

Increasing HDFS Performance with PowerScale for Google Cloud

Dell Technologies Cloud PowerScale for Google Cloud

Highlights

Fully integrated Google Cloud native service operated by Dell Services and backed by enterprise SLAs

Filesystem scales up to a massive 50 PB in a single namespace

Hadoop applications continue to use HDFS instead of using cloud storage connectors for object storage

Every PowerScale for Google Cloud node is a Namenode and Datanode providing enhanced HDFS resiliency with integrated load balancing

Scale HDFS storage independently from Google Compute Engine (Hadoop nodes)

Enterprise-class data features include multiprotocol access, Microsoft® Active Directory integration, native replication, and snapshots

POSIX-compliant file system

Easily combined with Google Cloud Compute Engine and analytics services

Guaranteed and predictable pricing based on four performance-optimized tiers

24x7 proactive monitoring, automated provisioning, and maintenance

Ordered and managed on Google Cloud

Billed and supported by Google

Business challenge

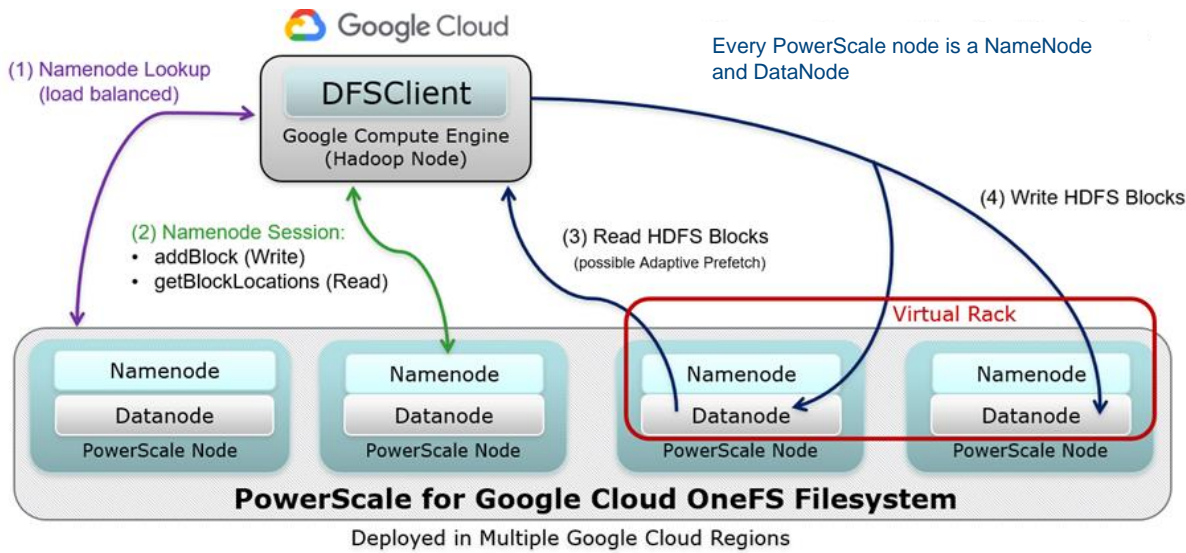
Apache® Hadoop® is a well-known big data platform that can scale across multiple compute nodes to process large datasets. The elastic capabilities of cloud computing make the cloud an excellent environment to run Hadoop. However, running the Hadoop Distributed File System (HDFS) in the cloud has its limitations. Cloud-native Hadoop services use cloud object storage by default instead of HDFS. Object storage has more I/O variance, no support for file appends or truncates, and is not POSIX-compliant. These factors can be a problem for many customers who want to migrate Hadoop workloads to the cloud. Also, as datasets increase in size, cloud object storage becomes more latent, often causing extensive runtimes.

Dell Technologies™ Cloud PowerScale™ for Google Cloud is a Google native scalable file service that provides high-speed file access over SMB, NFS, and HDFS. It is powered by the Dell EMC™ PowerScale family which includes Dell EMC Isilon™ and PowerScale nodes, and it includes PowerScale OneFS™ which runs across these systems.

Disaggregating HDFS in the cloud

PowerScale for Google Cloud enables customers to separate and tier HDFS storage from the Hadoop compute infrastructure. Each PowerScale node boosts performance and expands the Hadoop cluster storage capacity. For Hadoop analytics, PowerScale for Google Cloud minimizes bottlenecks, rapidly serves big data, and optimizes performance with integrated NameNode and DataNode load balancing.

PowerScale for Google Cloud serves as the file system for Hadoop. With PowerScale for Google Cloud, HDFS is supported as a lightweight protocol at the storage layer. Hadoop compute clients (DFS clients) can access data that is stored in PowerScale for Google Cloud by connecting to any PowerScale node over the HDFS protocol. All PowerScale nodes configured for HDFS provide NameNode and DataNode functionality as shown in the following figure.



High-performance HDFS access in the cloud

Hadoop distributions supported by PowerScale for Google Cloud

You can run Apache Hadoop, Hortonworks HDP, or Cloudera CDH with PowerScale for Google Cloud. Cloudera CDP will be supported with release 7.2. Go to DellTechnologies.com to view the compatibility matrix.


