

# Collaborative Video Editing with Dell Technologies Cloud PowerScale for Google Cloud

## Abstract

This paper describes using Dell Technologies™ Cloud PowerScale™ for Google Cloud in a collaborative video edit workflow using virtualized video-edit workstations.

June 2020

## Revisions

Date	Description
June 2020	Initial release

## Acknowledgments

Author: Gregory Shiff

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [6/12/2020] [Technical White Paper] [H18114]

# Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents .....	3
Executive summary.....	4
1 Introduction.....	5
2 PowerScale for Google Cloud overview.....	6
2.1 Testing goals .....	6
2.2 OneFS configuration.....	6
2.3 Google Cloud configuration .....	7
2.4 Windows configuration .....	8
2.5 Testing methodology .....	8
3 Testing results .....	10
3.1 Single-client Isilon H500 test .....	10
3.2 Two clients with single Isilon H500 node.....	12
3.3 Four clients with single Isilon H500 node .....	12
3.4 Four clients connected to four Isilon H500 nodes (one client per node).....	12
3.5 Single-client Isilon F800 test.....	13
3.6 Two clients connected to single Isilon F800 node.....	13
3.7 Eight clients connected to four Isilon F800 nodes.....	13
4 Conclusions .....	14
A Technical support and resources .....	15
A.1 Related resources .....	15

## Executive summary

This paper describes using Dell Technologies™ Cloud OneFS™ for Google Cloud in a collaborative video edit workflow using virtualized video edit workstations. It includes configuration details for the solution components and the testing methodology. The paper also provides testing results that demonstrate how PowerScale for Google Cloud is an ideal storage platform for collaborative video editing in the cloud, including 4K uncompressed workflows.

# 1 Introduction

A core strength of Dell EMC™ PowerScale™ storage in media post production is its ability to provide performance to multiple users simultaneously as a part of a collaborative workflow. The testing that is performed for this paper confirms that Dell Technologies Cloud PowerScale for Google Cloud maintains this same core functionality. It provides the same level of performance for virtual workstations that are hosted on Google Cloud compared to local workstations that are connected to on-premises PowerScale storage.

The solution that is tested in this document consists of PowerScale for Google Cloud, virtual workstations that are hosted on Google Cloud, and Teradici® PCoIP® remote desktop software. It also consists of various video-editing applications such as Adobe® Premiere® and Blackmagic Design DaVinci Resolve. With ping times of less than 15 ms between the local thin client and virtual workstation, Teradici PCoIP software makes these systems indistinguishable from local workstations at 4K resolution, even over the public WAN.

This architecture enables collaborative workflows for geographically disbursed teams or teams without local infrastructure that need to use the cloud for collaboration. It also allows teams to take advantage of the scale, performance, and reliability of the Dell EMC PowerScale OneFS™ platform. PowerScale for Google Cloud uses PowerScale or Dell EMC Isilon™ hardware for the storage back end. The underlying storage clusters are not shared between customers.

The following list summarizes the key benefits of Dell Technologies Cloud PowerScale for Google Cloud.

- Performance is assured. Each customer uses dedicated PowerScale or Isilon hardware.
- Data tenancy is defined. Storage hardware is not shared between customers.
- The solution is not impacted by cloud-data-center maintenance. PowerScale or Isilon hardware is located in dedicated data centers.
- Front- and back-end storage network traffic is on segregated network interfaces.
- Sizing can use existing PowerScale and Isilon performance guidelines and expertise.

