

Technical Review

Huawei FusionStorage Cloud Storage

Date: November 2018 **Authors:** Tony Palmer, Senior Validation Analyst and Alex Arcilla, Validation Analyst


Abstract

This ESG Technical Review documents hands-on testing of Huawei FusionStorage Cloud Storage. Testing focused on ease of deployment and management; consistent, high performance for real world workloads; availability and fault tolerance; and non-disruptive, online scalability.

The Challenges

Software-defined storage (SDS) can transform data centers, giving IT infrastructure the agility businesses need to remain competitive in this increasingly digitally-defined economy. SDS is a key enabler for software-defined data centers, hyperconverged infrastructure (HCI), and private cloud storage services. To better understand how SDS technology helps transform data centers, ESG conducted a research study that surveyed 303 storage decision makers. To qualify, respondent organizations had to currently be utilizing or evaluating SDS or interested in SDS as a long-term strategy. As part of this study, storage decision makers currently utilizing SDS technology were asked to identify the benefits that their organizations realized because of deploying SDS technology (see Figure 1).¹

Figure 1. Benefits Realized Resulting from Deploying SDS Technology



Source: Enterprise Strategy Group

¹ Source: ESG Research Report, [Software-defined Storage \(SDS\) Market Trends](#), February 2017.

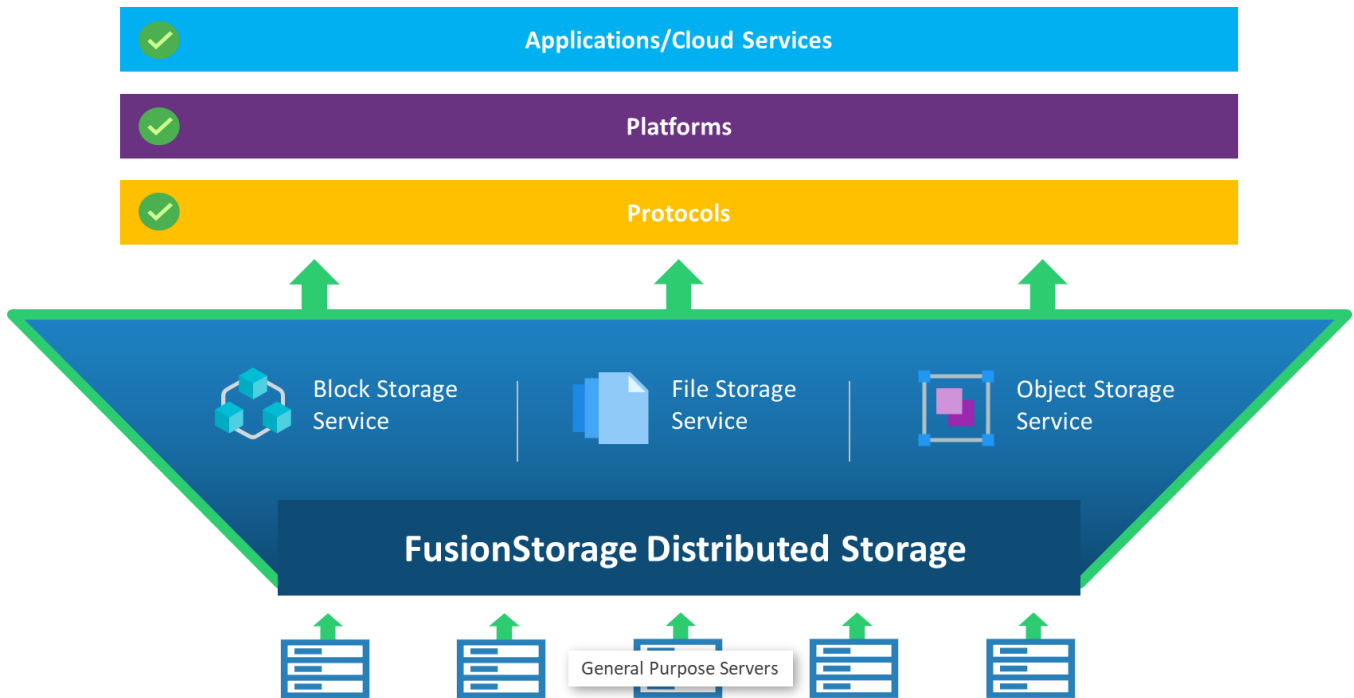
The most commonly identified benefits aligned along two main themes critical for IT transformation. Expedited deployment (34%) and simplified management (32%) enable organizations to optimize time and personnel, so that revenue generating IT initiatives are not delayed by complex, manual processes. The ability to cost-effectively scale capacity and performance up and down as demands evolve allows companies to leverage SDS as a critical component that can enable a web-scale architecture, delivering the benefits of speed and flexibility, while keeping the cost of ownership—both capital and operational—down.

