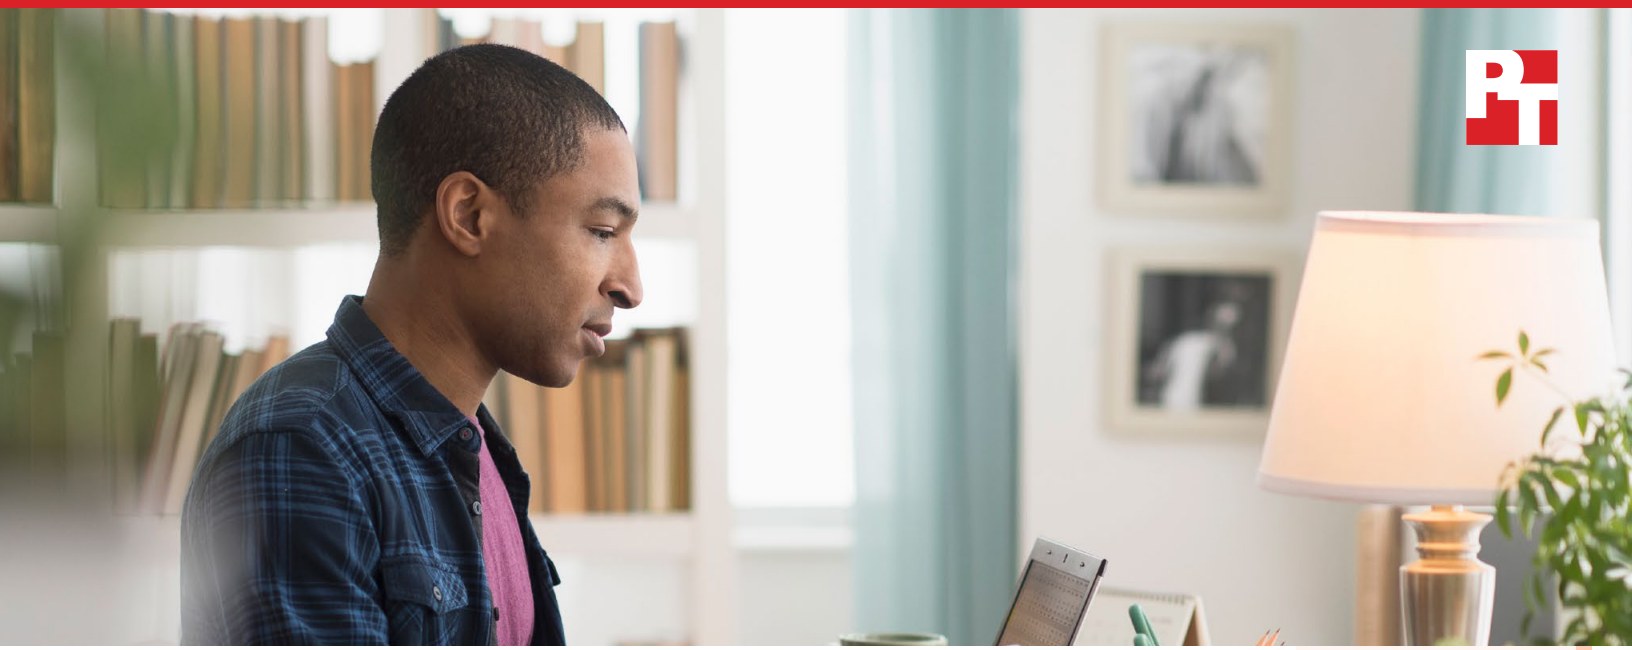

Complete Microsoft SQL Server queries faster with new Microsoft Azure VMs powered by 2nd Generation Intel Xeon Scalable processors – Cascade Lake

New VMs powered by 2nd Generation Intel Xeon Scalable (8272CL) Cascade Lake processors delivered consistently better data warehouse performance than older VMs with Intel Xeon (E5-2673 v4) Broadwell processors

Your organization likely has a great deal of data and is generating more every day. To make good use of that data, you must first be able to make sense of it: extract insights, discover trends, and look at it from multiple angles. You can get to that value faster with a speedy data warehouse platform. If you're analyzing your data in the cloud, however, there are a large number of VM types to choose from; you must decide which will best meet your performance needs.