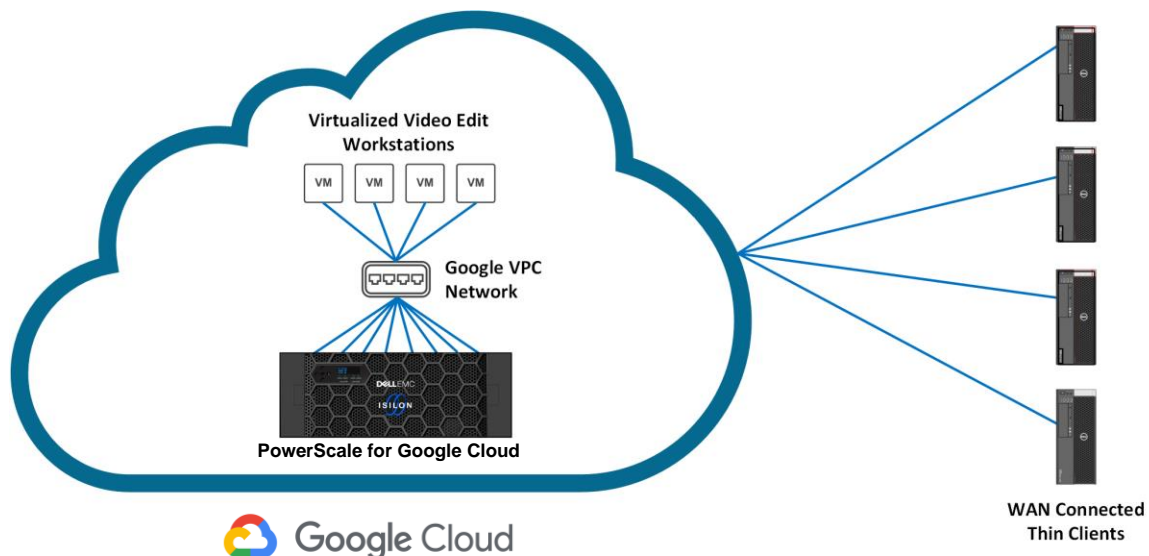


Figure 1 Collaborative video editing in the cloud with PowerScale for Google Cloud

## 2 PowerScale for Google Cloud overview

PowerScale for Google Cloud is a native offering within the Google Cloud intranet, is based on PowerScale technology, and is managed by Dell Enterprise Services. It is powered by the PowerScale family which includes Dell EMC Isilon™ and PowerScale nodes as well as Dell EMC PowerScale OneFS™ which runs across these systems. PowerScale for Google Cloud provides terabit networking and sub-millisecond latency between the PowerScale hardware and Google Cloud compute. The service is billed through Google Cloud like other Google services, and it can count toward any committed Google Cloud spend. Since PowerScale for Google Cloud is native within Google Cloud, there are no egress charges associated with moving data between this storage platform and other Google Cloud services.

Each customer is provided with dedicated PowerScale hardware. Management of the PowerScale cluster is accomplished through the Google Cloud Console for a native cloud experience. Since each PowerScale cluster is dedicated to a particular customer, performance is predictable and data tenancy is assured. Customers can use native OneFS data services for moving data between their on-premises and PowerScale for Google Cloud instances.

Media applications, particularly video editing, require high-throughput and low-latency file storage, making OneFS a good match to these workflows. By contrast, most cloud storage is based on object storage. Object storage has numerous benefits in hyperscale environments. This is a deep topic beyond the scope of this paper. In summary, while it provides fault tolerance and extreme scalability, object storage does not provide the performance or POSIX-style file-system interface that is required by real-time media applications.

PowerScale for Google Cloud compliments native Google Cloud object storage in several ways. It uses dedicated hardware for each customer. The potential for performance degradation that is caused by multiple users of shared infrastructure is eliminated. Since PowerScale for Google Cloud uses PowerScale or Isilon hardware, the seamless multi-petabyte scalability of the platform extends to this offering. Standard OneFS features such as multi-protocol file access and SMB3 multichannel create a mature file offering in Google Cloud.

### 2.1 Testing goals

The primary goal of the testing in this paper is to confirm the performance of PowerScale for Google Cloud when attached to virtualized workstations that are running media applications. Virtualizing media applications is a hot topic for many content creators who increasingly require flexible infrastructure, whether local or cloud based.

Video-edit workflows that use local PowerScale storage are well established. PowerScale for Google Cloud should provide the same level of responsiveness and performance; it uses the same hardware and is connected by low-latency, high-throughput links. Even so, given the additional layers of complexity that are involved in all cloud topologies, additional testing to validate this assumption is required.

