

Enhance Databricks Data Warehouse Workloads with Azure E8ds_v4 VMs and Databricks Photon Engine



Authors

Swastik Chakraborty

Cloud Solution Architect
swastik.chakraborty@intel.com

Lakshman Chari

Cloud ISV Partner Manager
lakshman.chari@intel.com



**Up to 80% Less
Time to Run Decision
Support Queries with
E8ds_v4 VMs vs. older
E8s_v3 VMs**

(see Figure 3, 10TB dataset)

**Up to 31% Less
Time to Run Decision
Support Queries with
E8ds_v4 VMs vs. E8as_v4
VMs with AMD EPYC
Processors**

(see Figure 5, 1TB dataset)

**Up to 65% Less
Time to Run Decision
Support Queries with
Photon Enabled vs.
Without Photon on
E8ds_v4 VMs**

(see Figure 1, 1TB dataset)

Featuring 2nd Generation Intel® Xeon® Scalable processors, E8ds_v4 VMs outperformed both E8s_v3 and E8as_v4 VMs

As IOT devices, websites, databases, and other sources provide increasing quantities of data to companies, the ability to store and analyze that data becomes increasingly important. Data lakes and data warehouses provide large-scale storage infrastructure for unstructured data and structured data, respectively. Databricks combines features from both data warehouses and data lakes to store and analyze vast amounts of structured and unstructured data with their Lakehouse Platform. The Lakehouse platform is built on Delta Lake, an open source project that delivers reliability, security and performance on the data lakes. The Databricks Lakehouse platform delivers performance at scale with optimizations such as caching, indexing and data compaction. Additionally, The the Databricks Lakehouse platform has Photon Engine, a vectorized query engine, that for SQL, further speeds SQL query performance at low cost, delivering business insights even sooner.

In our tests here at Intel, we ran a decision support benchmark against an Azure Databricks data engineering workload (Spark Cluster) using Databricks runtime version 9.0. The benchmark ran 99 queries against a 20-node Azure VM cluster. In addition to highlighting the Databricks Photon Engine, our goal was to show how quickly an E8ds_v4 VM cluster backed by 2nd Generation Intel Xeon Scalable processors could complete the data queries compared to an older E8s_v3 VM cluster. We also compared the E8ds_v4 cluster to an E8as_v4 cluster backed by AMD EPYC processors to show the performance and value advantages of Edsv4 VMs. For all of these comparisons, we tested both a 1TB and a 10TB dataset to illustrate how you can get better performance regardless of your dataset size.

Reduce your database query time with Databricks Photon Engine

The sooner data analytics queries complete, the faster you can implement the insights to improve and expand your business. Photon engine, built in C++, improves query performance in SQL and Spark, accelerating the analysis of your data. To demonstrate how well Photon can enhance query performance, we tested our Edsv4 cluster with Photon disabled and enabled. Figure 1 shows how the E8ds_v4 cluster with Photon enabled completed a 1TB dataset 2.82 in 65% less time than the same cluster without Photon, and completed a 10TB dataset in 62% less time.

Normalized processing time to complete queries on Databricks with vs. without Photon

