

Technical Validation

# Stellus Data Platform

## Accelerating Consistent Data Access Performance

By Tony Palmer, Senior Validation Analyst; and Alex Arcilla, Validation Analyst  
February 2020

This ESG Technical Validation was commissioned by Stellus and is distributed under license from ESG.

## Contents

Introduction .....	3
Background .....	3
Stellus Data Platform .....	4
Flexible Integration with SDS .....	5
ESG Technical Validation .....	5
Scale-through Performance .....	5
ESG Testing.....	5
Scalability, Availability, and Management .....	9
ESG Testing.....	9
The Bigger Truth.....	10

### ESG Technical Validations

The goal of ESG Technical Validations is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Technical Validations are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objectives are to explore some of the more valuable features and functions of IT solutions, show how they can be used to solve real customer problems, and identify any areas needing improvement. The ESG Validation Team’s expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.

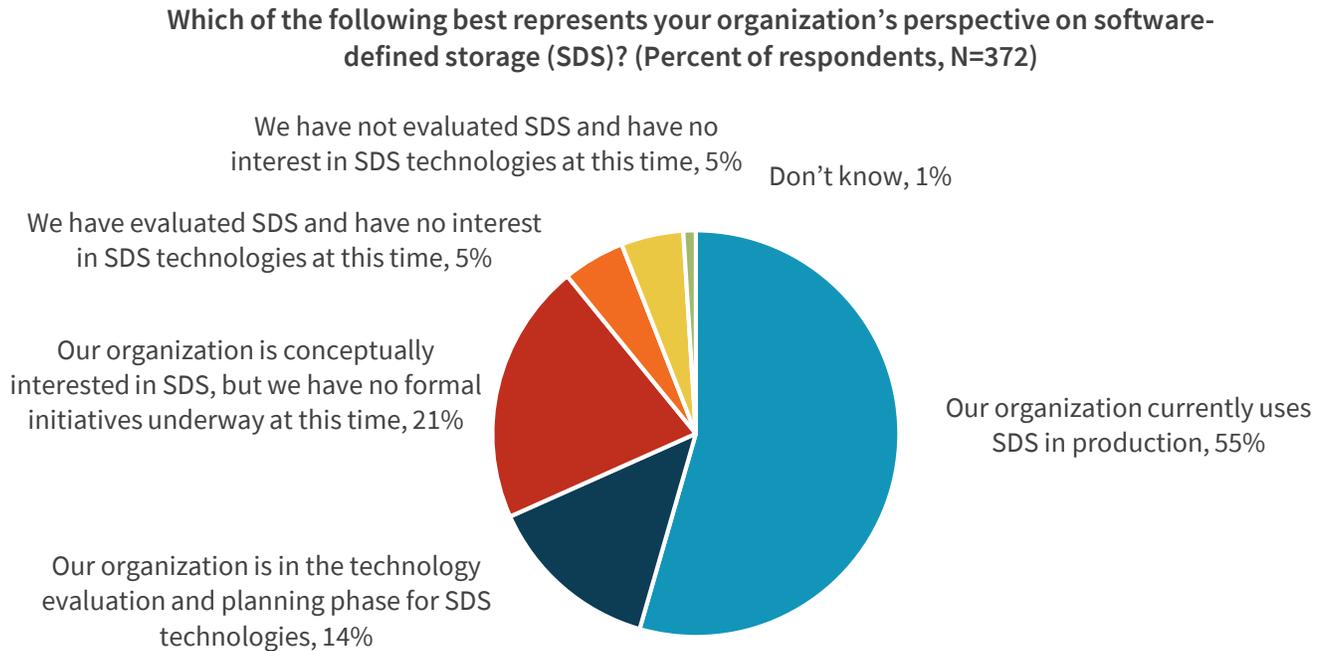
## Introduction

This ESG Technical Validation documents hands-on testing of the Stellus Data Platform (SDP). We evaluated how the SDP achieves high and consistent read and write performance via its “scale-through” architecture and looked at system availability.

## Background

ESG recently uncovered that more than four in ten organizations have at least 1 PB of primary storage capacity.<sup>1</sup> It is not surprising, since 56% of these organizations anticipate their on-premises data growth to be at least 21%, with one in four expecting greater than 50% year over year growth over the next three years. With organizations facing this level of growth, software-defined storage (SDS) technologies are being considered as a viable way to grow storage capacity in a simple and flexible manner. In fact, 55% of surveyed organizations are already using SDS technology in production and another 14% are formally evaluating it (see Figure 1). Forty-six percent of those using or evaluating SDS technology already achieve or expect to achieve increased performance.

