Technical Review

# Solving Big Data Challenges with Multi-Cloud Data Services for Dell EMC PowerScale

**Date:** April 2021 **Author:** Alex Arcilla, Validation Analyst

## Executive Summary

This Technical Review documents ESG's audit of testing performed with Dell EMC PowerScale for Multi-Cloud, a part of the Multi-Cloud Data Services for Dell EMC PowerScale, enabled by Faction, a multi-cloud data services provider. Testing was designed to demonstrate how the solution maximizes performance and scalability of cloud-based workloads requiring the use of large data sets. We focused on genomics processing as our use case to illustrate how the solution can maximize workload performance and availability while controlling overall cloud-related costs.

***What we found***. ESG validated that the Dell EMC PowerScale for Multi-Cloud solution can provide cost-effective storage performance at scale and scalable compute performance, leveraging multiple public clouds simultaneously. While the solution can be leveraged for various workload types, ESG verified the solution's benefits for large datasets using genomics processing as a use case, as it is a well-known data-intensive workload.

PowerScale for Multi-Cloud, designed to scale up to 90GB/s per file system, was observed by ESG to be one of the most performant cloud file services in today's market. The solution is especially useful for today's modern, data-intensive workloads.

Since the solution leverages close network proximity and efficiency delivered from cloud-adjacent data centers, organizations can utilize compute capacity or cloud-native services across multiple clouds to maximize application performance and availability while optimizing overall costs. No longer do organizations have to relegate workloads to any single public cloud service provider (CSP).

With PowerScale for Multi-Cloud, organizations gain advantages of Dell EMC Storage such as high availability for business continuity, data resiliency, and flexible scalability, coupled with the scalability and economic benefits of public cloud compute capacity. If your organization needs to efficiently leverage multiple CSPs to solve your big data challenges, ESG recommends taking a closer look at Dell EMC PowerScale for Multi-Cloud.

### ESG-validated Benefits

- **Scale:** One million CUDA cores used simultaneously across several public clouds with multi-petabyte file system scale.
- **Performance:** up to 160 Gb/sec aggregate throughput for processing a 70TB data set.
- **Compute Cost Savings:** up to 59% lower costs for GPU-based compute instances than on-premises GPU-enabled servers over a two-year period.
- **Storage Cost savings**: 89% lower costs when using PowerScale for Multi-Cloud with a single data set versus native file storage with the same data set on multiple clouds.
- **Analytics Efficiency**: Easily leverage cloud-native AI/ML services across multiple CSPs while accessing a single data repository.

## Background

Organizations have increasingly leveraged public cloud infrastructure to run a myriad of applications, such as file-based storage. In fact, ESG research has uncovered that 78% of organizations surveyed are currently using public cloud infrastructure services, up from 67% in 2020[1]. However, many find numerous challenges in using the public cloud effectively, including egress fees (31%), data copies located in multiple locations (25%), sizing workloads properly for cost (24%), and sizing workloads properly for performance (23%), as shown in Figure 1.[2]

**Figure 1. Top Ten Challenges of Leveraging Public Cloud Infrastructure Services**


*Source: Enterprise Strategy Group*

Modern workloads with large file-based data sets, such as big data analytics, artificial intelligence (AI), and machine learning (ML), typically require large amounts of compute power and storage. However, the performance and scale limitations of current public cloud infrastructure services for handling such workloads have prohibited organizations from moving these applications to the cloud. Maximizing application performance and availability while optimizing costs become real concerns.

While organizations can choose compute and storage services from any single CSP, the capabilities differ based on the underlying hardware and software infrastructure and the global regions in which those services are offered. These differences can lead to varying levels of workload performance and overall cost.

A good example is the graphical processing unit (GPU)-based compute instances (virtual machines or VMs) that are typically used for compute processing intensive workloads. While all CSPs offer GPU-based instances, today's CSPs do not use the same GPU hardware. These instance types may also not be available in every region in which these CSPs operate. Thus, the number and type of GPU-based instances used for a given workload will differ between CSPs. It becomes difficult to forecast which CSP will provide the desired performance and scalability, and subsequently, the optimal cost profile.

Organizations have to also closely consider the amount and type of cloud storage that these compute-intensive workloads use for data access and storage. Genomics processing presents a unique case in which more data is better for research.

[1] Source: ESG Master Survey Results, *2021 Technology Spending Intentions Survey*, December 2020.
[2] Source: ESG Master Survey Results, *2019 Data Storage Trends*, November 2019.

Such data sets can easily reach the petabyte (PB) range. Organizations can either store data in a single cloud and repatriate to process it on premises or move volumes of data stored on premises to the cloud and process it using compute capacity and cloud-native services. In both cases, it is difficult to move data in a timely fashion, thus affecting overall application performance. Additionally, organizations that want to leverage the capacity and capabilities of multiple clouds must move and replicate data between multiple CSPs, resulting in additional storage costs and egress charges, along with the added time and complexity of data management.