The Solution: Huawei FusionStorage Cloud Storage

Huawei FusionStorage is a software-defined cloud storage solution designed for cloud-based architectures. The on-board storage system software combines the local storage resources of general-purpose servers into fully distributed storage pools, providing block, object, or file storage services to the upper layers of the stack via industry-standard protocols—SCSI, iSCSI, NFS, SMB, FTP, HTTP, S3, Swift, Cinder, RESTful API, and Manila, for example. FusionStorage provides IOPS, bandwidth, and expansion capabilities for structured, unstructured, and semi-structured data. FusionStorage is engineered to maximize performance, availability, and efficiency. Active-active clustering, redirect-on-write snapshots, thin provisioning, remote replication, erasure coding, and other functionality work together to meet the needs of enterprises and service providers. Huawei provides open APIs via standardized protocols, interfacing seamlessly into Open-stack cloud-based architectures and Hadoop big data ecosystems.

Active-active cluster architecture and automatic load balancing enable FusionStorage to provide seamlessly scalable performance and capacity, which makes the product an excellent fit for deployments in enterprise private and hybrid clouds, test and development clouds, government clouds, public security clouds, carrier public clouds, and other environments that require converged storage capabilities.

Figure 2. Huawei FusionStorage Cloud Storage



Source: Enterprise Strategy Group

ESG Tested

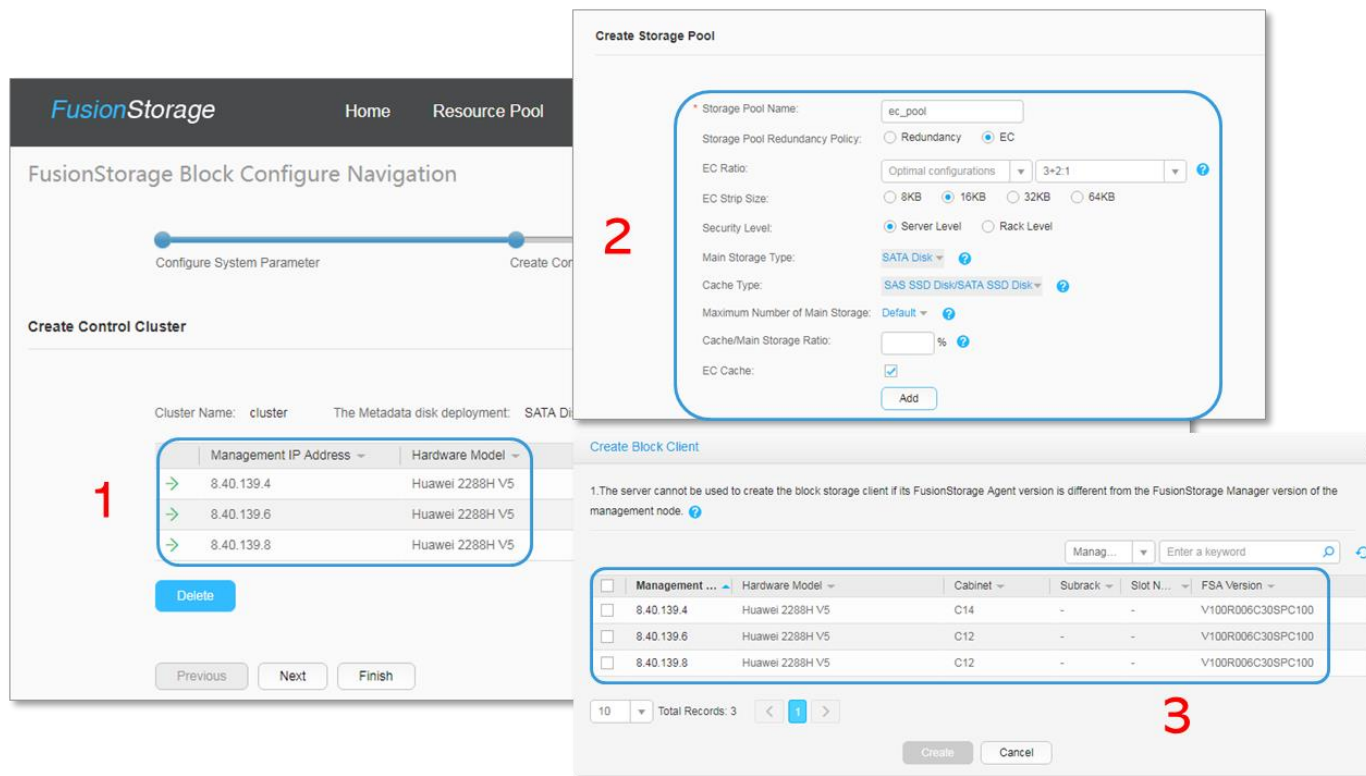
ESG tested Huawei FusionStorage Block to evaluate it as a fully distributed, software-defined cloud storage platform. ESG was interested in examining the ease of deploying FusionStorage on general purpose servers, the performance of the platform under multiple real-world workloads, and the resilience and availability of the system, including using an active-active cluster to extend an Oracle RAC implementation across multiple data centers.