To help you make this decision, we measured the online analytical processing (OLAP) performance of two Microsoft Azure VM series: older Ds_v3 VMs, powered by Intel Xeon (E5-2673 v4) Broadwell processors, and newer Dds_v4 VMs, powered by 2nd Generation Intel Xeon Scalable (8272CL) processors. At both sizes we tested (16 vCPUs and 64 vCPUs), the newer Dds_v4 VMs completed the analytics workload faster. If you're seeking to gain value from your data in Microsoft Azure, these newer VMs can help you do it sooner.

Medium VMs:
Finish data warehouse analysis up to
1.49x
as fast*

Large VMs:
Finish data warehouse analysis up to
1.35x
as fast**

*D16ds_v4 instances with 16 vCPUs vs. D16s_v3 instances with 16 vCPUs
**D64ds_v4 instances with 64 vCPUs vs. D64s_v3 instances with 64 vCPUs

How we tested

Using a HammerDB data warehouse workload, we tested Microsoft SQL Server performance on older Ds_v3 Microsoft Azure VMs, powered by Intel Xeon (E5-2673 v4) Broadwell processors, and newer Dds_v4 Microsoft Azure VMs, powered by 2nd Generation Intel Xeon Scalable (8272CL) Cascade Lake processors. (Microsoft Azure Ds_v3 VMs are available in a few different CPU configurations, but we used only Intel Xeon E5-2673 v4 processors for our testing.)