Hadoop benchmarks

TeraGen, TeraSort, and TeraValidate are well-known Hadoop benchmarks. TeraGen generates random data, Terasort sorts the data, and TeraValidate validates the sorted output.

We compared the performance benefits of using PowerScale for Google Cloud HDFS over using cloud object storage. The exact same number and type of Google compute instances were used to create two Hadoop clusters. One cluster used PowerScale for Google Cloud HDFS storage, the other cluster used Google Cloud Storage.

The Google Compute Engine VM instances used for both Hadoop clusters were N1-STANDARD-96. These instances have 96 CPUs and 360 GB RAM. The Hadoop clusters were configured with one master node and eight worker nodes.

The Hadoop replication count with the Google storage cluster defaults to two when using the Google Dataproc Hadoop service. The Cloudera HDP 3.1 cluster used with Cloud OneFS has a replication count of three which is normal for production clusters.




1 TB Hadoop benchmarks (runtimes logarithmically scaled)

Scaling HDFS performance with PowerScale for Google Cloud

To test the scalability of PowerScale for Google Cloud, the worker instances on the Cloudera HDP 3.1 cluster were doubled without changing the configuration or node counts of PowerScale for Google Cloud. These tests demonstrate how well PowerScale for Google Cloud can accommodate the additional load and sustain performance with TeraGen, TeraSort, and TeraValidate benchmarks.

The logarithmic graph shows that with PowerScale for Google Cloud, the TeraGen benchmark completed in roughly half the amount of time compared to running the benchmark using Google Cloud Storage. The TeraSort benchmark is also better with PowerScale for Google Cloud and completed in roughly one third of the time. The size of Google Cloud Storage is not publically disclosed. The number of tier 1 PowerScale nodes used for this test was only eight. The Google Cloud region used for the testing is US-EAST4.

The next graph shows that as more compute instances are added to the Hadoop cluster, benchmark runtimes are improved with PowerScale for Google Cloud. Customers can grow the number of PowerScale nodes to meet growing Hadoop and storage requirements as needed.



PowerScale for Google Cloud 1 TB benchmarks

PowerScale for Google Cloud scales out horizontally like Hadoop. The more PowerScale nodes that are used, the more bandwidth and storage capacity is possible.

Hadoop can use substantial compute cycles for I/O processing, including running HDFS and writing and reading files. With PowerScale for Google Cloud, the HDFS processing functions are offloaded, and more compute cycles are available to run Hadoop services and related compute functions.

Why this matters

For data-intensive workloads such as IoT and big data analytics, most cloud solutions cannot adequately scale the file system size or achieve sufficient performance. Also, the storage costs that are required to accommodate these workloads are often prohibitive. Using the industry-standard IOzone benchmark utility, ESG validated that Dell Technologies Cloud PowerScale for Google Cloud achieved a maximum read performance of 200 GB/sec and a maximum write performance of 120 GB/sec against a 2 PB storage volume. Because this is a scale-out service, expanding the file system would result in even higher performance. PowerScale for Google Cloud supports file-system sizes of up to 50 PB.

PowerScale for Google Cloud enterprise features for Hadoop

Since all PowerScale nodes function as a NameNode, a secondary NameNode or NameNode High Availability (HA) is not required. Any PowerScale node functions as a NameNode and provides continued data access without complex HA configuration. This increases operational availability and decreases administrative overhead.

PowerScale for Google Cloud comes with SmartConnect which provides a DNS namespace to access PowerScale for Google Cloud file system. SmartConnect transparently manages the load balancing of client connections and failover of IP addresses to maintain connectivity and data access for Hadoop.

DataNode load balancing is integrated with PowerScale for Google Cloud. HDFS clients are automatically connected to the PowerScale node that has the least load to achieve optimal performance and throughput. This removes the requirement to define racks to improve load balancing. Also, PowerScale for Google Cloud supports Pipeline Recovery. This feature allows Hadoop clients to continue writing data to the PowerScale for Google Cloud file system if an active PowerScale node becomes unreachable or returns an error, which improves HDFS resiliency.

PowerScale for Google Cloud allows the configuration of multiple authentication sources. Customers can use Active Directory, MIT Kerberos, NIS, and LDAP if needed.

For customers using PowerScale clusters on premises, SyncIQ can be used to provide high-performance asynchronous replication of data to and from the cloud.

PowerScale for Google Cloud: The ultimate file solution for the cloud

While many block and object storage workloads have found a home in the cloud, file workloads have been missing due to performance, scalability, and cost limitations. With the right capabilities, the cloud should work well for large file workloads that depend on a lot of data. Some of these workloads need only occasional bursts to scalable infrastructure—a common cloud use case—but only if performance objectives can be met.

PowerScale for Google Cloud brings an industry-leading file system to the cloud with the levels of scalability and performance that big data platforms demand. Organizations can take advantage of the cost efficiency and flexibility aspects of the cloud, saving money while making data-driven decisions with predictable I/O performance. PowerScale for Google Cloud provides high performance, high scalability, and enterprise features that enable customers to use the same protocols and applications in the cloud that they use on-premises.

See raw data and sources at <https://github.com/bonibruno/dataproc>.



[Learn more](#) about Dell
EMC storage



[Contact](#) a Dell
Technologies expert