The test environment included two FusionStorage Block clusters, one running on four Huawei FusionServer RH2288H servers and one running on four Dell PowerEdge R910 servers. Eight Huawei FusionServer E9000 blade servers ran VMware vSphere and VMware Horizon 7 desktop virtualization software. Two FusionServer RH2288H servers ran Oracle Real Application Clusters (RAC). Connectivity was provided by two Huawei S6720-HI Series Agile 10 GbE Switches. Server connectivity to the FusionStorage cluster was over 10GbE. Workloads were generated using multiple industry-standard benchmarking utilities: SLOB and Swingbench were used to generate database workloads on Oracle RAC, and Horizon View was used to generate workloads in the virtual desktop environment.

Simplicity and Flexibility

First ESG performed a deployment on four Huawei FusionServer RH2288H servers. Each server had dual Intel Xeon E5-2670 CPUs, 128GB of RAM, 2x 1.6TB NVMe SSDs, and 12x 4TB SATA HDDs. We added servers manually using their IP addresses. This part of the installation can be automated to make it easier to deploy larger clusters. We selected the software package to install and clicked next. Software installation completed in less than three minutes and the deployment wizard walked us through system parameter configuration, control cluster creation, storage pool creation, and finally creating and configuring block storage clients.

Figure 3. Huawei FusionStorage Installation and Deployment



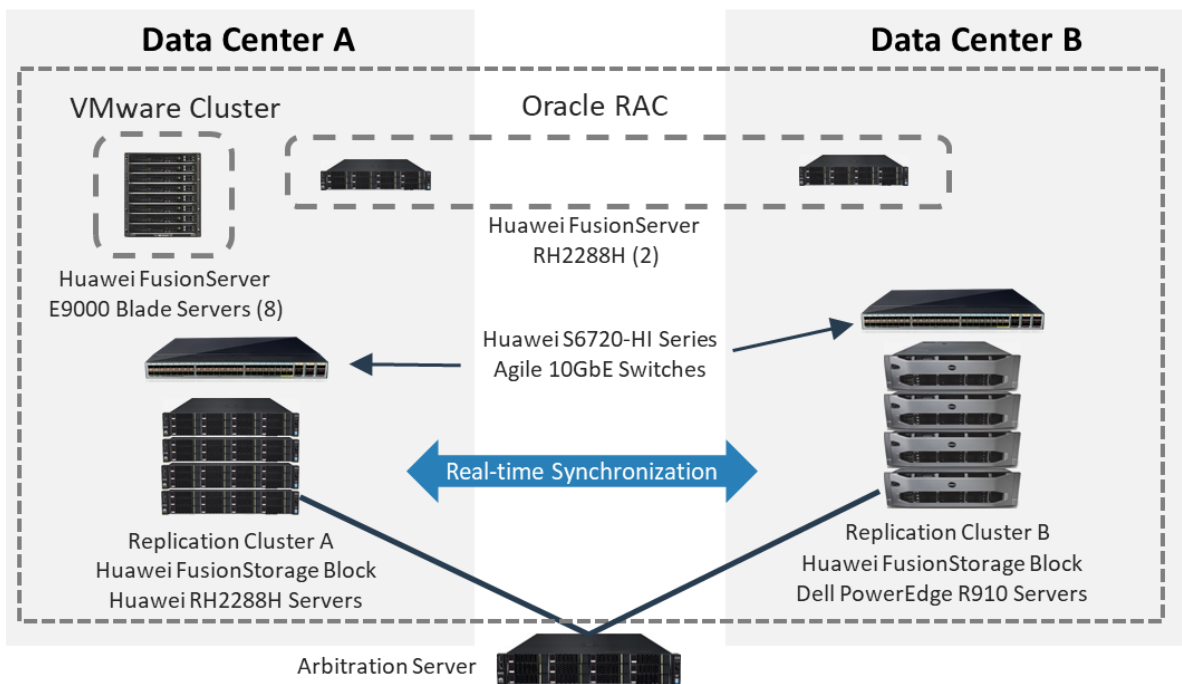
Source: Enterprise Strategy Group

For our three-node cluster, the entire process—from bare metal servers to FusionStorage Block cluster—took about 10 minutes.

Resilience and Availability

Next, ESG looked at the resilience and high availability features of the FusionStorage system, examining Huawei’s implementation of active-active clustering, snapshots, and erasure coding. To test active-active cluster availability, we deployed Oracle RAC on FusionStorage clusters stretched across two simulated data centers. Replication cluster A and B were attached to different network segments in different parts of the network data center connected by a single network connection simulating a metropolitan area network connection.

Figure 4. The FusionStorage Block Active-active Cluster Test Bed



Source: Enterprise Strategy Group