**Figure 1. Software-defined Storage Usage**



Source: Enterprise Strategy Group

As larger amounts of both structured and unstructured data are generated and collected, organizations have turned to SDS to build out their storage infrastructure more efficiently; one primary goal is to scale both processing power and capacity independently of each other to meet their current needs without overprovisioning either element. Now that organizations are employing data-intensive applications such as those used for life sciences, post-production editing, the internet of things (IoT), artificial intelligence (AI), and machine learning, which utilize unstructured data, organizations seek out SDS solutions that will provide fast and consistent storage performance—for both reads and writes—as their storage grows and

<sup>1</sup> Source: ESG Master Survey Results, [2019 Data Storage Trends](#), November 2019. All ESG research references and charts in this technical validation were taken from this master survey results set, unless otherwise noted.

support the large amount of application processing required so they can process data and extract value without unnecessary delay.

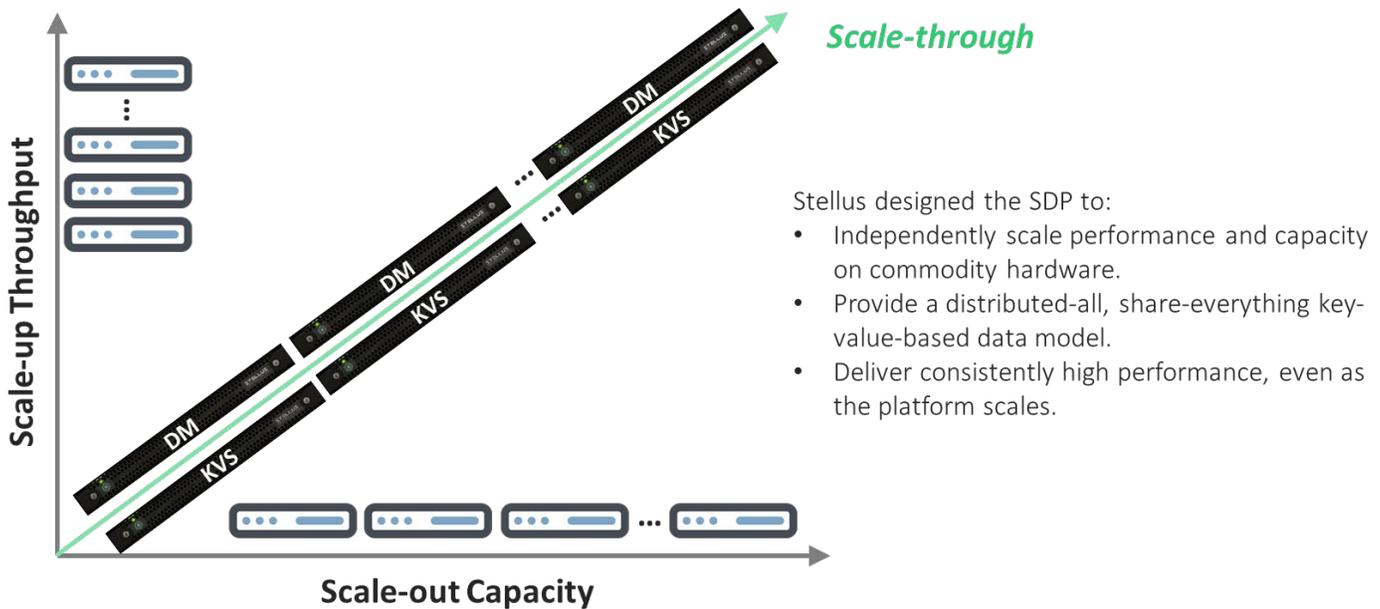
### Stellus Data Platform

Stellus designed the Stellus Data Platform to support and manage the exponential growth of unstructured data that organizations need to leverage data-intensive applications for competitive advantage. The Stellus Data Platform can deliver fast, consistent, and scalable read and write performance. Depending on the specific application use case, the SDP can deliver performance of 20-80+ GB/s read throughput and 19-70+ GB/s write throughput.

Leveraging SDS, organizations can scale processing power and capacity independently via its “scale-through” architecture, consisting of Data Managers (DMs) and Key-Value Stores (KVSs) (see Figure 2). With both scale-up and scale-out architectures of traditional storage solutions, organizations either need to scale up processing power by adding smaller storage systems at the expense of adding unwanted capacity or scale out by adding larger systems to gain additional capacity at the expense of purchasing unwanted processing power. Adding systems to scale processing power or capacity did not translate into consistent increases in overall performance. With the Stellus “scale-through” architecture, which is described as a separate category from “scale-up” and “scale-out” architectures, organizations have the option of adding DMs for increasing processing power or KVSs for increasing capacity. As each component is added, performance increases at a consistent rate.