### 2.2 OneFS configuration

OneFS does not require extensive configuration to deliver optimal performance. However, OneFS provides multiple file-layout and access strategies. Typically, the default concurrency strategy is sufficient. However, to maximize playback of high-throughput media files (particularly on clusters with hard disks), **streaming mode** is recommended. For playback of uncompressed image sequences, use of **file-name prefetch** should be enabled on the directories containing only those image sequences. As with all tunings, testing is key.

While file-layout and access strategies are beyond the scope of this paper, both of these topics are discussed in the document [Isilon Filename Based Prefetch](#).

For the testing performed in this paper, compressed media files are placed in directories with streaming mode enabled. Image sequences are located in directories with the file-name prefetch access strategy assigned.

## 2.3 Google Cloud configuration

Each workstation is configured with 32 virtual cores and two NVIDIA® GPUs. At the time of publication, the GPU-enabled compute available at the Google Cloud us-east4 data center includes 2.0 GHz and 2.2 GHz processors.

All virtual workstations run Microsoft® Windows Server® 2016 Datacenter (one of the default operating system offerings of Google Cloud) and run the latest NVIDIA drivers available. Access to virtual workstations is accomplished using Teradici Cloud Access Software using the Teradici PCoIP protocol. Connection to PowerScale for Google Cloud is through the SMB 3.1.1 protocol.

Virtual workstations are on a single subnet within Google Cloud with access to the PowerScale for Google Cloud subnet. See Figure 2 for Google Cloud workstation configuration information.

The screenshot shows the configuration interface for a Google Cloud workstation named 'workstation-2'. The configuration is as follows:

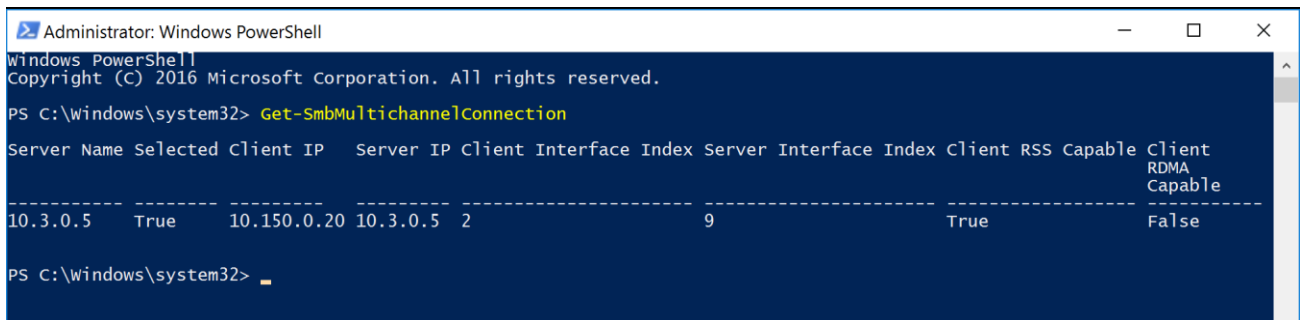
- Region:** us-east4 (Northern Virginia)
- Zone:** us-east4-c
- Machine configuration:**
  - Machine family:** General-purpose
  - Series:** N1
  - Machine type:** Custom
  - Cores:** 32 vCPU (range 1 - 96)
  - Memory:** 64 GB (range 29 - 208)
  - Extend memory
- CPU platform:** Intel Skylake or later
- GPU type:** NVIDIA Tesla P4
- Number of GPUs:** 2
- Enable Virtual Workstation (NVIDIA GRID)

Figure 2 Google Cloud workstation configuration

## 2.4 Windows configuration

Microsoft® Windows Server® 2016 Datacenter is one of the default operating system offerings available in Google Cloud. While an install of Microsoft Windows® 10 Pro may be the best operating system for virtualized media workstations, Windows Server 2016 works well for the performance testing in this paper. No special tuning or configuration is required beyond confirmation that SMB 3 multichannel is active and being used.

