



Run Databricks Queries in up to 76% Less Time and Reduce Costs With Amazon® R5d Instances Featuring 2nd Gen Intel® Xeon® Scalable Processors

With Photon Vectorized Query Engine Enabled, These Instances Dramatically Outperformed r5a.2xlarge Instances Featuring AMD EPYC™ Processors on Decision Support Workloads and Delivered Better Value

Many organizations rely on Databricks' Lakehouse Platform for storing and analyzing data, both structured and unstructured. To run your decision support queries quickly, it is important to select cloud instances backed by powerful hardware. But determining which instances meet this criterion can be a challenge.

We conducted tests to assist companies that are shopping for cloud instances for their decision support workloads. Specifically, we looked at AWS instance series: R5d instances enabled by 2nd Gen Intel® Xeon® Scalable processors and R5a instances with AMD EPYC processors. We created Databricks Runtime 9.0 clusters of these two instance types to run a decision support workload. On the R5d cluster, we used VMs that enabled a vectorized query engine called Photon designed to improve SQL query performance. At the time of this testing, Databricks' Photon engine is not supported on R5a instances.

R5d Instances Completed Decision Support Workloads in Less Time

We tested the two AWS instances with a decision support benchmark that generates a lower-is-better score that reflects the amount of time needed to execute a given set of queries. Selecting an instance that takes less time can help your company two ways: first, by getting valuable information sooner and second, reducing instance uptime and associated costs, which can help you spend less. As Figure 1 shows, r5d.2xlarge instances with 2nd Gen Intel Xeon Scalable processors and Photon enabled completed queries on a 1TB data set in 74% less time than r5a.2xlarge instances with AMD EPYC processors did. With a 10TB data set, query completion time of the r5d.2xlarge cluster was 76% shorter than that of the r5a.2xlarge cluster.

Normalized Processing Time to Complete Databricks Queries

Time | Lower is better




Figure 1. Relative processing time to complete a set of benchmark queries on a Photon-enabled r5d.2xlarge instance cluster with 2nd Gen Intel Xeon Scalable processors and an r5a.2xlarge cluster with AMD EPYC processors on both 1TB and 10TB data sets.

Intel Workload Proof Series: Databricks on Amazon R5d instances vs. R5a instances

Run Decision Support Queries in up to 76% Less Time with r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors with Photon enabled

vs. r5a.2xlarge instances with AMD EPYC processors

Spend up to 51% Less to Run Decision Support Queries with r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors with Photon enabled

vs. r5a.2xlarge instances with AMD EPYC processors

How Shorter Query Times Can Help Your Bottom Line

As is the case with any resource in which your company is investing, getting good value for your dollar is a priority. We calculated how much it would cost a company to perform the test scenarios we discussed on the previous page. We used the price per hour for each instance, storage, and Databricks DBUs at time of testing along with the times in Figure 1 to determine the price per TB for all four scenarios. As Figure 2 shows, a company would spend much less if they ran decision support workloads on Photon-enabled r5d.2xlarge instances. For the 1TB dataset, the r5d.2xlarge cluster enabled by 2nd Gen Intel® Xeon® Scalable processors could provide 46% lower price/performance than the r5a.2xlarge cluster with AMD EPYC processors did. For the 10TB dataset, the Photon-enabled r5d.2xlarge cluster would reduce price/performance costs by 51%.

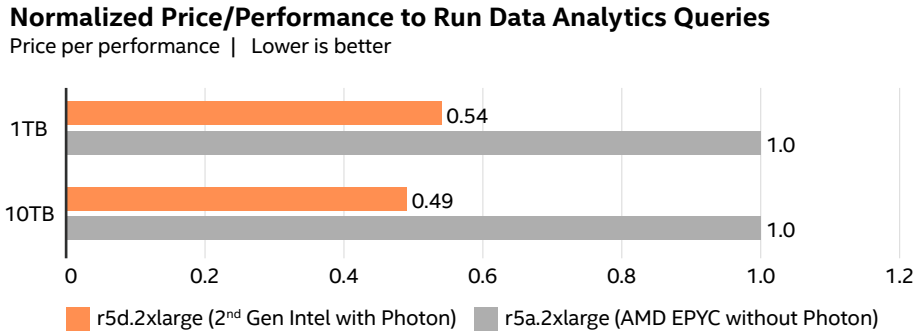


Figure 2. Normalized price/performance to run a decision support workload against a Databricks environment on Photon-enabled Amazon r5d.2xlarge instances compared to r5a.2xlarge instances on both 1TB and 10TB datasets.

Conclusion

We measured the time to complete a set of Databricks queries for two different data set sizes on Photon-enabled AWS r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors and r5a.2xlarge instances with AMD EPYC processors. The r5d.2xlarge instances completed sets of queries in up to 76% less time. When we combined these times with the hourly pricing for the two instances, we found that the r5d.2xlarge instances cost considerably less to execute the same amount of work—a cost savings up to 51%. If your company wants to get actionable insights earlier and reduce spending on AWS instances, choose Photon-enabled r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors.

Learn More

To begin running your Databricks clusters on Photon-enabled Amazon R5d instances with 2nd Gen Intel Xeon Scalable processors, visit <https://aws.amazon.com/quickstart/architecture/databricks/>.

To learn more about Databricks' Photon Vectorized Query Engine, visit <https://databricks.com/product/photon> and <https://docs.databricks.com/runtime/photon.html>.

For all of the results in this report, we used a decision support workload derived from TPC-DS. All tests were conducted in December 2021 on the us-east-1 AWS region. All tests used 20-node clusters with Ubuntu 18.04.1, kernel version 5.4.0-1059-AWS, Databricks 9.0, Apache Spark 3.1.2, Scala 2.12. Both instance types had 8 vCPUs and 64GB RAM. The r5d.2xlarge had a 300GB NVMe SSD, 10 Gbps Network BW, and 4,750 Mbps Storage BW. The r5a.2xlarge instances had a 250GB EBS volume, 10Gbps Network BW, and 2,880 Mbps Storage BW.



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 0322/JO/PT/PDF US003