To deal with unstructured data, Stellus has developed the KVS, leveraging the approach used by companies such as Google and Facebook. Unlike traditional block storage, a Key-Value Store allows for searching and identifying unstructured data via metadata. The metadata also supports the SDP’s high read and write performance as it enables fast unstructured data retrieval and storage of processed data. Global data maps and data caches are no longer required.

**Figure 2. Stellus’s ‘Scale-through’ Architecture**



Source: Enterprise Strategy Group

The SDP leverages 8TB hot-swappable, non-volatile memory express (NVMe) solid state drives (SSD) that help to reduce latency, achieve higher overall performance, and increase storage capacity in smaller form factors. Because they are hot swappable, critical applications can continue to run, minimizing any downtime due to hardware failure. The base Stellus

SDP configuration is delivered in a 5RU<sup>2</sup> rack mount package and has 184 TB of storage capacity; the largest configuration contains a maximum of eight DMs and eight KVSs, contained in 17RU with 1.475 PB of storage capacity.

Stellus also delivers a management service via the cloud so that organizations can manage and monitor the SDP. With client-level access analytics, organizations can gain insight and control over application performance and overall system reliability.

### Flexible Integration with SDS

As the Stellus Data Platform leverages SDS technology, it is designed to integrate easily into existing data center environments, offering support for a wide breadth of hardware options, thus having the flexibility to integrate new or different hardware technologies as they become available or demands shift. More than a third (37%) of organizations using or expecting to leverage SDS technology—such as the Stellus Data Platform—identify this ability to offer greater flexibility and choice of hardware as a key benefit driving interest in or use of SDS technology.

Simplifying technology integration is essential for modern IT environments, as nearly two-thirds of IT decision makers (64%) say IT is more complex compared with two years ago. There is more impact on organizations with mature digital transformation initiatives, as they are three times likelier than those with no digital transformation initiatives (29% versus 9%) to say IT is significantly more complex today. Of those experiencing this increase in complexity, 37% of organizations identify increases in data volumes as a top-five driver of that complexity, while 27% view digital transformation initiatives to leverage new technology to change operations as another top driver.<sup>3</sup> Organizations need technologies that can help address the massive growth in digital business demands, while simplifying the integration of those technologies into the data center.

### ESG Technical Validation

ESG performed evaluation and testing of the Stellus Data Platform at the Stellus facility in San Jose, California. Testing was designed to demonstrate the scalable, consistent performance of the Stellus “scale-through” architecture and validate the tangible business value offered by the platform.

#### Scale-through Performance

ESG tested using industry-standard tools and methodologies, and focused on validating Stellus performance using high-throughput streaming real-world workloads, as are used in multiple industries: media and entertainment, life sciences, seismic analysis, and IoT data acquisition and analysis. The goal of this phase of testing was to validate that Stellus can deliver tens of gigabytes per second of throughput with high availability while scaling performance and throughput independently.

#### ESG Testing

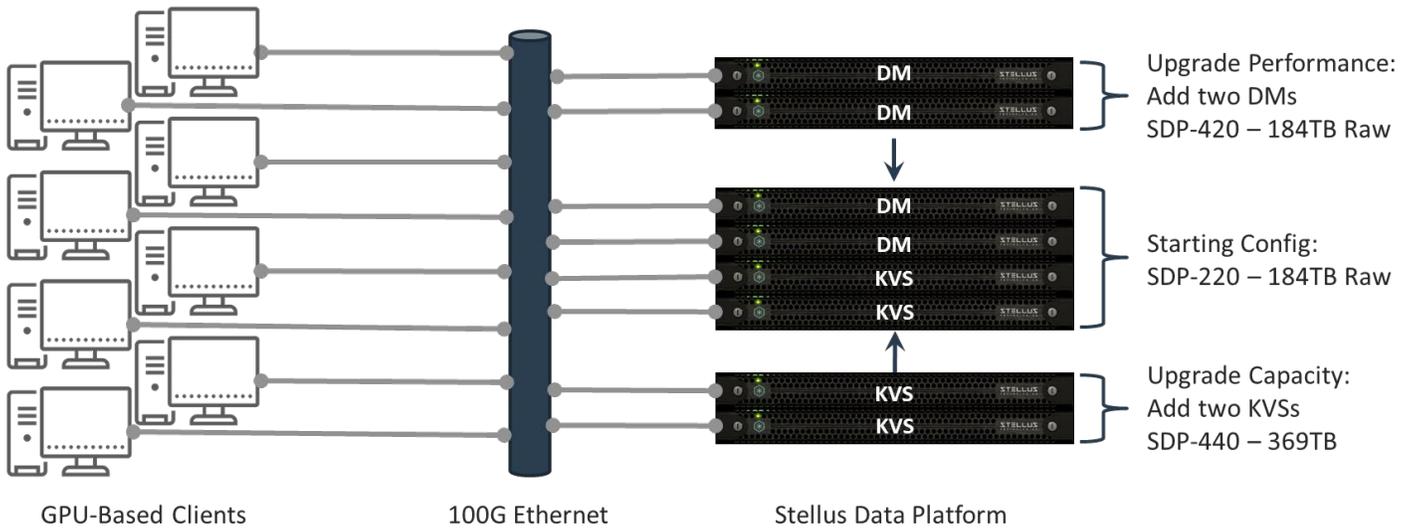
The test bed utilized by ESG consisted of 16 GPU-based clients, each running Davinci Resolve 16 video editing software. The workstations were connected to a Stellus SDP-220 via 100Gb Ethernet. The Stellus SDP-220 was populated with two DMs and two KVS enclosures, each populated with 24x 8TB NVMe drives. We had 32 different uncompressed 4K videos stored on the Stellus SDP.