A full discussion of SMB3 multichannel is beyond the scope of this paper. However, SMB 3 multichannel does an effective job of multithreading networked-storage operations and distributing those threads among cores on the workstation. See Figure 3 for SMB 3 multichannel configuration in Windows PowerShell®.



```

Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) 2016 Microsoft Corporation. All rights reserved.

PS C:\windows\system32> Get-SmbMultichannelConnection

Server Name Selected Client IP      Server IP Client Interface Index Server Interface Index Client RSS Capable Client
----- RDMA Capable
-----
10.3.0.5    True    10.150.0.20 10.3.0.5 2          9          True          False
PS C:\windows\system32>

```

Figure 3 Windows SMB 3 multichannel configuration

## 2.5 Testing methodology

The testing configuration in this paper includes multiple GPU-enabled virtual workstations that are connected to a PowerScale for Google Cloud instance. Two different PowerScale for Google Cloud cluster tiers are tested: Tier I is based on Isilon F800 all-flash hardware, and Tier II is based on Isilon H500 hybrid hardware. The virtual workstations and PowerScale for Google Cloud are located in the same physical Google Cloud region, which is **us-east4** in this case. Latency between the virtual workstations and storage is sub-millisecond, as expected from PowerScale for Google Cloud.

Teradici PCoIP Cloud Access Software is used to log in to the virtual workstations and view the displays at 4K resolution. Cloud Access Software uses the PCoIP protocol, a technology that delivers a secure, high-definition, and responsive remote experience. It uses a proprietary display protocol to compress, encrypt, and transmit only pixels to the client so no data ever leaves the cloud. PCoIP allows for full-frame-rate video across multiple displays, offers accelerated interactivity using pen tablets and displays, and retains color accuracy.

PCoIP features precise, build-to-lossless capability, even at 4K/UHD resolution. This means that the entire image is reproduced in pixel-perfect clarity. Teradici can accomplish color accuracy by using a fully reversible color-space conversion. This conversion results in remote reproduction of image colors that are precisely as intended by creative applications. Administrators have control over PCoIP features including bandwidth optimizations, image quality, and protocol modes such as [PCoIP Ultra](#). PCoIP Ultra uses the hardware-acceleration features of modern servers including multi-core AVX2 instructions and NVIDIA NVENC encoding. These features help to deliver highly interactive, creative software applications as if run by a local workstation.

The Google Cloud data center that runs the virtual workstations is several hundred miles away with ping times of about 11 ms to 12 ms from the local system that is running the Teradici client software. With that amount of latency, the virtual workstations are fully responsive with no noticeable lag.



Four data sets are created for playback:

- Compressed high-definition (HD) clips at ~281 Mbps (27.3 MB/s)
- Compressed ultra-high-definition (UHD) clips at ~829 Mbps (103 MB/s)
- Uncompressed 10-bit color UHD image sequence at ~6,144 Mbps (768 MB/s)
- Uncompressed 12-bit color UHD image sequence at ~9,600 Mbps (1,200 MB/s)

Copies of the dataset are created so that testing of multiple clients that are accessing the same media and multiple clients accessing different media can both be assessed.

For the compressed media, multi-cam objects with various numbers of simultaneous streams are created and added to unique timelines. These timelines are played back and monitored for dropped frames on multiple client systems. Before each playback test, the L1, L2, and L3 cache on the PowerScale for Google Cloud cluster is flushed. This action ensures a worst-case scenario where all data is pulled directly from the underlying storage. Throughput from the storage and network throughput on the client are monitored to verify that the video streams are coming from the storage and not the cache.

A similar method is used for the uncompressed DPX sequences. However, these sequences are added to unique timelines for playback of one and two layers.

Besides watching the network throughput on the client, CPU and GPU utilization are monitored. In general, when CPU and GPU use begin to spike on a client, frames drop before the performance of the PowerScale for Google Cloud instance is saturated. This result is expected because it is challenging for a single client or small groups of clients to saturate OneFS.

OneFS is a clustered, node-based system. Each node contributes to the total cluster performance and capacity. Client connections are distributed between nodes using Dell EMC PowerScale SmartConnect or manual IP address assignments. For various client-connection situations to be simulated, these tests are run with four clients connecting either to the same node or spread out between the four nodes in this small cluster.

