

Analyze 1.51x the Data per Second for Apache Spark™ Workloads with Google Cloud N2 VM Instances

Get Insights from Data Faster with Google Cloud N2 VM Instances Featuring 2nd Gen Intel® Xeon® Scalable Processors

Data analytics, machine learning, and artificial intelligence workloads sort and make sense of massive amounts of company data. Because such workloads are compute-intensive, hosting them in the cloud requires hardware that can meet the unique challenges inherent in gaining quick insight from data. If your organization runs Apache Spark clusters—which provides a unified analytics platform for processing big data—the type of VM instance you choose directly affects the speed at which you can turn data into action. For demanding Apache Spark workloads, Google Cloud N2 VM instances enabled by 2nd Generation Intel Xeon Scalable processors can boost machine learning workloads for data analysis to ensure your organization can move forward with new initiatives quickly.

In a comparison of k-means clustering machine learning performance, newer Google Cloud N2 VM instances enabled by 2nd Generation Intel Xeon Scalable processors outperformed older N1 VM instances with Intel Xeon E5 v4 processors, analyzing up to 1.51x the data per second for Apache Spark advanced data analytics workloads.

Analyze More Data per Second with 1.51x the Throughput

Cloud instances increasingly support compute-intensive workloads such as data analytics, machine learning, and artificial intelligence (AI), which is a boon for organizations that desire the agility and flexibility that cloud computing brings to the rest of their business.

Tests comparing k-means clustering performance using the HiBench benchmark suite show that Google Cloud N2 VM instances featuring 2nd Generation Intel Xeon Scalable processors can boost Apache Spark machine learning workloads by up to 1.51x the data per second of N1 VM instances with Intel Xeon E5 v4 processors. This means that choosing Google Cloud N2 VM instances over N1 VM instances can dramatically improve the speed at which your organization can process data

Relative throughput of Google Cloud N1 vs. N2 VM instances




Figure 1. Relative throughput comparison of Google Cloud N1 and N2 VM instances (32 vCPU/128GB RAM) for k-means clustering workloads from the HiBench benchmark suite.

Apache Spark

Analyze more data per second with 1.51x the throughput

Optimize data analytics and machine learning in the cloud

Accelerate workloads with Intel Optane persistent memory



New Technologies in 2nd Gen Intel Xeon Scalable Processors Optimize Data Analytics

Selecting Google Cloud N2 VM instances means you access key technologies embedded in 2nd Generation Intel® Xeon® Scalable processors that work to optimize compute-intensive advanced data analytics and machine learning workloads. Intel Advanced Vector Extensions 512 (Intel AVX-512) processes more data with a single instruction, handling computational tasks more quickly and accelerating these advanced workloads. Google Cloud N2 VM instances featuring 2nd Generation Intel Xeon Scalable processors offer Intel Deep Learning Boost (Intel DL Boost), which optimizes neural network loops to improve AI operations by as much as 3.4x, making it a formidable tool for optimizing advanced data analytics and machine learning workloads. To further flexibility, Google Cloud N2 VM instances enabled by 2nd Gen Intel Xeon Scalable processors also utilize Intel field-programmable gate arrays (FPGAs), which can be reconfigured for different ML models where algorithms frequently change to speed up AI operations.

Accelerate Workloads with Intel Optane Persistent Memory

Google Cloud N2 VM instances enabled by 2nd Gen Intel Xeon Scalable processors also introduce the option to further accelerate data analytics workloads with Intel Optane persistent memory (PMem). Intel Optane PMem is a combination memory and storage technology that allows Google Cloud N2 VM instances to place more data closer to the processor to deliver lower latencies for data analytics and machine learning.

Learn More

To begin your Apache Spark machine learning workloads on Google Cloud VM instances featuring 2nd Generation Intel Xeon Scalable processors, visit <http://intel.com/GoogleCloud>.



Performance varies by use, configuration and other factors. Learn more at <https://intel.com/benchmarks>.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. 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 0321/JO/PT/PDF US001

Please Recycle