Workloads on Oracle RAC were generated with the Swingbench OrderEntry benchmark. Swingbench is a freely available load generator designed to stress test Oracle databases.² The OLTP workload was started and, once ramped up, was generating 16,000 transaction per minute. At this point we simulated a system outage by disabling the storage pool in replication cluster B. I/O paused for approximately 20 seconds, then resumed, being serviced by replication cluster A. The amount of time for the I/O pause depends on the failure and on the hosts. An active-active cluster failover usually pauses for 16-20 seconds. Host multi-path software will also impact the I/O pause time, with many hosts pausing automatically for 30 seconds. After re-enabling the storage pool in replication cluster B, the workload resumed on both sides of the cluster automatically, and data was resynced with no interruption.

ESG also tested the performance impact of snapshot functionality, creating snapshots of volumes while the system continued to run the OLTP workload. Snapshots were enabled and created instantly with just a few clicks. Snapshots can also be scheduled using the Huawei GUI and application consistency groups can be created. Once the snapshots were created, we edited a file and deleted multiple files on one Windows system. Rolling the volume back to the snapshot was also quite simple, and the entire process had no measurable impact on the performance of the system. This was not surprising, since FusionStorage uses the same redirect-on-write technology as OceanStor Dorado to provide snapshots without the impact of copy on first write. ESG tested snapshots under load on the Huawei OceanStor Dorado platform in 2017.³

² <http://www.dominicgiles.com/swingbench.html>

³ Source: ESG Lab Review, [Huawei OceanStor Dorado V3 All-flash Storage](#), September 2017

Next, we tested the impact of a failed disk, and the time required to recover. FusionStorage supports erasure coding data protection. Erasure coding protects data by breaking it into fragments and encoding them with a configurable number of redundant pieces of data, so they can be stored across different locations, disks, storage nodes, or geographical locations, for example. In this test we used a 5.96TB data volume, built from 1.4TB NVMe SSDs using a 3+2:1 erasure code pool. Commonly, 3+2 erasure coding requires five disks—For every three data blocks, two parity blocks are required. 3+2:1 uses only three disks for better capacity utilization, with lower redundancy. Note: FusionStorage Block provides a wide range of N+M or N+M:B EC redundancy ratio configurations for users to choose from. N represents the number of data copies, and values of two to 20 are supported; M represents the number of redundant, or parity, copies and can range from one to four. B represents use of the Huawei-developed low-density erasure coding (LDEC) algorithm, which reduces the total number of disks required to support a given level of EC redundancy. B is always equal to one.