In organizations that opt to use native cloud services, different application teams may want access to the data in different formats. While file protocols such as NFS and CIFS are common for many applications, other teams working on cloud-native or big data applications look to leverage protocols such as HDFS or object (e.g., S3) in their applications.

Optimizing the performance and availability of a given workload while minimizing overall costs requires a solution that will enable the usage of different instance types and cloud-native services across multiple CSPs simultaneously, while enabling access to a common data repository employing the desired file or object-based data format.

## The Solution: Multi-Cloud Data Services for Dell EMC PowerScale

Multi-Cloud Data Services for Dell EMC PowerScale combines Dell EMC PowerScale with Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) to offer enterprise-grade scale-out file storage for operational flexibility and scalability for a myriad of data-intensive, high I/O throughput workloads, including high-performance computing (HPC) applications such as genomics processing and AI/ML (see Figure 2).

Dell Technologies has partnered with Faction, a multi-cloud data services provider, to enable practical use of large data sets (file or object based) in single or multi-cloud scenarios, while improving overall performance and lowering costs. Faction's Cloud Control Volumes (CCVs) reside on Dell EMC PowerScale storage, which are Dell EMC Isilon nodes managed by Faction in data centers selected for low-latency connections to major cloud service providers. Organizations can access data via a single-namespace, multi-petabyte repository, supporting file and object-based data formats. Faction's technology resolves potential network conflicts, especially when compute instances from multiple CSPs are accessing the same data files and connecting the environments to on-premises data centers. Today, PowerScale for Multi-Cloud can offer up to tens of petabytes of usable storage capacity in a single, multi-protocol file system. The service also supports bi-directional replication with on-premises PowerScale.

**Figure 2. Architecture of Multi-Cloud Data Services for Dell EMC PowerScale**

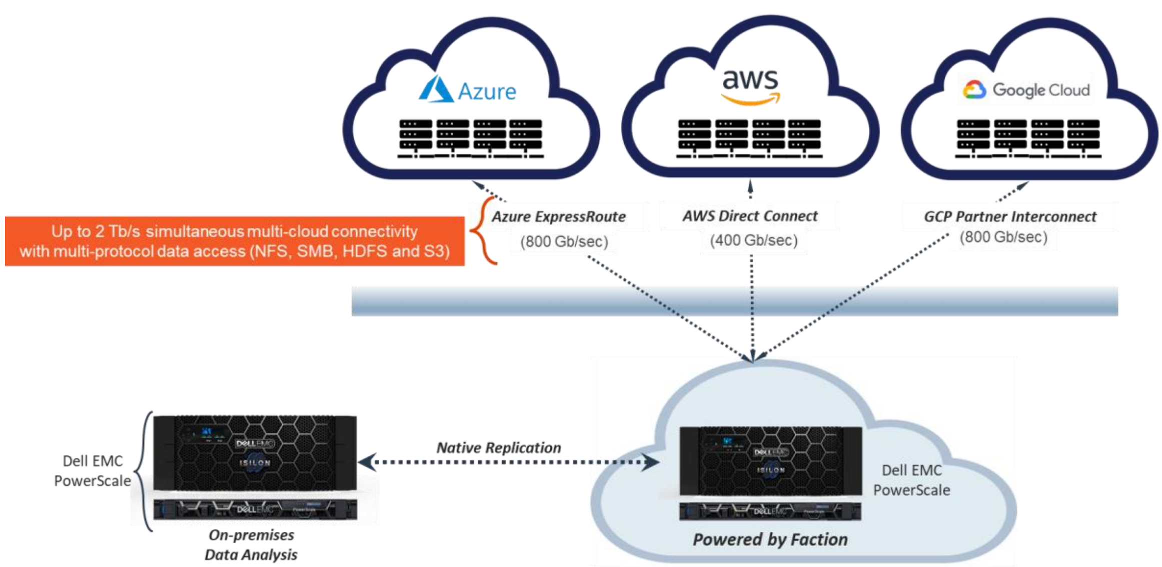


*Source: Enterprise Strategy Group*

By accessing data with one centralized namespace, updates are instantaneously visible to all instances across multiple CSPs and can be accessed quickly without incurring additional costs for replicating data, storing copies, and maintaining data consistency. This is especially important as data related to HPC applications tend to grow continually since data is not discarded.

PowerScale for Multi-Cloud also enables high-bandwidth, low-latency connectivity to each CSP using dedicated network links—AWS Direct Connect, Azure ExpressRoute Local, and GCP Partner Interconnect. Via the Faction Internetwork Exchange (FIX), the solution maintains LAN-level latency through physical proximity, capable of reaching up to 2 Tb/sec aggregate throughput from the three CSPs. The costs of the fully managed network connection and equipment are included in the PowerScale for Multi-Cloud service. Additionally, there are no charges for data movement between a customer's on-premises PowerScale and the Faction Cloud. Standard egress charges from CSPs apply, except for Microsoft Azure, which does not charge egress fees.

Additionally, the Faction platform supports self-service bandwidth re-provisioning through the service portal, enabling re-allocation of the supported service bandwidth between different CSPs as needs change. Organizations gain added flexibility to maximize the utility of their data across multiple public clouds.