The purpose of the tests in this paper is to show that PowerScale for Google Cloud connected to Google-Cloud-based compute performs at the same level as on-premises PowerScale hardware connected to on-premises compute. It is also to demonstrate the feasibility of doing editorial work on cloud workstations and PowerScale for Google Cloud. The purpose of this document is **not** to show the maximum performance of the tiered offerings and the results of this testing should not be taken as sizing guidance. As anyone who works with storage knows, sizing exercises require detailed inputs about the total environment and workflows involved.

## 3 Testing results

This section details the results of the solution that is tested in this paper. The configuration details are listed as follows.

### Google Cloud workstation configuration:

- 32-core, 2.2 GHz and 2.0 GHz Intel® processors
- Dual NVIDIA® T4 GPUs
- Windows Server 2016 Datacenter

### PowerScale for Google Cloud configurations:

- Tier II - Based on Isilon H500
  - SATA drives in H500
- Tier I - Based on Isilon F800
  - SSD drives in F800
- PowerScale OneFS - SMB 3 protocol

### Test media:

- Compressed HD: ProRes 422, 30 frames per second, 1920x1080 resolution
  - ~281 Mbps (27.3 MB/s)
- Compressed UHD: ProRes 422HQ, 30 frames per second, 3840x2160 resolution
  - ~829 Mbps (103 MB/s)
- Uncompressed DPX 8-bit: 24 frames per second, 3840x2160 resolution, 10-bit color
  - ~6,144 Mbps (768 MB/s)
- Uncompressed DPX 12-bit: 24 frames per second, 3840x2160 resolution, 12-bit color
  - ~9,600 Mbps (1,200 MB/s)

### 3.1 Single-client Isilon H500 test

- Compressed HD: 15 streams
- Compressed UHD: 7 streams
- 10-bit DPX: 2 image sequences
- 12-bit DPX: 1 image sequence

The results of single-client testing do not push the upper limits of PowerScale for Google Cloud. Depending on the media type, two factors limit the performance. For the compressed media, CPU usage of the virtual workstation starts to spike up to about 85% and causes dropped frames. For the uncompressed media, a throughput of ~16 Gbps is reached.

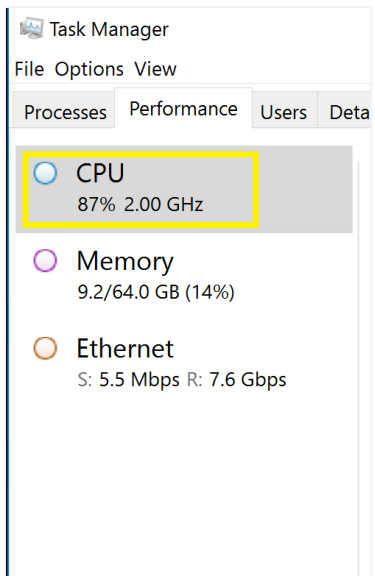


Figure 4 Compressed playback and high CPU

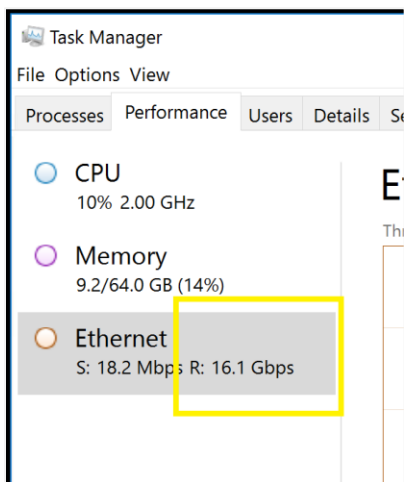


Figure 5 Uncompressed playback and high throughput

### 3.2 Two clients with single Isilon H500 node

- Compressed HD: 15 streams per client, 30 streams total
- Compressed UHD: 7 streams per client, 14 streams total
- 10-bit UHD DPX: 3 image sequences total, 2 from first client, 1 from second client
- 12-bit UHD DPX: 2 images sequences, 1 image sequence per client

In this test, two clients are attached to a single Isilon node. The same bottlenecks for compressed HD and UHD playback exist like the previous test, with GPU and CPU being the limiting factor. For the uncompressed playback, the upper limits of the single Isilon node are beginning to be reached, noting that the hardware for this solution is a four-node Isilon H500 cluster.