---

<sup>2</sup> Dimensions of an RU (rack unit) are 1.75” H x 17.2” W x 34” D.

<sup>3</sup> Source: ESG Master Survey Results, [2020 Technology Spending Intentions Survey](#), January 2020.

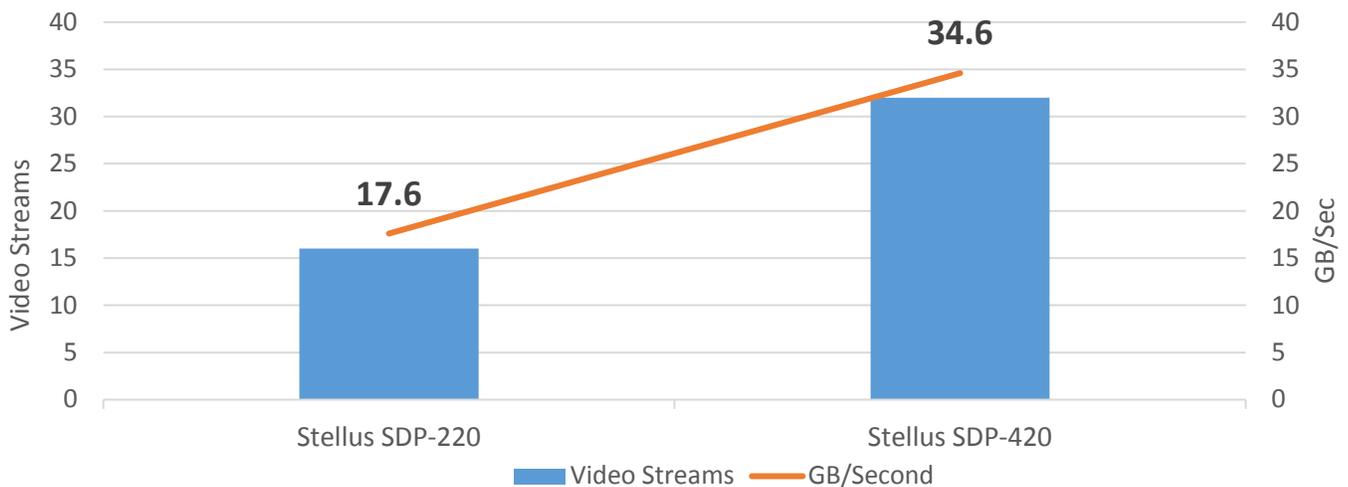
**Figure 3. The ESG Validation Test Bed**



Source: Enterprise Strategy Group

First, we started up sixteen simultaneous 4K video streams using eight of the workstations. The streams started up and began to run, and we verified that the Stellus SDP-220 was servicing 17.6 GB/second of video throughput in five RU without stutter, jitter, or dropped frames as we let this workload run for 30 minutes. Next, we upgraded the system to a Stellus SDP-420 by adding two DMs. The DMs were powered up and in the rack. Adding them to the cluster took just a couple of clicks, and they were online in less than 15 minutes. Once the upgrade was complete, we started up 32 simultaneous uncompressed 4K video streams. As seen in Figure 4, the Stellus Data Platform was then servicing 34.6 GB/sec—occupying only seven RU—and the video was playing smoothly on all clients without stutter, jitter, or dropped frames.