With these dedicated interconnects between the cloud data service and the CSPs, organizations can leverage resources and services from multiple CSPs for a workload without copying or moving data. Access to this single data repository helps organizations to better control overall performance, availability and, most importantly, their cloud expenses.

## ESG Tested

ESG first examined how PowerScale for Multi-Cloud can maximize performance and availability of HPC and AI/ML applications by focusing our tests on a specific use case, genomics processing. Due to the large data sets involved, we chose genomics processing to illustrate the value that organizations can extract using this solution for enterprise-level AI/ML applications. The goal of the testing was to observe increased performance and reduced time to insight when analyzing genomics data using resources from three CSPs while accessing the same data set simultaneously.

Very large genomics data sets require fast processing to help solve complex healthcare issues that can save or improve lives. While the speed of sequencing a human genome has grown by orders of magnitude in the past 15 years, primary analysis of the raw genome of a single human can take over 30 hours[3] with a modern, high-end server system. The challenge then lies in the secondary analysis of these genomes, which can be used over and over again for research purposes. As more genome sequences are generated, secondary analysis identifies the variants (differences) between these genomes and other population genomes, or reference DNA sequences. Variant analysis supports the development of novel treatments for diseases, along with diagnosis and prevention.

Without adequate compute and storage resources, secondary analysis cannot keep pace with the generation of genomic data. To speed up secondary analysis, graphical processing unit (GPU)-based cloud compute instances can accelerate time to insight.

ESG audited the results of tests performed in the environment shown in Figure 3. Located on the US East Coast (Virginia), the test bed consisted of Dell EMC PowerScale storage systems, as part of PowerScale for Multi-Cloud, containing a 70TB genomics data set[4] on a single file system. The storage volume was interconnected to virtual networks in AWS, Azure, and GCP with AWS Direct Connect, Azure ExpressRoute, and GCP Partner Interconnect via the FIX. Maximum throughput for the provisioned AWS and GCP direct links were 20 Gb/sec per virtual private cloud. For the Azure link, maximum throughput was up to 200 Gb/sec per virtual network. A single mount point was established using the NFS protocol.

---

[3] Source: Dell Technologies White Paper, *Accelerating Next Generation Sequencing Secondary Analysis with Dell EMC Isilon and NVIDIA Parabricks*, 2020.

[4] Dell Technologies and Faction wish to acknowledge the source of the data used for our testing - 1000 Genomes Cohort - generated at the New York Genome Center with funds provided by NHGRI Grant 3UM1HG008901-03S1 (available from ENA Project ID: PRJEB31736).

In each CSP, GPU-enabled instances were deployed. All instances were powered by NVIDIA Tesla multicore GPUs, leveraging proprietary NVIDIA technology, Compute Unified Device Architecture (CUDA), for parallel processing[5]. Because CSPs do not offer the same GPU instance type in all regions, the test bed used up to 100 NVIDIA P100 GPU-enabled Azure VMs, 20 NVIDIA T4 GPU-enabled GCP instances, and 30 NVIDIA V100 GPU-enabled AWS EC2 instances.[6] Both on-demand and spot instances were employed for this test to ensure that our analysis was completed within a certain cost threshold.[7] CSP autoscaling was enabled to rapidly scale up the number of instances if additional processing power (supplied by the spot instances) was required.

All compute instances ran NVIDIA Clara Parabricks software[8] to complete the secondary analysis. The test environment also utilized grid engine software, located in Microsoft Azure, to control distribution of the jobs associated with the analysis and spin up compute instances across the three CSPs as the workload required.