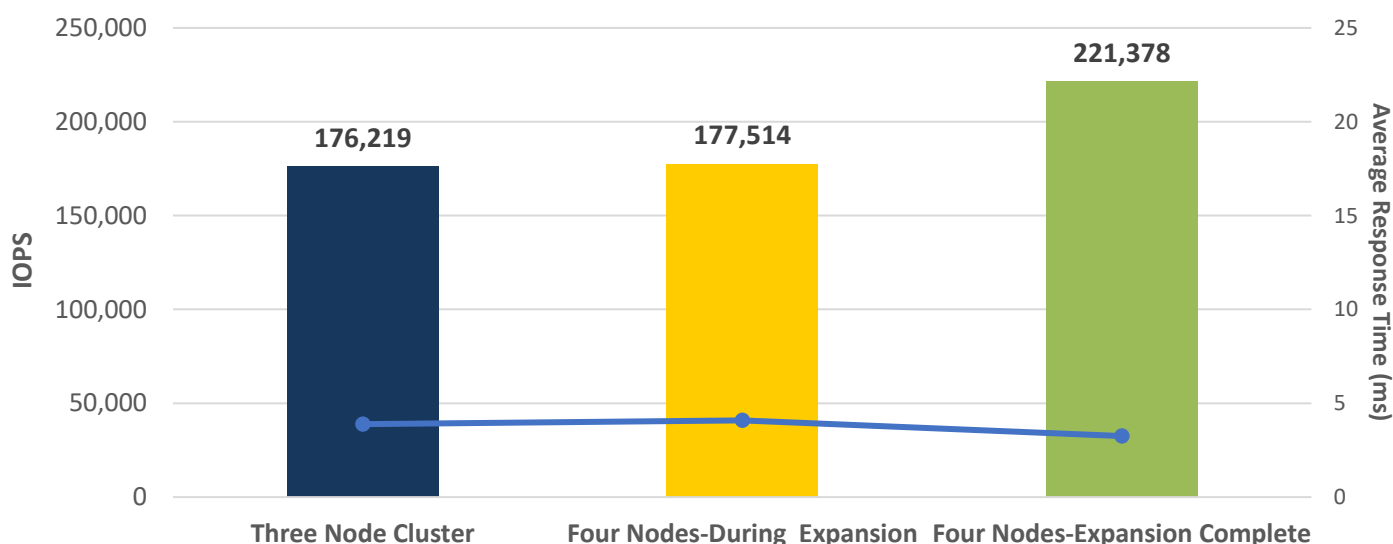
The volume contained 1.2TB of data. We failed a disk in the pool and the total capacity of the data volume shrunk to 5.3TB. Once the system confirmed that the drive was faulty, the rebuild of the volume took 9 minutes.

Performance and Scalability for Real-world Workloads

Finally, ESG examined the performance and scalability of the FusionStorage platform running multiple enterprise workloads. VMware Horizon 7 was used to create a 500-seat virtual desktop environment, and Horizon View Planner was used to generate workload. Swingbench and SLOB were used to generate OLTP database workloads, and ESG audited Huawei FusionStorage’s published SPC-1 results.⁴

First, ESG tested mixed workload performance. We started Horizon View Planner with 500 desktops. While desktops were powering on and the workload was ramping up, we started the OLTP workload. At this point the nodes were servicing 30,000 IOPS each, with a total of 120,000 IOPS for the whole cluster. These workloads ran for 90 minutes, and the average response time during this run was 3.0ms. We took multiple snapshots while the workload was running with no detectable impact on performance.

Figure 5. FusionStorage Block OLTP Workload Performance—Before, During and after Expansion



Source: Enterprise Strategy Group

Before the expansion, the three-node cluster was servicing an average of 176,219 OLTP IOPS with average latency of 3.89ms. When the fourth node was added, the system began to redistribute the data. Performance remained steady, averaging 177,514 OLTP IOPS with average latency of 4.09ms. Once the data redistribution was complete, performance