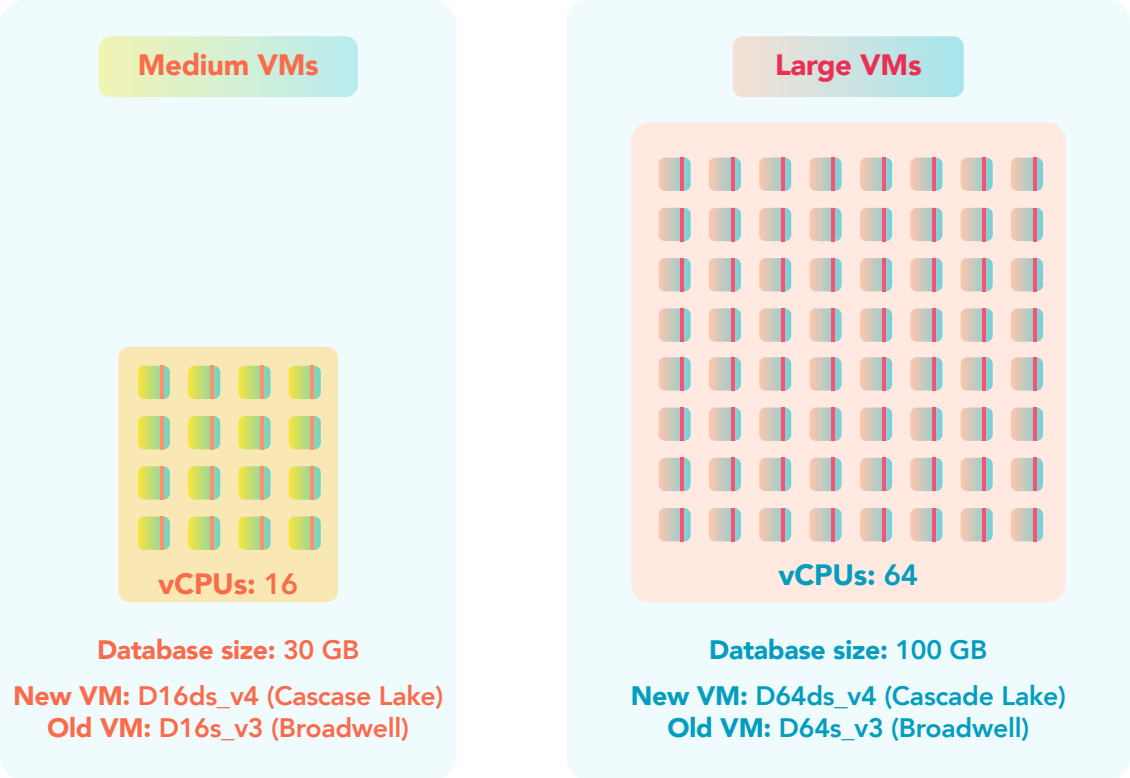

Figure 1: Configuration details for the VMs we tested. Source: Principled Technologies.

We wanted to assess how data analytics performance would scale from a smaller VM to a larger one. Running the VMs in the East US region, we tested two sizes of each VM type, 16-vCPU with a 30GB database and 64-vCPU with a 100GB database. Two factors determined how we sized each database: matching the database size to a likely workload scenario for each compute level, and keeping each database size within the limits of the RAM allocated to each VM to ensure we were fully representing the performance of each processor.

Note that the HammerDB developers derived the data warehouse workload from TPC-H benchmark specifications. Because the HammerDB test is not a full implementation of the TPC-H specification, our results are not directly comparable to published TPC-H results. Note also that prior to each test, we ran a single-user query stream to cache the database in memory.

About HammerDB

We tested both sets of VMs with an OLAP workload from the HammerDB suite of benchmarks. Their TPC-H-like workload measures how long it takes to complete sets, or streams, of 22 serialized database queries.

To learn more about HammerDB and their TPC-H-like benchmark, visit www.hammerdb.com.

Our results

Medium VMs (30GB database)

We measured the time it took for each 16-vCPU VM to complete a single stream of data warehouse queries. Next, we scaled up to two, three, and four concurrent streams. (The TPC-H specification recommends a maximum of four streams for this size of database.)