**Figure 3. Test Bed for Genomics Processing Use Case**

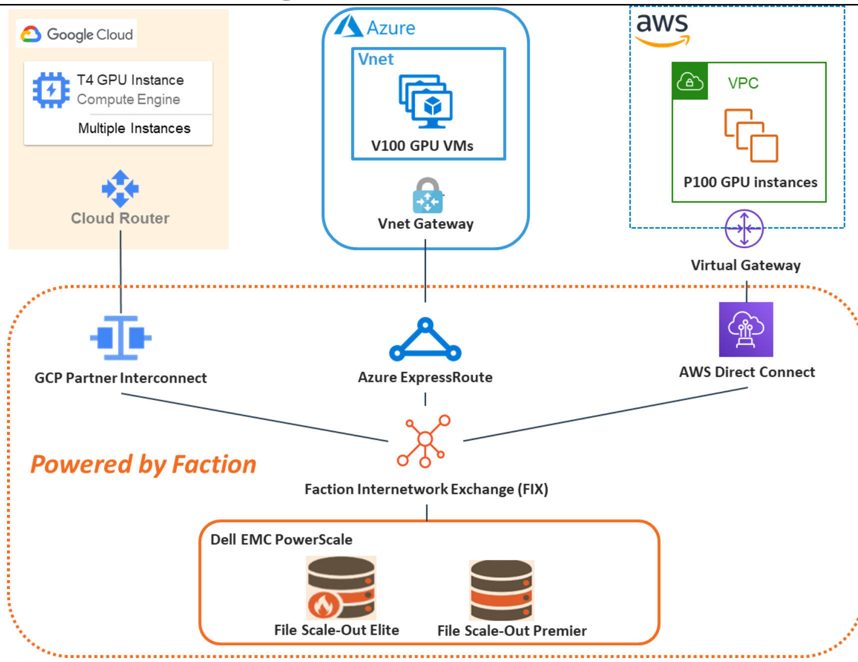


*Source: Enterprise Strategy Group*

ESG began our audit by observing the maximum number of instances that were required to process and analyze our 70 TB data set. Approximately 30 instances were initially used across the three cloud platforms. As the test continued, we saw that additional instances from Microsoft Azure, AWS, and GCP were added as more data was processed and analyzed. The initial and maximum number of instances that were used are shown in Figure 4. Peak cluster size across all three CSPs represented over one million GPU CUDA cores used.

ESG noted that the number and type of instances used varied during the test run, illustrating the importance of leveraging multiple clouds simultaneously in case instances from a single CSP became unavailable. At some points during the test, we observed that some Azure spot instances were reallocated to other customer workloads or a CSP experienced issues that affected the availability of on-demand instances. The solution then utilized compute instances available from other CSPs.

[5] While AWS and GCP refer to their compute VMs as instances, Azure uses the term VM. Throughout the rest of the paper, we will refer to any cloud compute VM as an “instance.”
[6] A T4 GPU contains 2560 CUDA Cores. A V100 GPU contains 5120 CUDA cores. A P100 GPU contains 3584 CUDA cores.
[7] While spot instance types are offered at a steep discount, there is no guarantee of availability for a given workload, as they can be terminated at any time by the CSP and revert them back to “on-demand” or reserved instances that have been previously allocated to other customers.
[8] NVIDIA Parabricks v3.0 Accelerated Burrows-Wheeler Aligner (BWA)-GATK equivalent TO GATK v4.0, the industry-standard toolkit used in genomics analysis.

**Figure 4. Scaling Computing Capacity for Processing 70 TB Genomics Data Set**

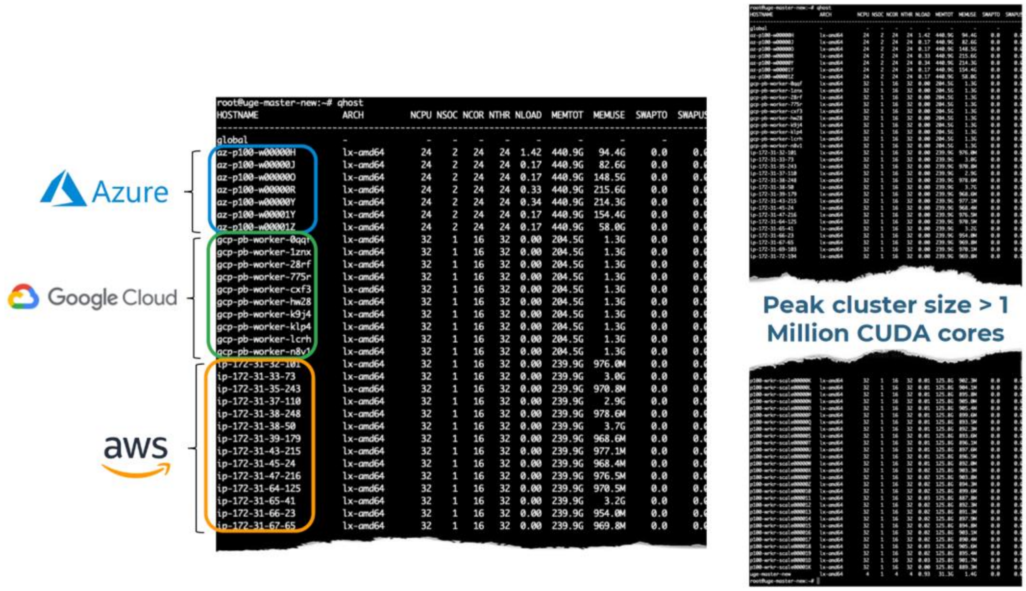


ESG then observed the maximum throughput achieved between the cluster of compute instances and PowerScale for Multi-Cloud. We measured the throughput of the individual dedicated links between the CSPs and the storage volume as well as the aggregate throughput achieved. As illustrated in Figure 5, we saw that the total maximum aggregate throughput was measured at slightly over 160 Gb/sec, while the maximum throughput achieved on Azure ExpressRoute was 140 Gb/sec.