Lower is better

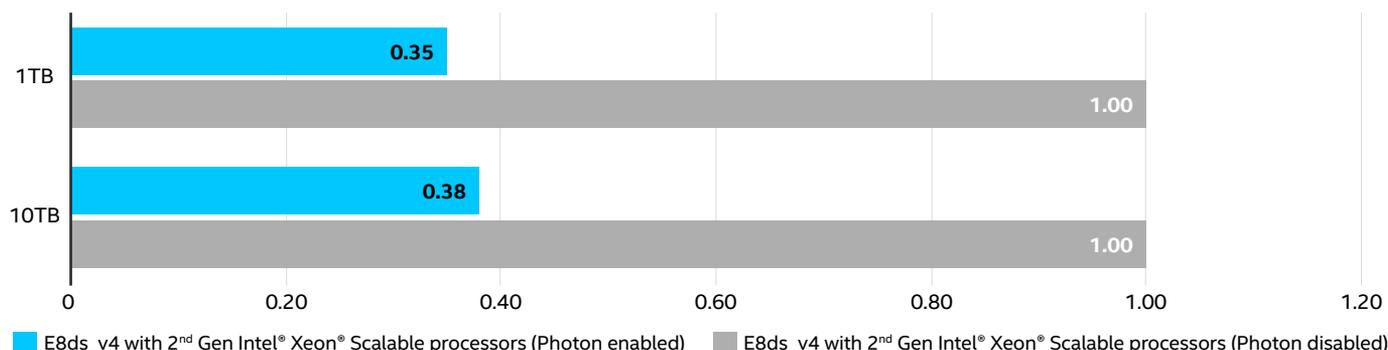


Figure 1. The relative processing time to complete the 99 decision support benchmark queries with Photon compared to without Photon on E8ds_v4 clusters on 1TB and 10TB datasets.

In addition to accelerating the time to insights, faster data querying also means less VM uptime. With Photon, the E8ds_v4 cluster would cost 35% less to run a 1TB dataset than the same cluster with Photon, and 30% less to run a 10TB dataset. Figure 2 compares the cost/performance for each cluster. As it shows, shorter run times translate to savings.

Normalized price/performance to run a decision support workload with vs. without Photon

Lower is better

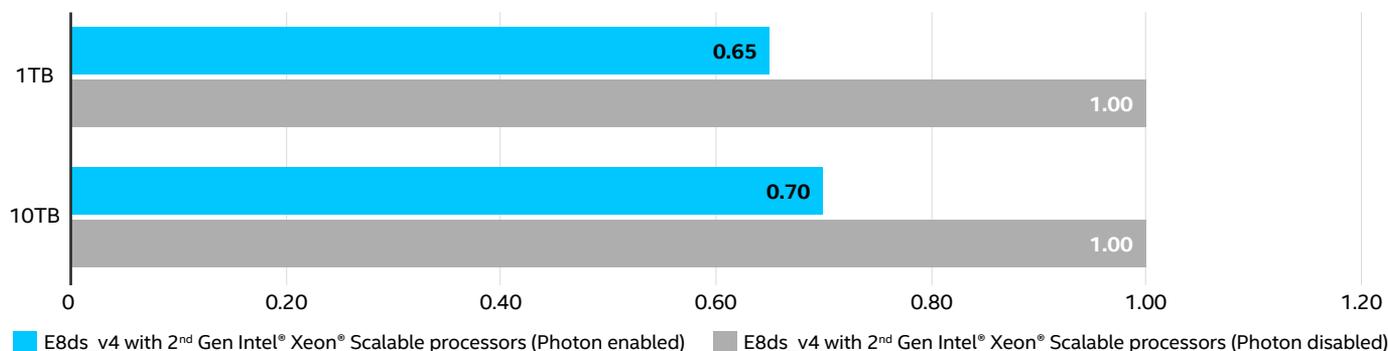


Figure 2. Normalized price/performance to run a decision support workload against a Databricks environment on Azure E8ds_v4 VMs on both 1TB and 10TB datasets.

Gen over Gen: Is Upgrading Worth It?

Gen over Gen Performance

Using the Databricks Photon Engine isn't the only way to ensure good performance. Choosing the right hardware also has an impact on the performance of your decision support workloads. While the concept that newer hardware leads to better performance is hardly novel, it's not always obvious how much of an improvement a workload will actually achieve on newer hardware. The improvement may be small enough that it doesn't seem worth the effort or cost increase of moving to newer hardware. While Azure removes a lot of the pain and effort from upgrading hardware – one of the many conveniences of running your workloads in the cloud – what about the cost? It's natural to assume that newer instances cost more. While is true in this case, the performance increase the newer hardware delivers yields a net win in terms of value, supporting the choice to upgrade.

