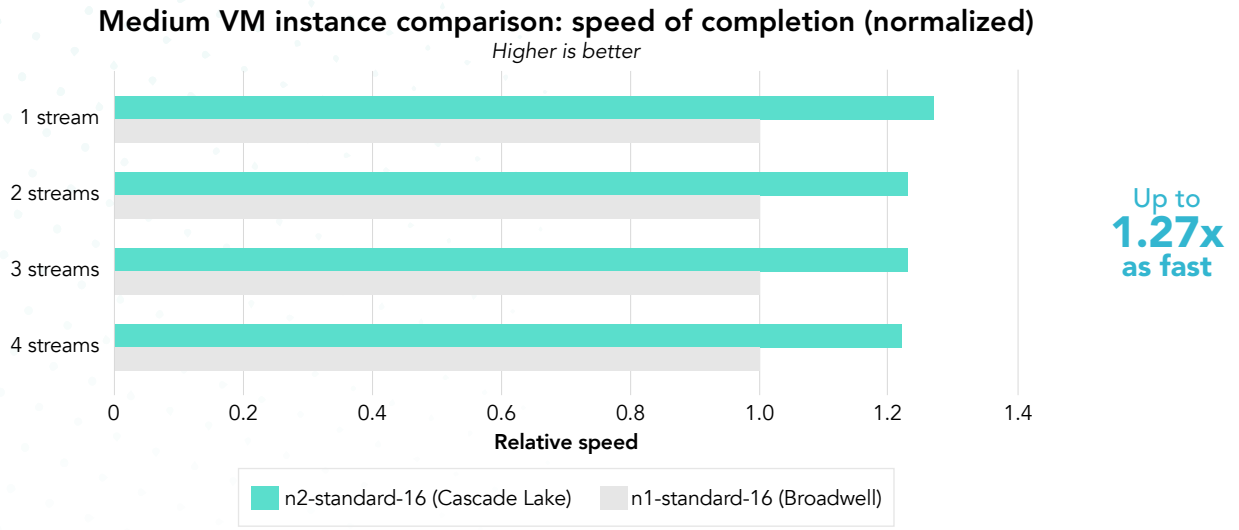

Figure 2: Comparison of the relative speed at which each of the medium VM instances completed the TPC-H-like workload from HammerDB, relative to the completion time of the n1-standard-16 VM instance. Greater speed is better. Source: Principled Technologies.