**Figure 5. Throughput of Individual Dedicated Network Connections and Aggregate Throughput**

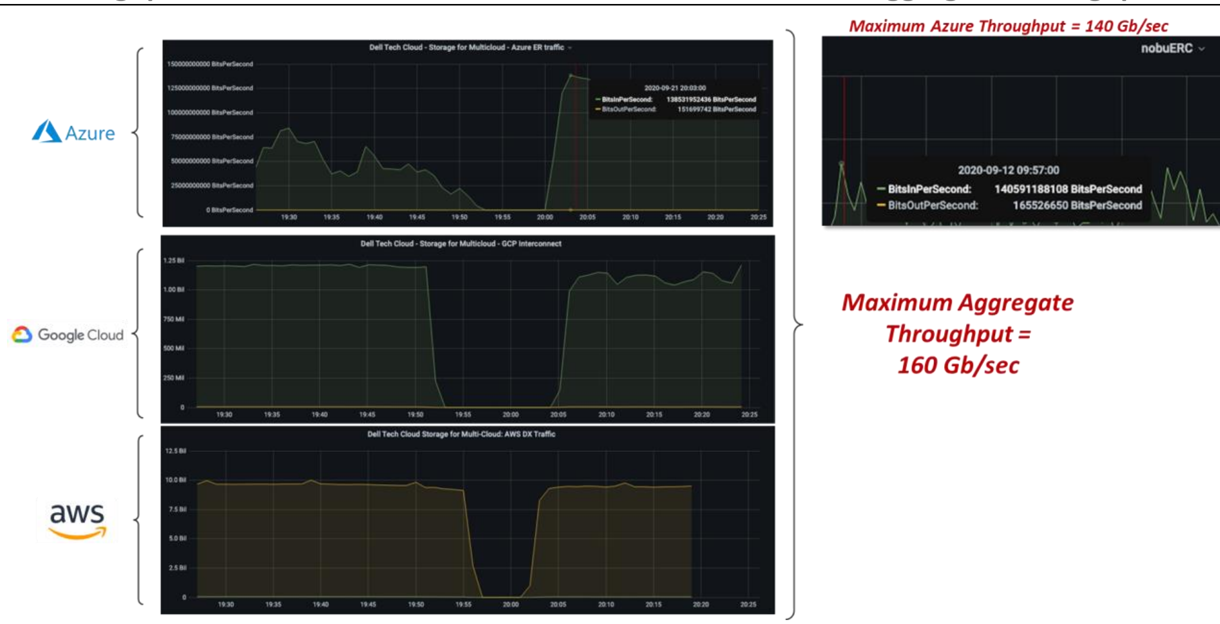


## What the Numbers Mean

- Aggregate throughput is a measure of the total amount of data flowing between the three cloud platforms simultaneously and PowerScale for Multi-Cloud.
- Higher throughput means more data can be ingested or processed in a shorter amount of time.
- The maximum of one million GPU CUDA cores simultaneously used for processing the 70TB data set illustrated how PowerScale for Multi-Cloud can maximize processing power, even though the underlying hardware infrastructure for the GPU instances differs across all CSPs.

- The aggregate throughput of 160 Gb/sec demonstrated how the solution achieves high performance for a single NFS volume containing our 70 TB data set.
- The maximum throughput of 140 Gb/sec on the Azure ExpressRoute link was observed only because the workload had utilized 70-80 Azure compute instances. If more instances were utilized, we could have achieved higher throughput.

With PowerScale for Multi-Cloud's scale-out architecture and the maximum throughput limits offered by each CSP, ESG is confident that higher levels of aggregate throughput can be achieved with workloads larger than those that were used during our testing.

ESG then examined the potential cost savings in computing capacity that organizations can achieve with PowerScale for Multi-Cloud. We first compared the estimated costs for similar genomic analysis workloads using both on-demand and spot compute instances with estimated costs for purchasing servers that supply comparable computing power over a two-year period. For this scenario, we compared the cumulative costs of GCP T4-based instances–ten on-demand instances and 20 spot instances–with the costs to purchase and operate 30 T4 GPU-based Dell servers. We assumed that spot instances operate an average of four hours per day. Both configurations represent a total of 120 GPUs. As seen in Figure 6, the costs for running our workload in the cloud is 59% less than the total costs for on-premises servers.