⁴ <https://spcresults.org/benchmarks/results/top10/performance/spc1/3>

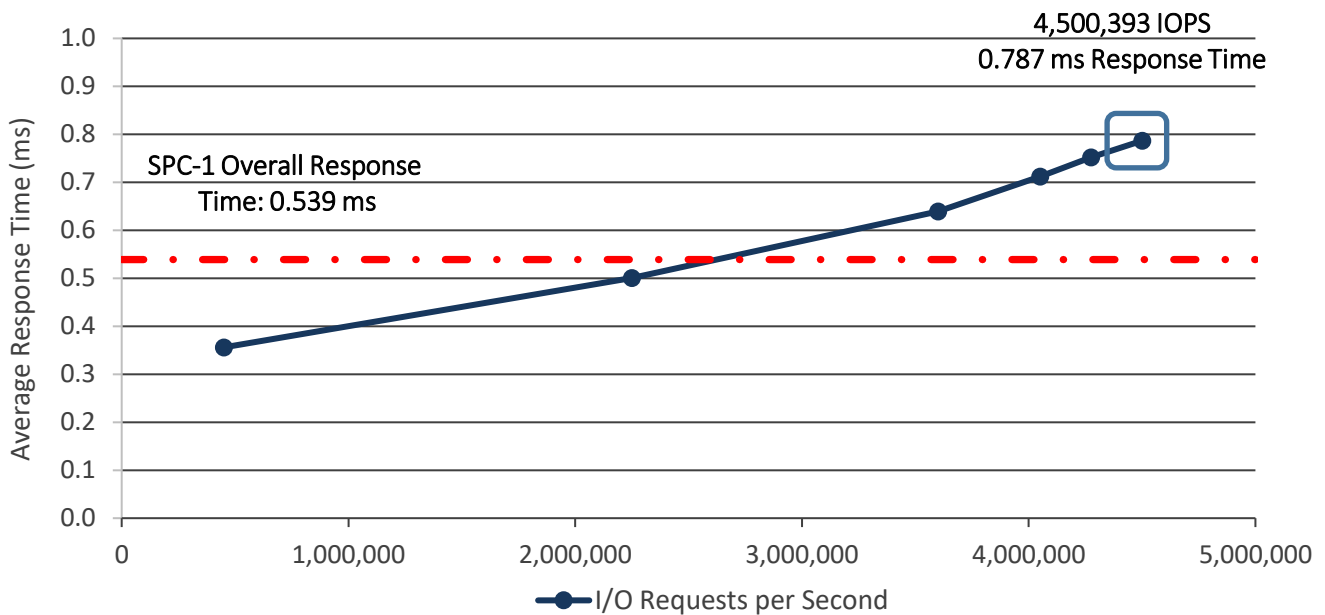
increased to 221,378 OLTP IOPS and response time dropped to 3.25ms. With no disruption, performance increased 25.63%, showing near-linear scalability. Multiple tests were run, and the cluster scaled up and down with little impact or overhead. Huawei reports that the FusionStorage platform supports up to 4096 nodes. The platform is currently running in multiple sizable public and hybrid cloud environments including Huawei’s own public cloud service with thousands of nodes and hundreds of petabytes of data as of this writing.

SPC-1 Results

ESG reviewed Huawei’s published results of the SPC-1 application-level benchmark suite maintained by the Storage Performance Council (SPC), a vendor-neutral storage industry standards body. SPC-1 testing generates a series of workloads designed to emulate the typical functions of transaction-oriented, real-world database applications. These applications are typically characterized by random I/O and generate both queries (reads) and updates (writes). Real-world examples of this application type include OLTP, database operations, and mail server implementations. SPC results can be roughly mapped into metrics such as the number of credit card authorizations executed per second. We should note that the SPC-1 benchmark consists of over 60% writes, a mix of random and sequential I/O, and a variety of block sizes. Results should not be compared with marketing performance numbers consisting of 100% random reads with a homogeneous block size.

At the time of publication, Huawei occupies the top three spots in SPC-1, Version 3 testing, ranked by the maximum number of IOPS supported.⁵ FusionStorage Block—all-flash—SPC-1 test results currently hold the third spot, achieving 4,500,393 SPC-1 I/O requests per second at 100% load with an average response time of only 0.787 milliseconds. Figure 6 shows a response time/throughput curve, which visually represents the performance of the system under test as load is increased. A long, flat curve indicates better performance, as this denotes that response time stays low as IOPS increase.

Figure 6. Huawei FusionStorage All-flash SPC-1 Throughput and Response Time



Source: Enterprise Strategy Group

⁵ <https://spcresults.org/benchmarks/results/top10/performance/spc1/3>

Why This Matters

Software-defined storage is reported to deliver many benefits to organizations surveyed by ESG.⁶ Expedited deployment (34%) and simplified management (32%) enable organizations to optimize time and personnel and the agility to cost-effectively scale capacity and performance to align with evolving business requirements. In addition, consolidating workloads driven by physical and virtualized systems onto a single storage platform can help drive higher levels of infrastructure efficiency through improved resource utilization, but when multiple applications share the same underlying storage platform, contention for resources can significantly impact all applications, leading to poor response times, lost productivity, and, in the worst case, lost revenue.