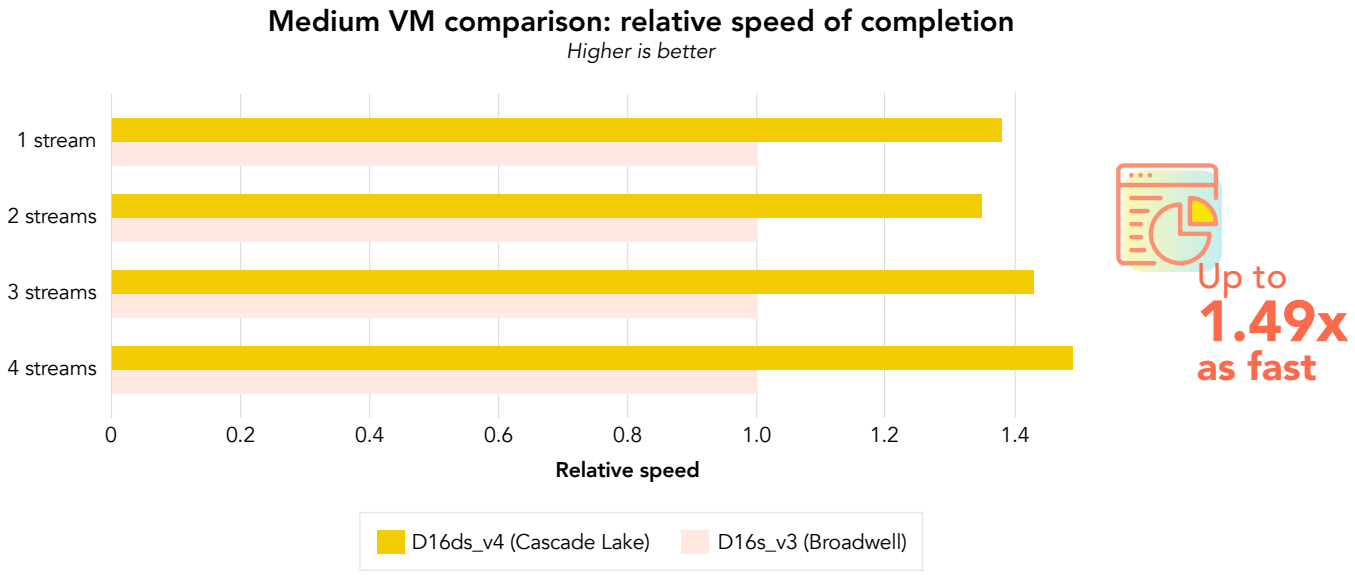

Figure 2: Comparison of the relative speed at which the medium VMs completed one, two, three, and four query streams from the TPC-H-like HammerDB workload. Greater speed is better. Source: Principled Technologies.

The new Dds_v4 Microsoft Azure VM, powered by Intel Xeon Platinum 8272CL processors, consistently completed the queries in less time than the older Intel Xeon E5-2673 v4 processor-powered Ds_v3 VM. When we tested four simultaneous streams, the new VM completed the work 1.49 times as fast as the older VM.

Series' advantages callout

According to Microsoft, Dds_v4 series VMs offer several advantages that may contribute to better performance than the Ds_v3 series:[3,4]

- Guaranteed 2nd Generation Intel Xeon Scalable processors for every VM
- 50 percent larger default disk drives
- Higher IOPS on default disk drives

More performance, only slightly higher price

Increased performance does come with a small cost increase—but only a small one. While completing data warehouse queries up to 1.49 times as fast as the older Ds_v3 VM, the newer Dds_v4 VMs were just 1.17 times the cost. That increase is a comparatively small price to pay for significantly higher performance.[1]

Large VMs (100GB database)

Though Transaction Processing Performance Council (TPC) mandates at least five streams for a VM with a 100GB database[2], we assessed the performance of the 64-vCPU VMs while running one, two, three, four, and five concurrent streams to provide additional scenarios for readers. Again, the new Dds_v4 VM powered by 2nd Generation Intel Xeon Scalable (8272CL) Cascade Lake processors offered consistently faster performance than the older Ds_v3 VM with Intel Xeon E5-2673 v4 processors. When processing three, four, and five streams simultaneously, the Dds_v4 VM was 1.35 times as fast as the Ds_v3 VM.