**Figure 6. Compute Cost Savings from Utilizing On-demand and Spot Instances versus On-premises Servers**

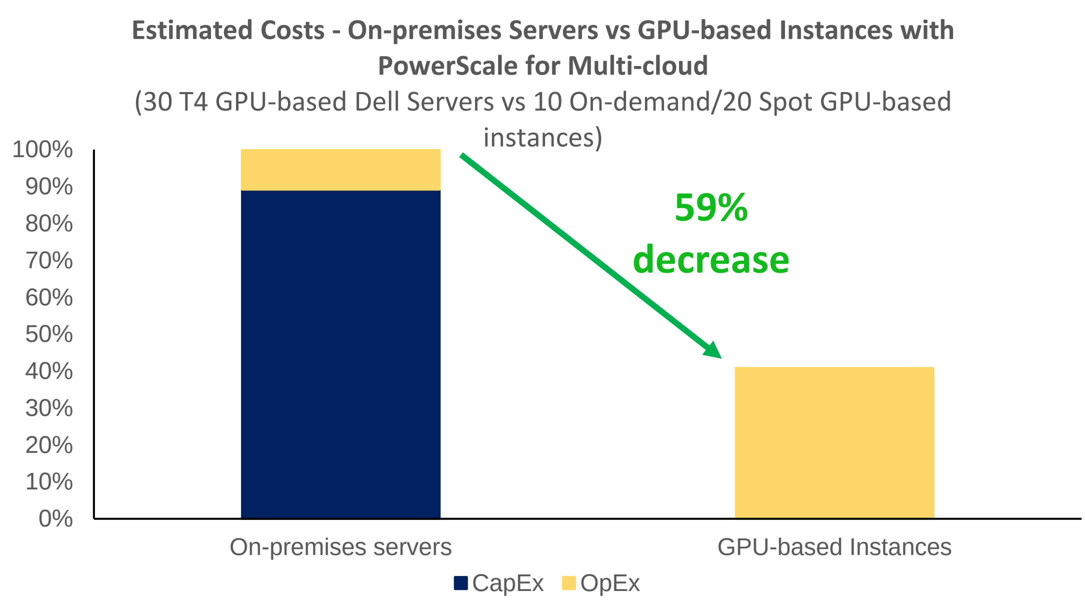


*Source: Enterprise Strategy Group*

While we only assumed the use of instances from a single cloud, ESG can see overall savings achieved when using instances across multiple clouds simultaneously leveraging PowerScale for Multi-Cloud. While the savings can vary depending on hardware specifications and instance types offered by individual CSPs, we found the cost of cloud GPUs could be less expensive than buying and operating GPU-based hardware. This savings was more pronounced the more variable and bursty the workload and the shorter the planned lifespan for the hardware were. We should note also that the use of spot instances can help organizations achieve optimal cost profiles.[9]

To evaluate the potential storage cost savings in using PowerScale for Multi-Cloud versus file storage across multiple CSPs, we reviewed a scenario for storing a 432TB file-based data set. We assumed using three cloud-based storage services: AWS Elastic File System (EFS) Standard, Microsoft Azure Files Premium, and Google Basic SSD. We compared the cumulative costs of storing multiple data set copies with the three CSPs with the costs of using PowerScale for Multi-Cloud's Faction File

[9] While spot instance types are offered at a steep discount, there is no guarantee of availability for a given workload, as they can be terminated at any time by the CSP and reverted back to "on-demand" or reserved instances that have been previously allocated to other customers.

Scale-Out Premier service tier. We did not consider the operational costs for maintaining consistent data set copies across the three CSPs in this analysis. The results (displayed in Figure 7) revealed an 89% decrease in overall monthly storage costs.

**Figure 7. Cost Savings from Utilizing PowerScale for Multi-Cloud versus Cloud-native Storage**

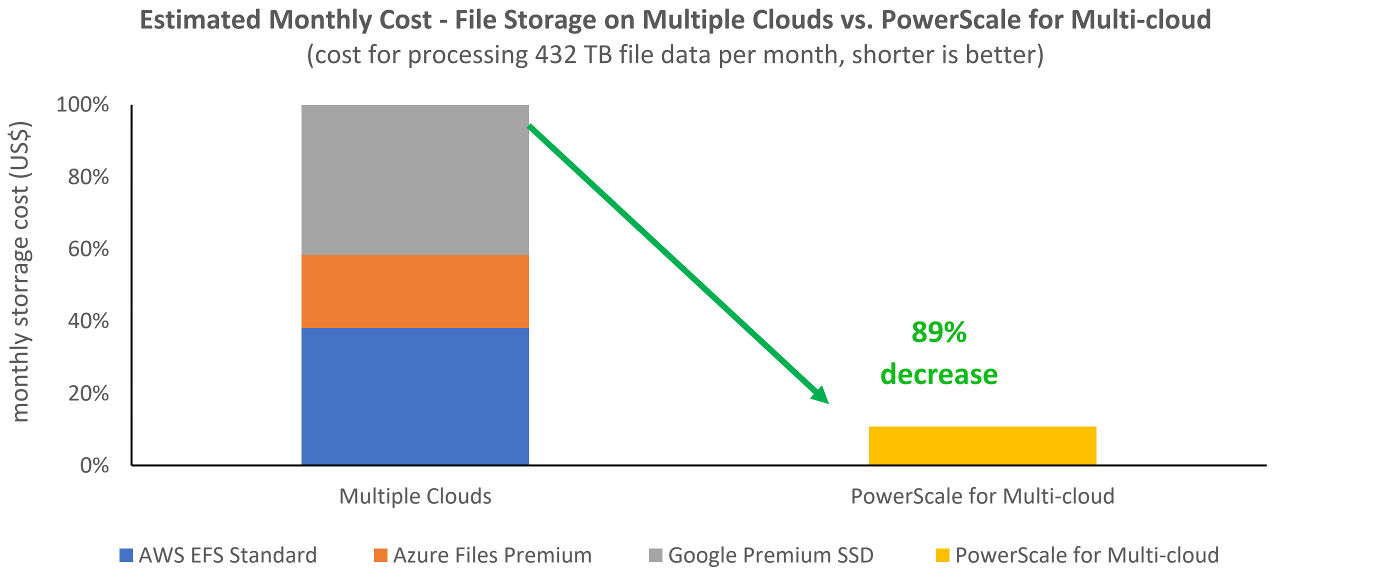


*Source: Enterprise Strategy Group*

Based on this scenario, ESG noted that storing separate data sets with each CSP can quickly become cost prohibitive when compared to storing a single master data set in the PowerScale for Multi-Cloud solution. For scenarios in which organizations continue to grow a master data set, ensuring that data copies are consistent across the CSPs would incur administration overhead and egress costs. The amount of egress costs could easily accrue when having to move data to maintain data consistency across all cloud platforms.