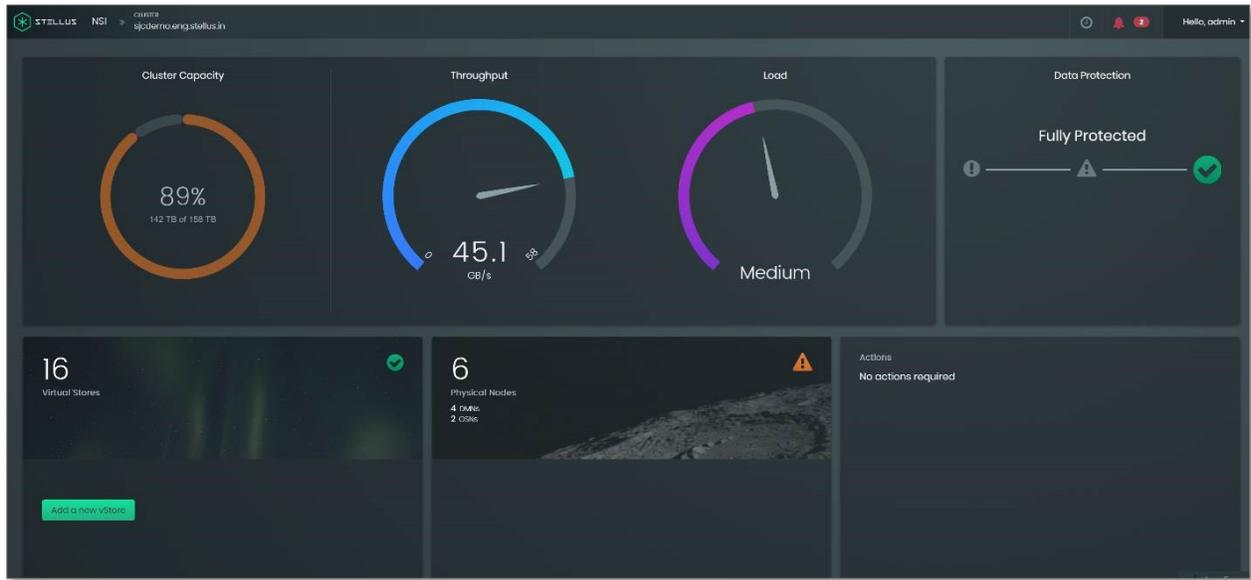
**Figure 4. Stellus Data Platform Streaming Performance**



Source: Enterprise Strategy Group

Next, we examined single-client, multi-threaded read and write performance on the SDP-420 and 440. The FIO benchmark—v. 3.13 for Windows—was used to generate the workloads, running on a single workstation connected via multiple 100Gbps Ethernet connections. As seen in Figure 5, the SDP-420 was able to sustain 45 GB/sec of reads.

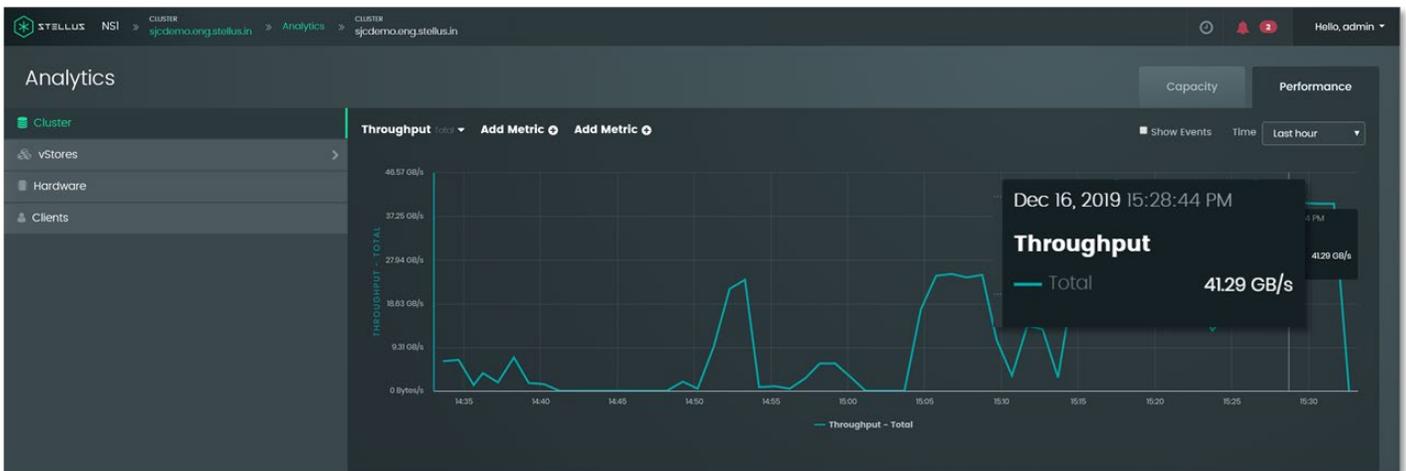
**Figure 5. Stellus Data Platform Single Client Read Performance**