To show readers what they may be missing by staying with older hardware, we created a hypothetical situation. We tested our decision support workload on a 20-node E8s_v3 cluster with Databricks Runtime 9.0 to get a baseline performance metric. The Esv3 series from Azure offer VMs with a range of processors from the Intel® Xeon® E5-2673 v4 to the Intel Xeon 8272CL. Which CPU you get when you spin up a VM is random, meaning a 20-node cluster could have a mix of CPU types, including some that are three Intel CPU generations older than the newest processors. For our tests, we ensured that all the E8s_v3 VMs listed the same Intel Xeon Platinum 8171M processor when we started our tests for consistency. We then tested the same workload on a 20-node E8ds_v4 cluster with Photon enabled. Azure guarantees that every Edsv4 VM uses an Intel Xeon Platinum 8272CL processor, offering consistent performance.

Enhance Databricks Data Warehouse Workloads with Azure E8ds_v4 VMs and Databricks Photon Engine

When we compare the performance of the two VM clusters, it's clear that upgrading to the newer E8ds_v4 VMs is a boon for performance regardless of dataset size. With the 1TB data set, the E8ds_v4 cluster query reduced completion time to only 26% of that of the E8s_v3 cluster. With the 10TB data set, the E8ds_v4 cluster query completion time was even lower, one-fifth that of the E8s_v3 cluster (see Figure 3).

Normalized processing time to complete queries on Databricks E8ds_v4 vs. E8s_v3

Lower is better

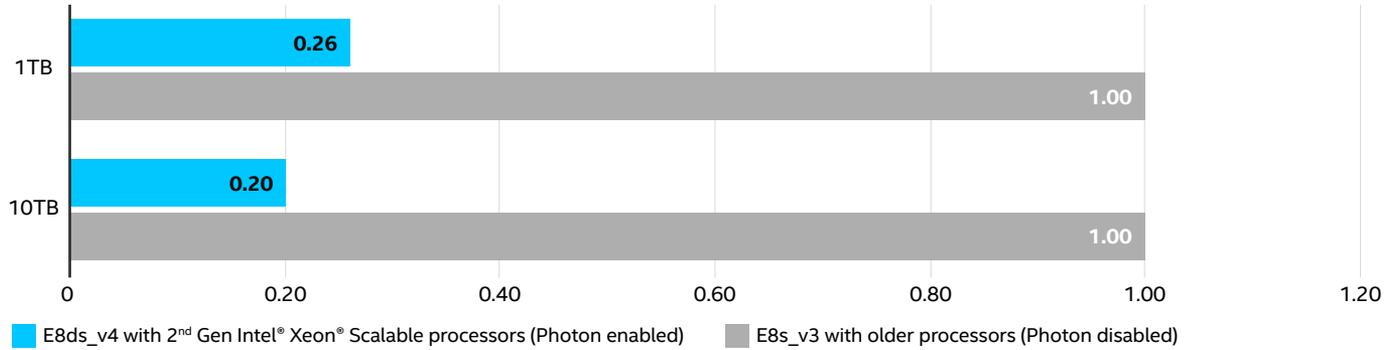


Figure 3. The relative processing time to complete the 99 decision support benchmark queries on an E8ds_v4 VM cluster with 2nd Gen Intel Xeon Scalable processors compared to an older E8s_v3 VM cluster on both 1TB and 10TB datasets.

Better Gen Over Gen Value

With performance wins such as those illustrated in Figure 3, it seems likely that the extra cost of upgrading from E8s_v3 VMs to E8ds_v4 VMs would be worth paying. Using the public price per hour at the time of testing, we determined the cost to execute each workload scenario. We converted the total query processing time from milliseconds to hours, combined the hourly cost of the instances, storage, and Databricks DBUs, and calculated the price per TB run for all four scenarios. As Figure 4 shows, running a decision support workload with a 1TB dataset would cost almost twice as much on the E8s_v3 cluster as it would cost on the E8ds_v4 cluster. Even more impressive, running the 10TB dataset on the E8ds_v4 cluster would cost well under half as much as it would on the older E8s_v3 cluster, a savings of 61%.