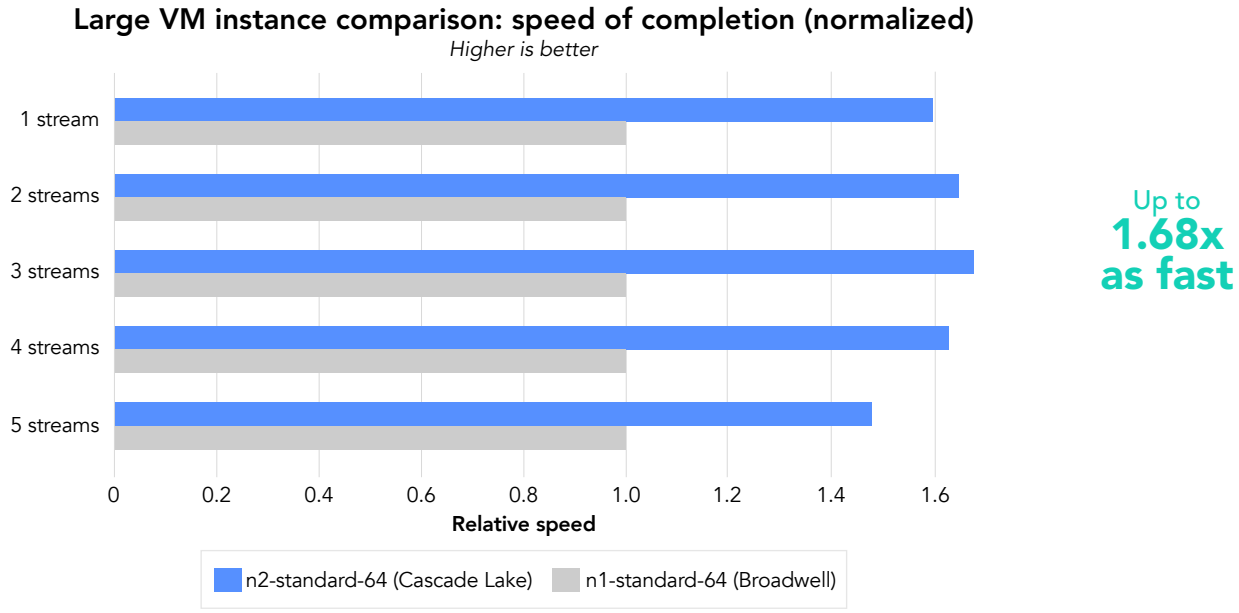

Figure 3: Comparison of the speed at which each of the large VM instances completed the TPC-H-like workload from HammerDB, relative to the completion time of the n1-standard-64 VM instance. Greater speed is better. Source: Principled Technologies.