Source: Enterprise Strategy Group

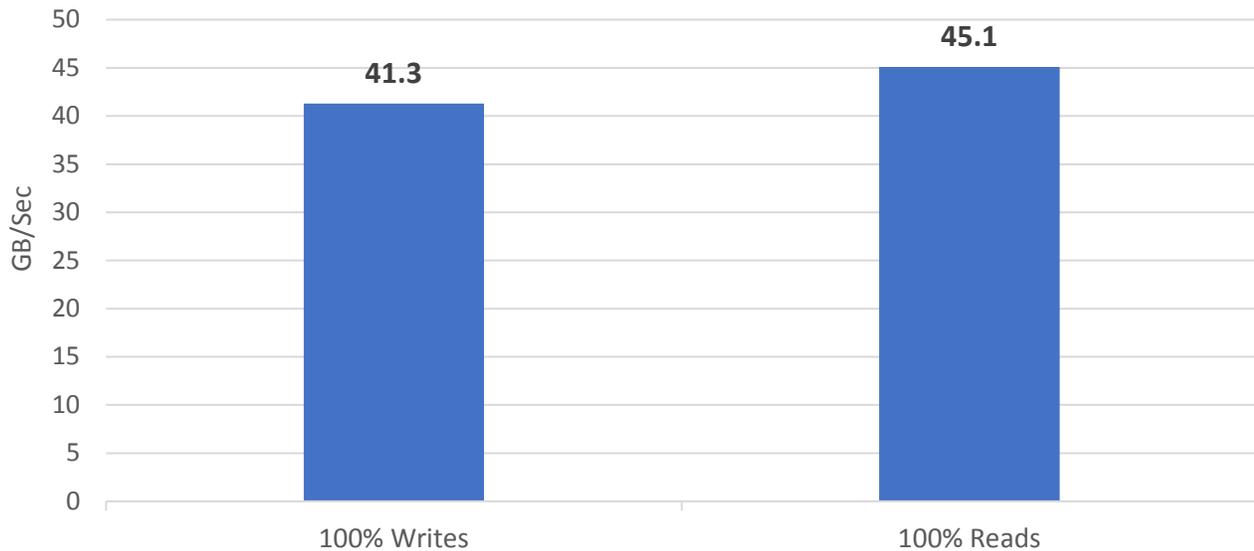
Next, we looked at writes. The same client was used to generate a 100% write workload against an SDP-440. Stellus was able to sustain 41.3GB/sec as seen in Figure 6.

**Figure 6. Stellus Data Platform Single Client Write Performance**



Source: Enterprise Strategy Group

The results of both tests are summarized in Figure 7.

**Figure 7. Stellus Data Platform Single Client Read and Write Performance**

Source: Enterprise Strategy Group

It's important to note that not only are read and write performance nearly equal, which is not the case for most file storage systems, but also both read and write performance in this test exceeded Stellus's guidance for performance of the SDP-420 and 440, which is 40 GB/sec. It's also important to note that response time across all tests never exceeded 500  $\mu$ sec.



### Why This Matters

The complexity and performance challenges associated with a modern storage environment can lead to overbuying storage capacity in order to achieve acceptable levels of performance. Because poor storage performance can result in lost sales, lost customer goodwill, lost productivity, and lost competitiveness, it's imperative that organizations keep their infrastructures performing to the best of their ability.

ESG has validated that the Stellus Data Platform delivers consistently impressive performance and sub-500  $\mu$ sec response times, and it does so in a high-density package. Stellus accomplishes this through a combination of deep file system and storage expertise that enabled the design of an architecture that leverages Key-Value Stores (KVSs), key-value-over-fabric (KVoF) networking, and algorithmic data locality, all working together to provide an enterprise-ready, scalable file system that has the potential to change the game.

The parity between read and write performance was particularly impressive, with the Stellus Data Platform exceeding the Stellus performance claims for the model under test, delivering 41.3 GB/sec write and 45.1 GB/sec read throughput, the fastest we have ever seen for a file storage system. These performance levels and near parity in reads and writes are essential to advancements in media and entertainment workloads like volumetric capture, life sciences research involving high-speed microscopy and genome sequencing, and all forms of the industrial internet of things (IIoT). Near parity reads and writes can create a competitive edge for digital enterprises that depend on write-intensive workloads. Stellus eliminates the latency inherent when existing NAS systems must force destage on writes.