Normalized price/performance to run a decision support workload on E8ds_v4 vs. E8s_v3

Lower is better

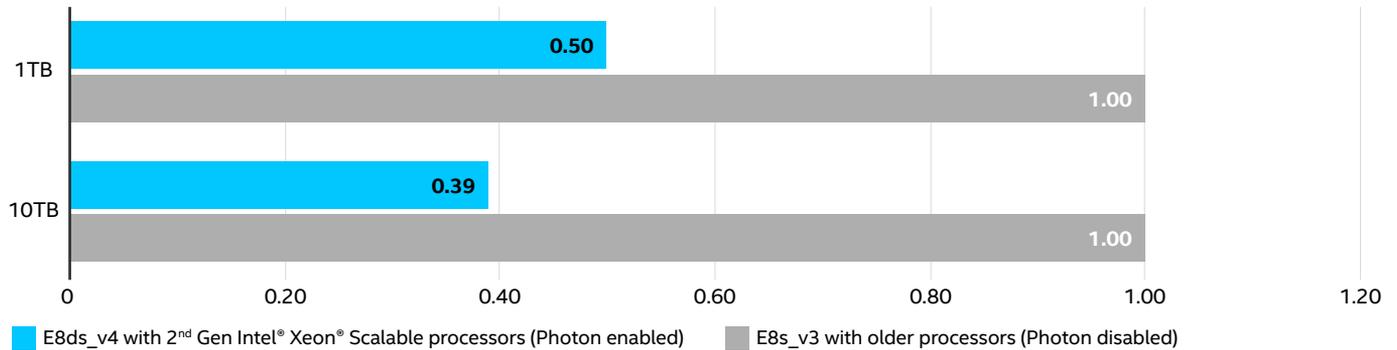


Figure 4. Normalized price/performance to run a decision support workload against a Databricks environment on Azure E8ds_v4 VMs compared to E8s_v3 VMs on both 1TB and 10TB datasets.

Competitive: But what about the E8as_v4 VMs?

Competitive performance

Okay, so we've convinced you to move on from the E8s_v3 VMs, but Azure also offers E8as_v4 VMs backed by AMD EPYC processors. How do those compare to the E8ds_v4 VMs backed by 2nd Generation Intel® Xeon® Scalable processors? We expanded our hypothetical scenario outlined in the previous section to answer that question. We tested the same 20-node Databricks Runtime 9.0 decision support workload on the E8as_v4 VMs and compared the results to the E8ds_v4 cluster. As you can see in Figure 5, the E8ds_v4 cluster backed by Intel processors outperformed the E8as_v4 cluster backed by AMD EPYC processors with both dataset sizes. The E8ds_v4 cluster completed the 99 queries on a 1TB dataset in 31% less time than the E8as_v4 cluster. And for the 10TB dataset, the E8ds_v4 cluster's time to complete was 23% lower than that of the E8as_v4 cluster.