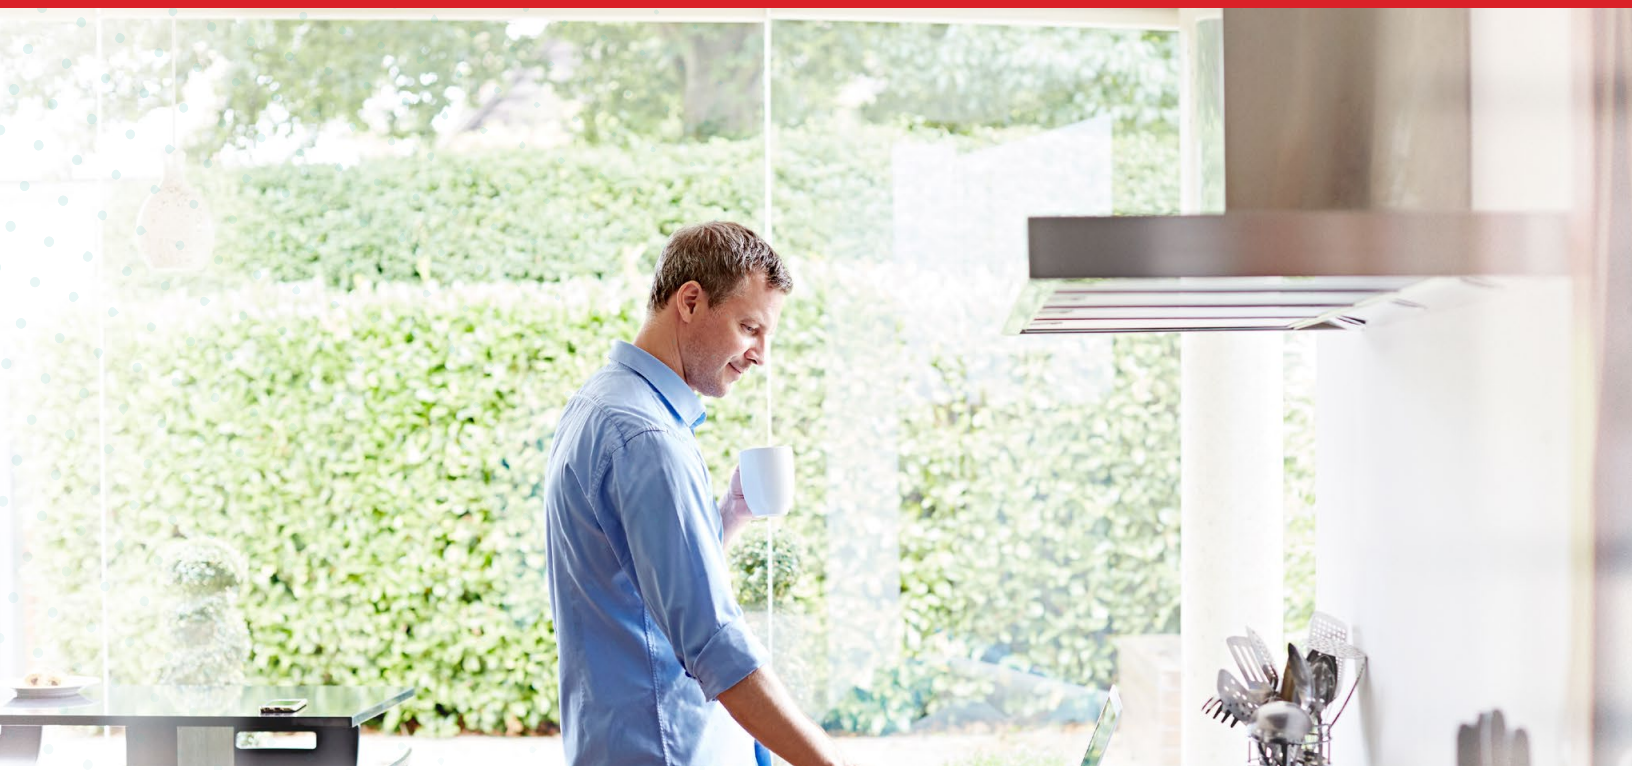

Real-world context

Imagine that each night, your company has a four-hour window in which to analyze data on a 100GB database. This analysis is quite important, as each day your company generates reports that help executives and other leaders to focus their efforts and provoke discussion about business improvements.

Based on the results of our single-stream OLAP tests, we calculate that the Large N2 standard VM instance would enable you to complete 580 query streams within your nightly analysis window. However, the Large N1 standard VM instance would allow time for only 300 query streams—a difference of 37 percent.

Now, let's say that each night, your company needs only complete 300 query streams. The large N1 standard VM instance would require four hours. But with the N2 standard VM instance, you would be able to complete this task in just 2.5 hours, shrinking your required window by an hour and a half compared to the large N1 standard VM instance. Over the course of a year, you would save 547.5 hours, or 22.8 full days' worth of analysis time per year. You would be achieving the same amount of analysis you could get with the large N1 standard VM instance, but with time to spare, potentially enabling you to save money on VM instance uptime.

Performance that scales

Medium N2 VM instances were 1.22 to 1.27 times as fast as medium N1 VM instances. The performance margins for the large 64 vCPU VM instances were even higher, with the N2 VM instances being 1.48 to 1.68 times as fast as the N1 series.

1.22 – 1.27 times as fast
for medium instances

1.48 – 1.68 times as fast
for large instances

Which is the better investment?

Sometimes, the cost associated with newer technologies can make you wonder whether opting for older tech might be a smarter financial move. Our test results suggest that wouldn't be the case with N2 standard series VM instances for OLAP work. The newer N2 standard series VM instances cost only 1.02 times as much as their older counterparts, yet they achieved 1.22 to 1.68 times the performance in our testing.[3]

Conclusion

Every day, your company makes decisions that affect the course of business. When your data analytics tools can generate insights quickly, you can devote more time to ensuring your next decision is a good one.

In our online analytics processing tests on Google Cloud Platform, newer N2 standard series VM instances powered by 2nd Generation Intel Xeon Scalable Cascade Lake processors analyzed query streams up to 1.68 times as fast as older N1 standard series VM instances powered by older processors. For this increased performance, the new instances cost just 1.02 times as much, making them a more cost-effective choice for running your data warehouse workloads.

1 Transaction Processing Performance Council (TPC), "TPC Benchmark H," accessed November 4, 2020, http://www.tpc.org/tpc_documents_current_versions/pdf/tpc-h_v2.18.0.pdf

2 Machine Types," accessed November 12, 2020, https://cloud.google.com/compute/docs/machine-types#n1_machine_types

3 "VM Instances Pricing," accessed November 4, 2020, https://cloud.google.com/compute/vm-instance-pricing#n2_predefined

Read the science behind this report at http://facts.pt/NZJ1Hqd ▶

Principled Technologies®

Facts matter.®

This project was commissioned by Intel.