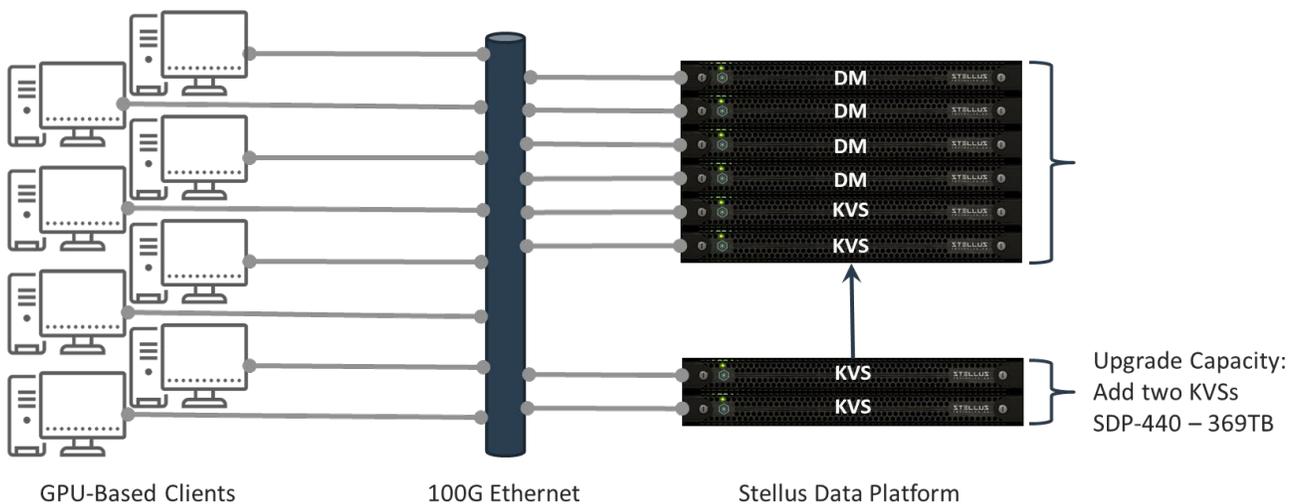
## Scalability, Availability, and Management

ESG examined scalability and availability of the platform by adding two Key-Value Stores to increase capacity and simulating a drive failure by pulling an SSD while the system was running the 16-stream workload, with 4K uncompressed video.

### ESG Testing

ESG observed the Stellus Data Platform’s performance when upgrading the system from an SDP-420 to an SDP-440 by adding two Key-Value Stores. With the new nodes powered on and connected to the network, adding them to the cluster took just a couple of clicks. Rebalancing the cluster—redistributing the data so it resided evenly across all four nodes—took 43 minutes, 12 seconds. Performance remained consistent as the cluster rebalanced.

**Figure 8. Scale-through Upgrade—Add Capacity**



Source: Enterprise Strategy Group

Next, ESG tested how the Stellus Data Platform handled an SSD failure while video continued to stream. We first tested how long the SDP would rebalance data after pulling one SSD. According to the Stellus Data Platform System Control, rebalancing took 8 minutes, 23 seconds. When we inserted a replacement drive and added it into the cluster, the rebalance completed in 16 minutes, 18 seconds.

We repeated the same test to verify that the rebalancing times were consistent. After removing a different drive from another KVS, time to rebalance was 8 minutes, 35 seconds. After replacing the drive, the cluster rebalanced in 16 minutes, 47 seconds. We should note that during both drive removals and rebalancing events, performance remained consistent at over 17GB/sec.

Finally, ESG looked at the manageability of the system. The Stellus System Control (SSC) provides a clean interface that we found easy to use to monitor and manage the system. Tasks like replacing SSDs and upgrading the system by adding KVSs and DMs took only a couple of clicks. Availability events like SSD failures triggered clear, concise alerts on the home screen. In addition, the Stellus Data Platform provides RESTful APIs with a Python add-on to provide data scientists with the ability to execute fine-grained automation and control.

## **i** Why This Matters

When asked to name their biggest challenges in terms of their on-premises storage environment, for file environments, both data protection (30%) and rapid data growth rates (25%) were among the most cited responses. A file storage system must address both challenges without compromising performance.