Normalized processing time to complete queries on Databricks E8ds_v4 vs. E8as_v4

Lower is better

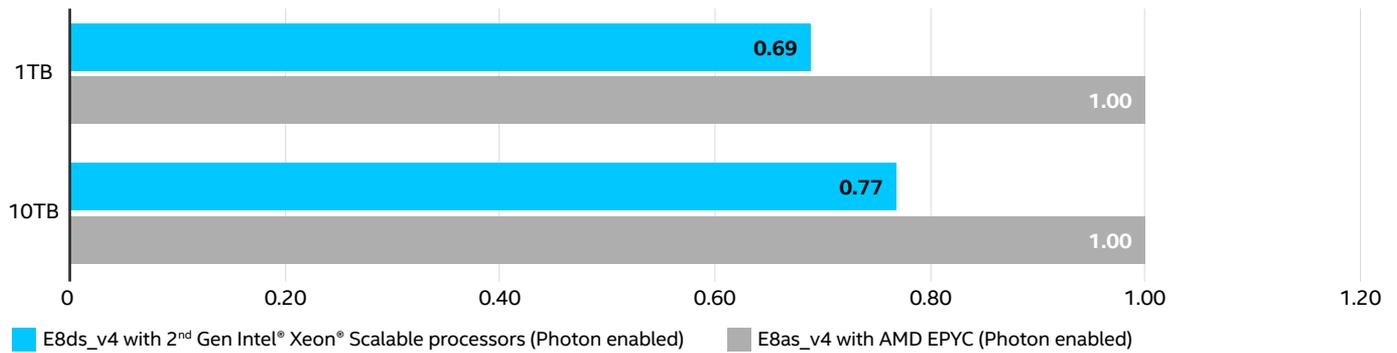


Figure 5: The relative processing time to complete 99 decision support queries on Databricks on an E8ds_v4 cluster backed by 2nd Gen Intel® Xeon® Scalable processors vs. an E8as_v4 cluster backed by AMD EPYC processors.

Competitive value

Finally, what about cost differences? Using the same methodology we outlined in the previous section, we looked at the price per TB run to get the relative value of the E8ds_v4 compared to the E8as_v4. Once again, as Figure 6 shows, the E8ds_v4 cluster backed by Intel processors offers the better value for Databricks decision support workloads. On 1TB datasets, the cost to run the workload on an E8ds_v4 cluster is 30% less than running the same workload on the E8as_v4 cluster. With a large 10TB dataset, the E8ds_v4 cost 22% less than the E8as_v4.

Normalized price/performance to run a decision support workload on E8ds_v4 vs. E8as_v4

Lower is better

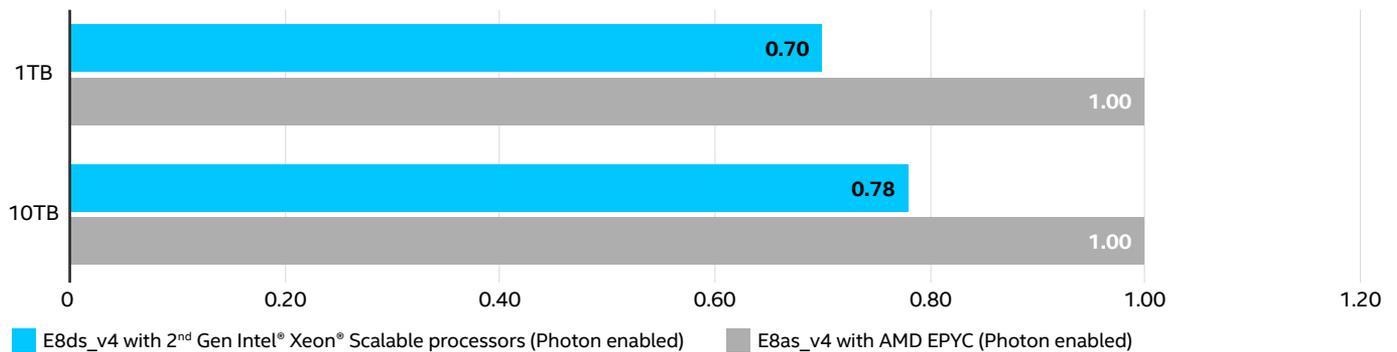


Figure 6. Normalized price/performance to run a decision support workload against a Databricks environment on Azure E8ds_v4 VMs compared to E8as_v4 VMs on both 1TB and 10TB datasets.

Additional competitive data: What about storage-optimized instances?

We performed one additional test comparing the E8ds_v4 instances featuring 2nd Gen Intel Xeon Scalable processors to the storage-optimized Lsv2 instances featuring AMD EPYC 7551 processors. Even here, the Intel-backed E8ds_v4 instances delivered advantages in both performance and price/performance. As Figure 7 shows, the E8ds_v4 cluster completed the queries in up to 38% less time than the L8s_v2 instances. When we calculate the price/performance difference, the E8ds_v4 cluster cost up to 39% less than the L8s_v2 cluster (see Figure 8).