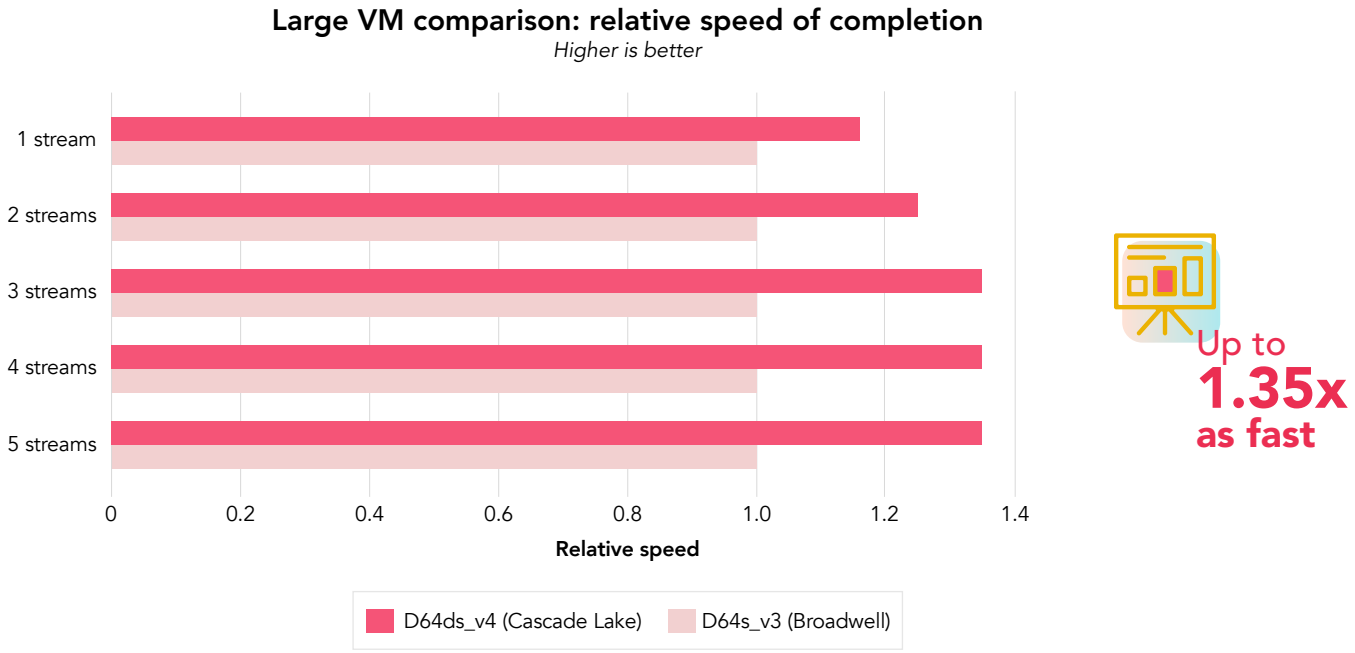

Figure 3: Comparison of the relative speed at which each large VM completed one, two, three, four, and five query streams from the TPC-H-like HammerDB workload. Greater speed is better. Source: Principled Technologies.

Hands-on testing, real-world benefit

For a better idea of how our results can relate to the real world, consider the following hypothetical scenario. Each night, a company has a one-hour window in which to analyze data on a 30GB database. The company uses this analysis to generate reports for executives each morning and to drive other business processes throughout the day.

Based on the results of our single-stream tests, we calculate that a new D16ds_v4 VM would enable this hypothetical company to run 171 query sets within their analysis window each night. By contrast, the D16s_v3 VM would complete just 124 query sets in the same time frame (27.5 percent fewer).

If this company needed to complete only 124 query sets each night, the new D64ds_v4 VM would enable them to finish their work in just 43 minutes, shrinking the required analysis window by 17 minutes compared to the D16s_v3 VM.

Over the course of a year, this hypothetical company would save 103.4 hours—or, 4.3 full days' worth—of analysis time while maintaining the same rate of analysis they could get with the D16s_v3 VM, enabling them to save money on VM uptime.

Conclusion

With speedy cloud VMs powering your online analytics processing work, your business could be well on its way to discovering the key insights that will transform it for the better.

In our tests, new Dds_v4 series VMs powered by 2nd Generation Intel Xeon Scalable (8272CL) Cascade Lake processors analyzed data warehouse query streams up to 1.49 times as fast as older Ds_v3 series VMs powered by Intel Xeon (E5-2673 v4) Broadwell processors. In addition to enabling your company to cut down on analysis times and start decision-making processes sooner, the Dds_v4 VMs could be a better investment than Ds_v3 VMs, giving you more performance for only a small cost increase.

---

1 "Windows Virtual Machine Pricing," accessed November 3, 2020, https://azure.microsoft.com/en-us/pricing/details/virtual-machines/windows/.

2 Transaction Processing Performance Council (TPC), "TPC Benchmark H," accessed November 3, 2020, http://www.tpc.org/tpc_documents_current_versions/pdf/tpc-h_v2.18.0.pdf.

3 "Ddv4 and Ddsv4 series," accessed November 12, 2020, https://docs.microsoft.com/en-us/azure/virtual-machines/ddv4-ddsv4-series

4 "Dv3 and Dsv3-series," accessed November 12, 2020, https://docs.microsoft.com/en-us/azure/virtual-machines/dv3-dsv3-series

Read the science behind this report at http://facts.pt/sjWlJxD ►

Principled Technologies®

Facts matter.®

This project was commissioned by Intel.