ESG validated that the Stellus Data Platform effectively addresses these issues. The fully redundant architecture is highly available and provides extremely low-latency performance during planned and unplanned outages. The platform rebalanced data after loss of a 7.68TB drive in a 92TB data space in less than nine minutes with minimal impact to system performance. System upgrades were fast and simple to execute, enabling ESG to double storage capacity with a couple of clicks and less than 45 minutes processing time. The system also offers impressive density, scaling to 1.475 PB in just 17 RU

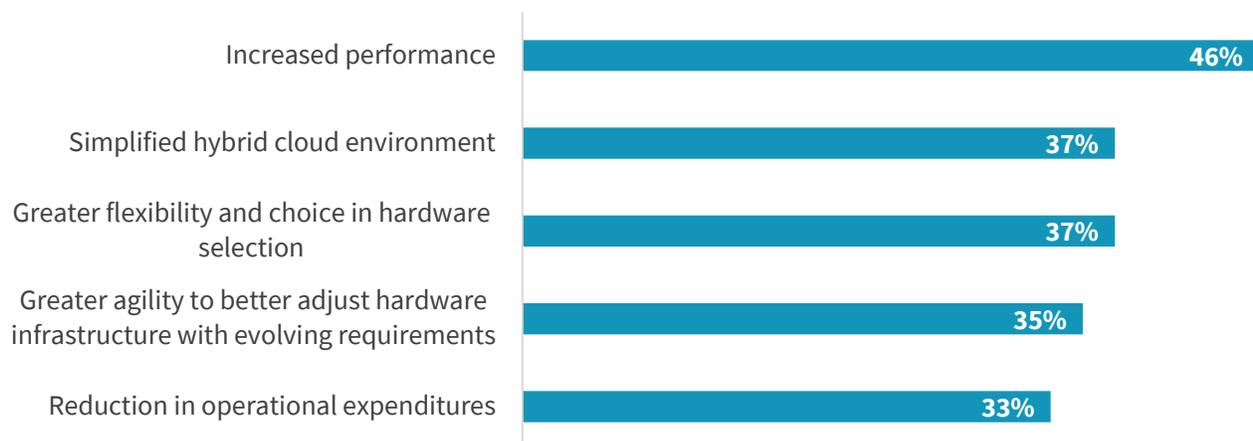
Business process improvement continues to be a high IT priority. One important area where businesses can improve process is system monitoring and management. ESG validated that the Stellus System Control makes monitoring and management of the Stellus Data Platform easy. Stellus's RESTful APIs and Python add-on enable data scientists to execute fine-grained automation and control.

## The Bigger Truth

Organizations are continuing to generate and store exceptionally large amounts of unstructured data. ESG uncovered that more than half of organizations expect their on-premises data to grow by at least 21% annually over the next three years. With the increasing adoption and use of data-intensive applications—the internet of things (IoT), artificial intelligence, and machine learning, to name just a few—organizations require a solution that can efficiently store and process data with consistently high read and write performance at parity. The solution should also scale in a manner that enables organizations to increase processing power and capacity independently. When scaling either processing power or capacity, read and write performance should increase consistently and predictably.

**Figure 9. Top Five Benefits of Software-defined Storage**

**What benefits has your organization realized—or does it expect to realize—as a result of deploying software-defined storage technology (SDS)? (Percent of respondents, N=334, five responses accepted)**



Source: Enterprise Strategy Group

The Stellus Data Platform has been designed to provide scalable and highly performant storage, addressing the continued growth of unstructured data in organizations. SDP can increase processing power or capacity independently by adding Data Managers and Key-Value Stores, respectively. The addition of either or both components results in consistent increases in both read and write performance. Read and write performance currently scale from 20 to 80 GB/s. The high, consistent performance of the SDP is also supported using key-value-over-fabric (KVof) and algorithmic data locality to simplify the storage and retrieval of data. NVMe SSDs enable the SDP to store large amounts of data, as the Stellus system can support up to 1.475 PB in a 17RU form factor. Stellus also supports advanced erasure coding, cloud integration, and the U.2 format.

ESG testing of the Stellus Data Platform validated that with more than 45 GB/sec of sustained streaming of read traffic in seven rack units and over 41GB/sec sustained write traffic in only nine rack units, it's the fastest we have ever seen for a file storage system. The platform also provided easy scalability, high availability through component failures, and simple management.

ESG looks forward to seeing how customers respond to this new composable storage platform. At this point, Stellus has invested heavily in design, testing, and integration to ensure high performance and smooth operation. The addition of features like automatic snapshots would provide even more value to customers.

If your organization is looking to gain a competitive edge as you use IT to transform the business into a digital enterprise, ESG believes that you should strongly consider the performance and operational advantages of fueling the data-driven enterprise with the Stellus Data Platform.

### Ease of Use and Maintenance

**Stellus System Control (SSC) provides a clean interface that we found easy to use to monitor and manage the system. Tasks like replacing SSDs and upgrading the system by adding KVSs and DMs took only a couple of clicks. Stellus scale-through upgrades are fast and easy: DMs and KVSs were added to the cluster with three clicks. The DMs were online in less than 15 minutes. With KVSs, rebalancing in the background completed in less than an hour.**

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



**Enterprise Strategy Group** is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.