Normalized processing time to complete queries on Databricks E8ds_v4 vs. L8s_v2

Lower is better

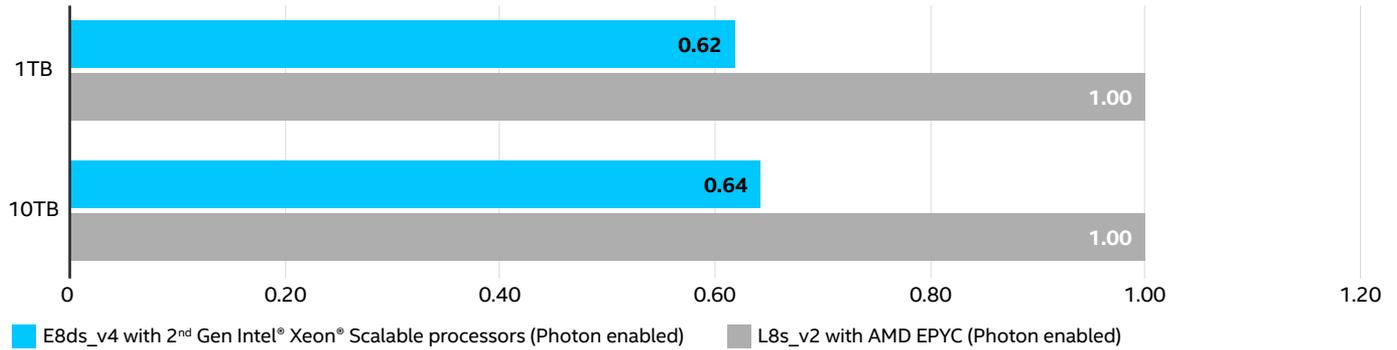


Figure 7: The relative processing time to complete 99 decision support queries on Databricks on an E8ds_v4 cluster backed by 2nd Gen Intel Xeon Scalable processors vs. an L8s_v2 cluster backed by AMD EPYC processors.

Normalized price/performance to run a decision support workload on E8ds_v4 vs. L8s_v2

Lower is better

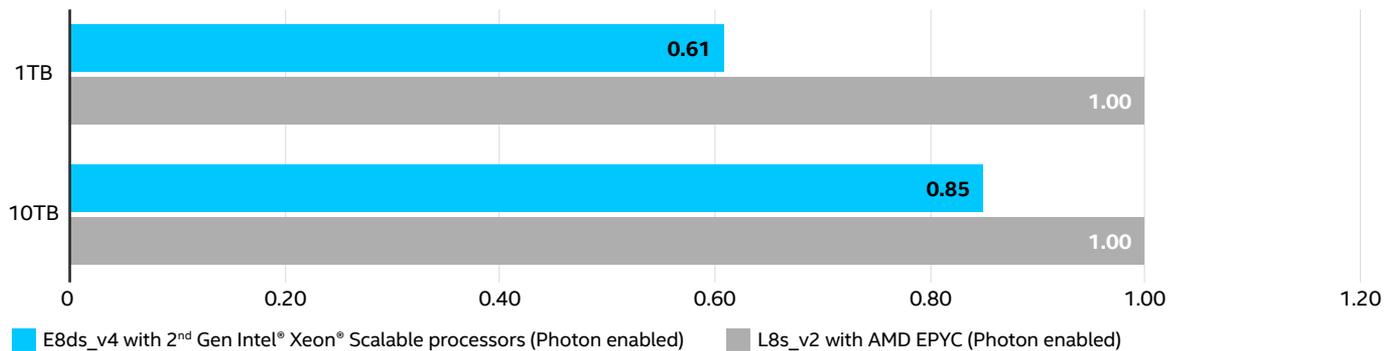


Figure 8. Normalized price/performance to run a decision support workload against a Databricks environment on Azure E8ds_v4 VMs compared to L8s_v2 VMs on both 1TB and 10TB datasets.