ESG then observed how the PowerScale for Multi-Cloud solution supported access from cloud-native services from multiple CSPs to the same data. In this case, we examined how an organization could leverage cloud-native services instead of spinning up instances to run a workload.

For this test, the test bed utilized three cloud-native analytics services—AWS Elastic MapReduce (EMR), Microsoft Azure HDInsight (powered by HDFS), and Google DataProc— to implement a data pipeline. We initiated a working session with a PySpark Jupyter Notebook,[10] a standalone Jupyter Notebook with Python hosted on Azure (see Figure 8). The Azure notebook was able to access two PowerScale for Multi-Cloud service tiers, supported by Dell Technologies Isilon H500 and Dell Technologies Isilon F810.

Each tier contained portions of a data set combining metadata from a GitHub repository (released by Tencent) and the Google Open Images repository. The complete data set consisted of 10 million images and their metadata for training image classification models.

[10] A Jupyter Notebook is an open-source web application that can be used to create and share documents containing live code, equations, visualizations, and text. It supports three programming languages - Julia, Python, and R.

## Figure 8. Test Bed for Accessing and Merging Data Files with Cloud-native Analytics Services

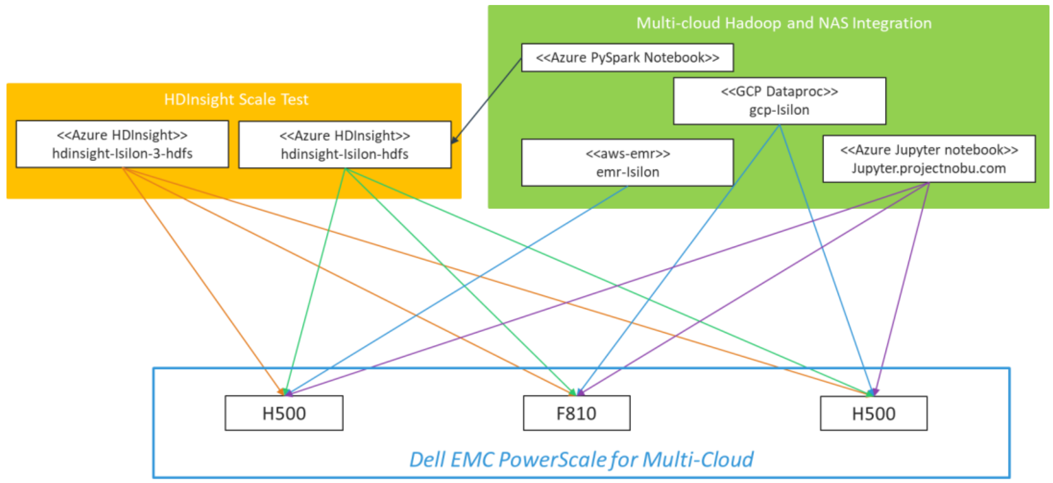


*Source: Enterprise Strategy Group*

To develop our ML model, custom scripts were developed and run via the Azure notebook. The primary steps were:

- Collect, join, and sort data from the cloud data services with Microsoft Azure HDInsight.
- Access and enrich the data using Spark on AWS EMR.
- Query image data with Spark on GCP Dataproc.
- Use selected output from the multi-cloud data pipeline to create a transfer learning image recognition model.
- Check and validate the custom model accuracy with AWS Rekognition and Azure Computer Vision.

In Figure 9, the script (1) counted images according to machine-generated labels after joining the data sets in the service, (2) attached human-readable labels to identify images within the compiled counts, and (3) stored a subset of 3,883 images based on chosen human labels stored in the service. Script execution time was measured at slightly over seven minutes. Via our Jupyter notebook on Azure, we then accessed the data subset on the service via AWS EMR to analyze and submit the data for ML model training.

