

Run Your Databricks Queries in up to 71% Less Time and Reduce Costs: Select Amazon[®] R5d Instances Featuring 2nd Gen Intel[®] Xeon[®] Scalable Processors



Databricks



Run Decision Support Queries in up to 71% Less Time with r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors with Photon enabled

vs. r6gd.2xlarge instances with Amazon Graviton2 processors



Spend up to 33% Less to Run Decision Support Queries with r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors with Photon enabled

vs. r6gd.2xlarge instances with Amazon Graviton2 processors

With Databricks' Photon Vectorized Query Engine Enabled, These Instances Delivered Faster Performance and Better Value than R6gd Instances Featuring Amazon Graviton2 Processors on Decision Support Workloads

Databricks' Lakehouse architecture allows organizations to store and analyze great volumes of structured and unstructured data. To make better decisions earlier, it is smart to select cloud instances that can complete queries quickly, which improves productivity and fosters better collaboration between data engineers, business analysts, and data scientists. But how can you know which instances can give you this advantage?

We've carried out a series of tests to answer this question. For companies that are seeking cloud instances on which to run their data warehousing/decision support workloads, we tested two AWS instance series: R5d instances enabled by 2nd Gen Intel[®] Xeon[®] Scalable processors and R6gd instances enabled by Amazon Graviton2 processors. We ran a decision support workload on clusters of these two instance types with Databricks Runtime 9.0. On the R5d cluster, we enabled the Databricks' Photon engine, designed to accelerate SQL query performance. At the time of this testing, Databricks' Photon engine is not supported on R6gd instances.

The R5d instances with 2nd Gen Intel Xeon processors and Photon enabled completed the queries in less time than the R6gd instances. Selecting these instances can deliver actionable information to decision-makers earlier and reduce infrastructure costs.

R5d Instances Needed a Fraction of the time to Execute Queries

To compare the performance of the two AWS instance series, we used a benchmark that measures how long each required to execute a set of database queries. 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 71% less time than r6gd.2xlarge instances with Amazon Graviton2 processors needed. With a 10TB data set, query completion time of the r5d.2xlarge cluster was 66% shorter.

This means that organizations looking to quickly gain insights from data can meet that goal by selecting Amazon R5d instances featuring updated 2nd Generation Intel Xeon Scalable processors.

Speedier Instances Can Save Your Company Money

To understand how the performance differences between these two AWS instance series affect your bottom line, we calculated the cost per terabyte to perform our test scenarios on each. We used the relative query completion times in Figure 1 along with the price per hour for each instance, storage, and Databricks DBUs at time of testing. As Figure 2 shows, a company could run decision support workloads on Photon-enabled r5d.2xlarge instances for considerably less. For the 1TB dataset, a company could spend 23% less for a given level of performance with the Photon-enabled r5d.2xlarge cluster enabled by 2nd Gen Intel® Xeon® Scalable processors than they would for the r6gd.2xlarge cluster with Amazon Graviton2 processors. For the 10TB dataset, the cost reductions would be 33%.

Conclusion

We performed a set of Databricks queries on two different-sized data sets on two AWS instances: Photon-enabled AWS r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors and r6gd.2xlarge instances featuring Amazon Graviton2 processors. The r5d.2xlarge needed up to 71% less time to carry out the workloads. Combining these times with the two instances' hourly pricing, we discovered that the r5d.2xlarge instances cost as much as 33% less to carry out a fixed amount of work. By opting for Photon-enabled r5d.2xlarge instances featuring 2nd Gen Intel Xeon Scalable processors, you organization can obtain vital insights sooner while also spending less.

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>.

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 r6gd.2xlarge cluster with Amazon Graviton2 processors on 1TB and 10TB data sets.

Normalized Price/Performance to Run Data Analytics Queries

Price per performance | Lower is better




Figure 2. Normalized price/performance to run a decision support workload against a Databricks environment on Photon-enabled Amazon r5d.2xlarge instances compared to r6gd.2xlarge instances on 1TB and 10TB datasets.

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 r6gd.2xlarge instances had a 474GB NVMe SSD, 10 Gbps Network BW, and 4,750 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 US004