Our test environment

For all of the results listed in this report, we used a decision support workload and the Databricks Runtime version 9.0 that includes Apache Spark 3.12 and Scala 2.12. Workload parameters for all tests included the following:

- spark databricks passthrough enabled: true
- spark databricks adaptive autoOptimizeShuffle enabled: true
- spark databricks io cache maxMetaDataCache 10g
- spark databricks io cache max DiskUsage 100g
- spark databricks delta preview enabled: true

All VMs under test ran Ubuntu 20.04.1 kernel v5.4.0-1056-azure. Intel conducted the E8s_v3 vs. E8ds_v4, the E8ds_v4 with and without Photon, and the E8ds_v4 vs. E8as_v4 testing in October 2021 on the Azure US- East-2 region. Intel conducted the E8ds_v4 vs. Ls_v2 testing on November 17, 2021.

E8s_v3 configuration

The baseline E8s_v3 cluster configuration consisted of 1x Standard_E8s_v3 VM for the controller VM and 20x Standard_E8s_v3 VMs for worker VMs. While the E8s_v3 VMs have multiple processor options, our VMs had the Intel Xeon Platinum 8171M CPU @ 2.60GHz. The VMs also had 64GB memory each, and 128GB storage disks. The network BW/Instance (Mbps) was 4/4000 and the storage BW/Instance (Mbps) was 16,000/128 (200).

E8ds_v4 configuration

The E8ds_v4 cluster configuration consisted of 1x Standard_E8ds_v4 VM for the controller VM and 20x Standard_E8ds_v4 VMs for the worker VMs. Each E8ds_v4 VM consisted of an Intel Xeon Platinum 8272CL CPU @2.60GHz, 64GB of memory, and a 300GB storage disk. The network BW/Instance (Mbps) was 4/4000, and the storage BW/Instance (Mbps) was 77000/485 (200). The throughput/BW (cache size) for the storage-optimized comparison was 38000/500.

E8as_v4 configuration

The E8as_v4 cluster configuration consisted of 1x Standard_E8as_v4 VM for the controller VM and 20x Standard_E8as_v4 VMs for the worker VMs. Each E8as_v4 VM consisted of an AMD EPYC 7452 32-core processor, 64GB of memory, and a 128GB storage disk. The network BW/Instance (Mbps) was 4/3200 and the storage BW/Instance (Mbps) was 16000/128 (200).

Ls_v2 configuration

The Ls_v2 cluster configuration consisted of 1x Standard_Ls_v2 instance for the controller VM and 20x Standard_Ls_v2 instances for the worker VMs. Each Ls_v2 VM consisted of an AMD EPYC 7551 32-core processor, 64GB of memory, and an 80GB disk. Additionally, the instance came with a 1.92TB NVMe disk (400000/2000 Read IOPS/MBPs). The Network BW/Instance (Mbps) was 4/3200 and the Storage Throughput/BW (cache size) was 8000/160 (200).

Pricing

Pricing details are from <https://azure.microsoft.com/en-us/pricing/calculator/> and <https://azure.microsoft.com/en-us/pricing/details/databricks/>. Pricing as of the time of testing.

Azure VM	Hourly cost
Standard_E8s_v3 - 1st Generation Intel Xeon Platinum 8171M processors	\$0.532
Standard_E8ds_v4 - 2nd Generation Intel Xeon Platinum 8272CL processors	\$0.576
Standard_E8as_v4 - 2nd Generation AMD EPYC™ 7452 processor	\$0.504
Standard_Ls_v2 – 1st Generation AMD EPYC™ 7551 processor	\$0.624

Conclusion

Our Databricks testing with a decision support benchmark provides information that can help companies that are facing decisions about performance and cost when running their workloads in the cloud. If you are trying to determine the best instances for your decision support database workloads, the many choices available can seem overwhelming. Our tests show that selecting newer VMs can provide both better performance and better value. Additionally, choosing Intel-backed VMs in our testing outperformed the same generation AMD VMs. For your decision support database workloads, choose the E8ds_v4 VMs to get the most for your money.



Performance varies by use, configuration and other factors. Learn more at www.intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure. Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Printed in USA 0422/JO/PT/PDF US004