### 3.3 Four clients with single Isilon H500 node

- Compressed HD: 12 streams per client, 48 streams total
- Compressed UHD: 4 to 5 streams per client, 19 streams total
- 10-bit UHD DPX: 3 image sequences total
- 12-bit UHD DPX: 2 image sequences total

In this test, four workstations are connected to a single Isilon node. For the HD compressed content, the total per-workstation streams drop slightly as the Isilon node begins to be taxed. For the UHD compressed media, the limits of the single Isilon H500 node also begin to be a factor. For the uncompressed playback, the single Isilon H500 node is the limiting factor.

### 3.4 Four clients connected to four Isilon H500 nodes (one client per node)

- Compressed HD: 16 streams per client, 64 streams total
- Compressed UHD: 7 per client, 28 streams total
- 8-bit DPX: 1 per client, 4 image sequences total
- 12-bit DPX: 1 per client, 4 images sequences total

OneFS is a clustered system, and client connections are usually distributed between nodes. For this behavior to be simulated, each client is connected to different Isilon node in the four-node cluster. Once again, HD playback is only limited by the workstation, which is the same with UHD playback. For DPX playback, the limits of the underlying sixty 7200 RPM hard drives are being pushed.

Figure 6 shows the results of this test. The SMB column shows traffic out to various clients, and the DiskOut column shows non-cached data streaming from drives.

Node	CPU	SMB	FTP	HTTP	NFS	HDFS	Total	NetIn	NetOut	DiskIn	DiskOut
All	20.5%	2.7G	0.0	0.0	0.0	0.0	2.7G	2.9M	2.8G	3.2M	3.1G
1	31.9%	1.1G	0.0	0.0	0.0	0.0	1.1G	1.7M	1.2G	1.5M	784.8M
2	15.0%	500.2M	0.0	0.0	0.0	0.0	500.2M	376.9k	527.7M	42.6k	729.9M
3	16.0%	518.1M	0.0	0.0	0.0	0.0	518.1M	413.7k	546.7M	1.6M	796.8M
4	19.2%	523.8M	0.0	0.0	0.0	0.0	523.8M	471.2k	552.9M	47.5k	754.3M

Figure 6 Mixed-edit workload on four-node H500 based PowerScale for Google Cloud cluster

### 3.5 Single-client Isilon F800 test

- 10-bit UHD DPX: 2 image sequences
- 12-bit UHD DPX: 2 image sequences

The F800 tests that focus on uncompressed media playback show that two 4K DPX sequences can play back from a single node. The Windows VM begins to bottleneck at >2 image sequences while additional bandwidth from the F800 node is possible.

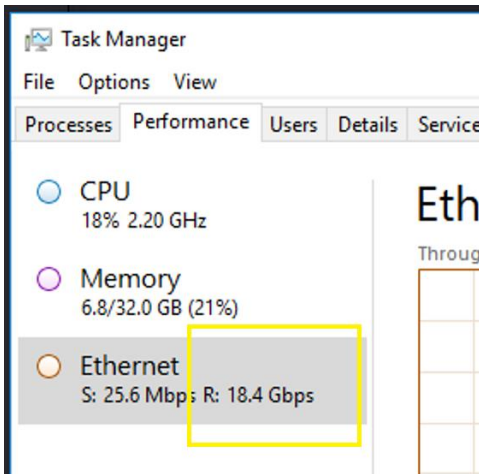


Figure 7 Two uncompressed image sequences playing from a single virtual workstation on PowerScale for Google Cloud

### 3.6 Two clients connected to single Isilon F800 node

- 10-bit UHD DPX: 4 image sequences total, 2 from first client, 2 from second client
- 10-bit UHD DPX: 4 image sequences total, 2 from first client, 2 from second client

This PowerScale for Google Cloud cluster is quite large, consisting of 56 nodes of F800 hardware. Therefore, it is possible to saturating a single node. The clustered nature of OneFS allows for each node to receive storage resources of the entire cluster.