ESG Lab validated that two enterprise application workloads were easily consolidated onto a single Huawei FusionStorage Block cluster without impacting one another. A demanding VDI infrastructure that supported 500 heavy users was serviced simultaneously with an OLTP application. ESG confirmed that a three-node FusionStorage Block cluster with a hybrid SDD/HDD configuration sustained more than 175,000 IOPS with an average response time of under 4 ms. ESG also validated the completely non-disruptive scalability of FusionStorage; a fourth node was added to the cluster with zero impact while the data was redistributed to the new node, and the additional performance resources were immediately available once the migration completed.

In its annual IT spending intentions survey, ESG asked organizations about their plans for data center modernization.⁷ Improving data backup and recovery was cited by 31% of respondents, the second most-cited response. Clearly resilience and data protection are top of mind. Features like active-active clustering, redirect-on-write snapshots, and fast, low impact rebuilds from drive failures are important, because they reduce or eliminate the need to recover data from a backup.

ESG was impressed with how easy it was to set up FusionStorage Block in an active-active cluster across two simulated data centers. This capability is intrinsic to the FusionStorage platform and provides active-active access to data by applications across metro cluster distances without requiring additional software. The process of connecting the nodes and configuring the cluster took just a few clicks—without downtime, professional services, or interruptions to productivity. Administrators could easily set it up with no specialized storage skills. The system was able to take point in time snapshots of active volumes with no performance impact, and recovery of 1.2TB of data in a 5.9TB volume after a hard disk failure was accomplished in just over nine minutes.

⁶ Source: ESG Research Report, [Software-defined Storage \(SDS\) Market Trends](#), February 2017.

⁷ Source: ESG Research Report, [2018 IT Spending Intentions Survey](#), February 2018.

The Bigger Truth

SDS technology plays a central role in achieving a cloud-like data center, offering far more value than simply the ability to leverage lower cost hardware. By transitioning away from manual, slow, and costly IT procedures to an agile, dynamic, and automated cloud-like services model, businesses are far better equipped to compete in this increasingly digitally defined economy. IT decision makers need to evaluate SDS based on its ability to deliver cloud-level capabilities in flexibility, scale, performance, and management. In addition, SDS should deliver predictability and transparency while enabling tremendous flexibility in both application and hardware infrastructure support. Constantly updating hardware and migrating data from one system to the next has become far too time-consuming and costly for businesses to stay competitive in this modern digital era. SDS can free IT from the constant manual complexities of day-to-day IT and deliver an automated cloud-like infrastructure.

Through its validation process, ESG has found that Huawei FusionStorage is a viable software-defined cloud storage solution well-suited to enable a cloud-like infrastructure within organizations. Huawei has engineered FusionStorage to provide IOPS, bandwidth, and expansion capabilities for structured, unstructured, and semi-structured data. FusionStorage supports active-active clustering, redirect-on-write snapshots, thin provisioning, remote replication, erasure coding, and other functionality to automate operational and administration activities, enabling organizations to respond quickly to business demands in a digital economy.

ESG tested Huawei FusionStorage for its ease of deployment on general-purpose servers, its resiliency and availability under various failure scenarios, and the performance of the platform under multiple real-world workloads. We found that organizations can deploy and configure a 52TB, three-node FusionStorage cluster, including bare metal servers and the FusionStorage Block cluster, in approximately 10 minutes. We observed how active-active clustering, snapshots, and erasure coding contribute to FusionStorage's resiliency, noting how they minimally impact user accessibility and system performance. Finally, ESG observed how the FusionStorage system can maintain high performance when expanding the storage cluster while continuing to run mixed workloads. FusionStorage's high performance has also been validated via SPC-1 v3 testing and is currently noted as the third highest IOPS achieved at the time of this publication.

ESG validated that the FusionStorage platform delivers high performance and availability at consistently low response times in a mixed workload, highly virtualized environment. Should your organization wish to leverage SDS in achieving a cloud-like storage infrastructure to efficiently satisfy business demands, ESG suggests that Huawei FusionStorage is worth serious consideration.

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. 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.

The goal of ESG Validation reports is to educate IT professionals about information technology solutions for companies of all types and sizes. ESG Validation reports 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.