## Figure 9. Joining Data Sets, Sorting Data, and Outputting Subset of Original Data Set

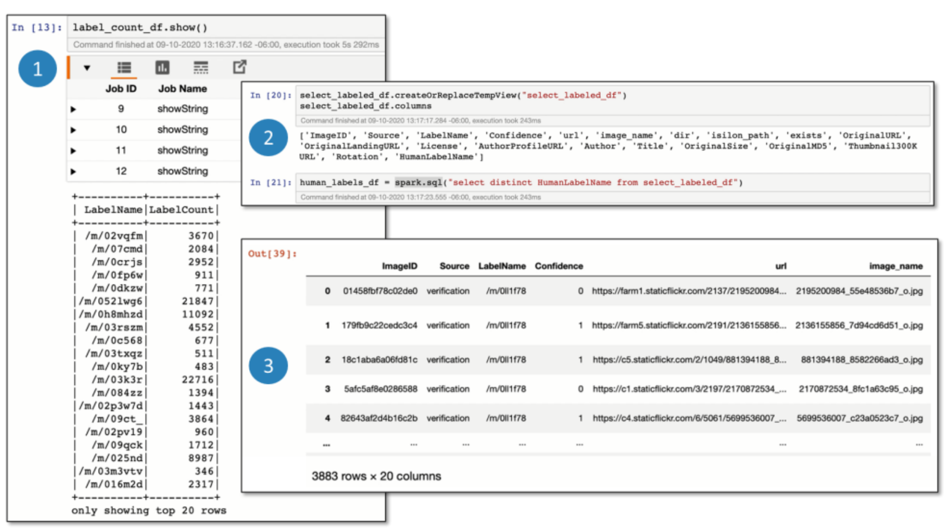

Based on the script execution, ESG observed that the data analysis pipeline was executed without the need to store multiple copies of the original image set and the subset for ML model training in all clouds. The setup enabled the joining, sorting, and analysis of the image data set using common storage. No costly data transfers between the CSPs occurred during this test. We could see how organizations can synchronize access automatically to large training data sets stored within multiple clouds. As in the previous test, ESG noted that costs for moving or storing copies of the training data would not be incurred.

## Why This Matters

Organizations continue to adopt both public cloud infrastructure and public cloud-native services to decrease overall on-premises infrastructure costs while reducing the time and effort to scale IT resources and respond to business needs. However, CSPs do not empower their customers to use services from multiple cloud platforms simultaneously on any single workload, leading to vendor lock-in. For those that want to take advantage of multiple cloud platforms to complete a workload with large data sets, organizations would need to store multiple copies of those data sets with each CSP. However, this leads to increased operational complexity and added costs.

ESG's audit of the PowerScale for Multi-Cloud solution revealed that organizations can employ both compute instances and cloud-native services from multiple CSPs simultaneously without the burden of maintaining multiple copies of file data. This is especially crucial for AI/ML applications that consume massive amounts of file data, as the size of the data set amplifies the challenges associated with data management, specifically when maintaining data consistency.

After testing PowerScale for Multi-Cloud to complete a genomics processing workload with a 70TB data set, ESG observed that the solution achieved maximum aggregate throughput of 160 Gb/sec leveraging over one million GPU cores across all three cloud platforms—Microsoft Azure, AWS, and GCP (with the given number of instances purchased). This performance was greater than the performance that could be achieved using any single CSP. In another test, ESG also observed that processing time for large data sets can be reduced by leveraging the interoperability of PowerScale for Multi-Cloud with multiple cloud-based analytics services. In the use case of creating and training an ML model, we again accessed and manipulated large data sets without transferring data between clouds, illustrating the solution's ability to harness the capabilities of multiple cloud platforms simultaneously.

## The Bigger Truth

The public cloud is expanding from serving as simply an extension of the corporate data center to a flexible IT asset for running specialized workloads such as AI/ML capabilities. Instead of building out AI/ML capabilities in house, organizations can leverage IT resources and cloud-native services on demand. However, organizations cannot leverage and scale cloud compute resources from multiple CSPs concurrently without adding operational complexity and slowing their internal processes due to the time required to move and copy data between cloud platforms, ultimately making the endeavor cost prohibitive.

The PowerScale for Multi-Cloud solution combines Dell EMC PowerScale with Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) to help organizations run data-intensive, high I/O throughput, file-based workloads, especially high-performance computing (HPC) applications such as AI/ML. To maximize workload performance and availability while minimizing overall cloud expenses, the solution enables the use of multiple compute resources or cloud-native services across multiple CSPs simultaneously while consolidating data into a single repository. Distributing a workload across multiple cloud platforms becomes easier as that single repository eliminates complex data management requirements for maintaining data consistency.

ESG validated that the PowerScale for Multi-Cloud can simplify the simultaneous use of multiple cloud compute resources and cloud-native AI/ML services without the need to maintain multiple copies of master data within individual cloud platforms. We used a genomics processing use case—a workload that requires large data files to obtain appropriate insights—to demonstrate the use of multiple compute instances (GPU-based instances) across cloud platforms to generate the desired output. We observed that the ability to shift compute capacity between cloud platforms helped in running a more cost-effective workloads while compensating for varying availability of GPU instances. When we observed another test for training an ML model to recognize images, we found that we could allocate different parts of the data processing pipeline to cloud-native services of AWS, Microsoft Azure, and GCP. In both tests, overall processing time was minimized.

Multi-Cloud Data Services for Dell EMC PowerScale, enabled by Faction, allows organizations to run workloads concurrently on multiple cloud platforms, achieving high performance while controlling overall cloud spend. We recommend that you take a close look at this solution should you need to process data-intensive HPC workloads with large data sets.



www.esg-global.com contact@esg-global.com 508.482.0188