### 3.7 Eight clients connected to four Isilon F800 nodes

- 8-bit DPX: 16x total image sequences, 2 per client
- 12-bit DPX: 10 total image sequences, mix of 1 or 2 image sequences per client

This large F800 cluster is able to saturate the performance of single nodes. Two of the four nodes are saturated while two of the nodes delivered close to saturated performance.

Node	CPU	SMB	FTP	HTTP	NFS	HDFS	Total	NetIn	NetOut	DiskIn	DiskOut
All	4.3%	14.3G	0.0	0.0	0.0	0.0	14.3G	17.7M	14.8G	1.2M	15.2G
1	51.7%	4.4G	0.0	0.0	0.0	0.0	4.4G	5.3M	4.5G	67.2k	296.8M
2	50.4%	4.4G	0.0	0.0	0.0	0.0	4.4G	6.1M	4.9G	0.0	288.7M
3	23.1%	1.9G	0.0	0.0	0.0	0.0	1.9G	2.3M	2.0G	0.0	254.5M
4	39.6%	3.6G	0.0	0.0	0.0	0.0	3.6G	4.0M	3.4G	14.7k	250.0M

Figure 8 14.3 GB/s of uncompressed video playback from PowerScale for Google Cloud

## 4 Conclusions

The testing in this paper confirms that the performance of PowerScale for Google Cloud is minimally impacted by being part of a cloud architecture. It validates PowerScale for Google Cloud as a storage platform for virtualized video-edit workstations within Google Cloud.

The compressed media that is used in these tests has high bit rate by design. The total number of virtual workstations is limited, and using higher-bit-rate material is required to push performance on the storage layer. Therefore, the available processor speeds is a limiting factor. This factor would be less of an issue with lighter-weight, lower-bit-rate, compressed material and the availability of higher-speed, GPU-enabled compute instances at other Google Cloud data centers.

For uncompressed playback, it is recommended to use flash-based OneFS systems. PowerScale for Google Cloud with an Isilon all-flash F800 system as the underlying storage layer can play many simultaneous uncompressed image sequences. Even so, the SATA-drive-based Isilon H500 system is also able to handle 4K uncompressed image sequence playback as well.

The default layout and access strategy for OneFS is **concurrency**. This use is acceptable for most workflows. However, to push the performance of the Isilon H500 cluster that underpins this PowerScale for Google Cloud instance, **streaming** mode and **file-name prefetch** are employed. Most OneFS clusters are doing more than playing back large video files and have far more than the four clients used in this test. For those clusters, using the default concurrency mode is likely the best choice. However, selective use of streaming and file-name prefetch on a per-directory basis can be considered.

With ping times of 11 ms to 12 ms between the local client and virtual workstation, Teradici Cloud Access Software can drive a 27-inch 4K display with no noticeable lag or image degradation. Also, audio and video play back in sync. Workstations at a data center that is ~3,000 miles away are created where ping times are considerably higher (~75 ms). In this case, latency is noticeable, but the workstations are usable.

While this testing is focused on the PowerScale for Google Cloud offering, the lessons demonstrated also apply to virtualized edit seats in on-premises environments. Some of these lessons include the following:

- High clock speeds are preferable to high core counts.
- PowerScale storage works as well for virtualized creative seats as it does for non-virtual workstations.
- Keeping ping times from the virtual instance to the local system below 15 ms is critical.
- SMB 3 multichannel makes a big difference in Windows-based workflows.

## A Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

### A.1 Related resources

- [PowerScale and Isilon technical documents and videos](#)
- [Dell EMC Isilon H500](#)
- [Dell EMC Isilon OneFS Best Practices](#)
- [Dell EMC Isilon OneFS Technical Overview](#)
- [Teradici Cloud Access Software](#)
- [Isilon Filename Based Prefetch](#)
- Dell Technologies Cloud PowerScale for Google Cloud on [delltechnologies.com](https://delltechnologies.com).