

DATA DRIVEN GLOBAL VISION CLOUD PLATFORM STRATE  
ON POWERFUL RELEVANT PERFORMANCE SOLUTION CLO  
VIRTUAL BIG DATA SOLUTION ROI FLEXIBLE DATA DRIVEN

WHITE PAPER

# Hitachi Virtual Storage Platform G1000 All Flash Array With Hitachi NAS Platform 4100 8-Node Cluster: SPECsfs2008 Performance Analysis

By Hitachi Data Systems

April 2014

## Contents

Executive Summary .....	2
Performance Summary .....	3
Test Setup and Methodology .....	4
Detailed Component Summary and High-Level Overview.....	4
Appendix A: Contributors .....	6

## Executive Summary

IP-based network storage, for files and objects, is growing at exponential rates in the enterprise today. Ethernet network infrastructures continue to increase in performance, while new advances in server virtualization, software-defined data centers (SDDC) and cloud services merely reflect where enterprise users are headed. The Hitachi Virtual Storage Platform G1000 and Hitachi NAS (HNAS) Platform family of products provide the capabilities and features for consolidation. They help IT managers consolidate many, if not all, of their server-based filers and NAS appliances into fewer Hitachi systems. This approach leads to not only immediate capital expenditure (capex) and operating expense (opex) savings. It also delivers power, cooling and space savings in the data center; and it provides a storage platform that is scalable for future needs.

Hitachi Virtual Storage Platform G1000 and Hitachi NAS Platform models 4100, 4080, 4060 and 4040 deliver best-in-class performance and scalability. They provide smart primary deduplication, active-active clustering with automated failover, 99.999% availability and nondisruptive upgrades. They provide intelligent file tiering with policy-based automated migration, pointer-based snapshots, file and directory cloning, and large 256TB file system pools, a single namespace up to the maximum usable capacity. The HNAS family offers a universal migrator feature and both the HNAS family and VSP G1000 integrate with the Hitachi Command Suite of management and data protection software.

HNAS file module uses a hardware-accelerated "hybrid-core" architecture that accelerates network and file protocol processing to achieve the industry's best performance in terms of both throughput and operations per second (IOPS). Additionally, the architecture also accelerates primary dedupe to avoid file sharing workload degradation. HNAS file module uses the object-based Hitachi NAS Silicon File System and virtualization to deliver the highest scalability in the market. HNAS enables organizations to consolidate file servers and other NAS devices into fewer nodes and storage arrays for simplified management, improved space efficiency, and lower energy consumption. Each Hitachi NAS Platform 4100 node or cluster can scale up 32PB of usable data storage.

Hitachi Virtual Storage Platform G1000 offers the highest performance and scalability in its class and includes flash-optimized system software and patented Hitachi Accelerated Flash storage to improve application performance. Utilizing external storage virtualization and automated tiering, VSP G1000 centralizes storage management of multiple storage tiers for the highest economic value. VSP G1000 all flash storage system, when combined with the HNAS file module, delivers excellent performance and dramatically lowers response times.

Hitachi Data Systems tested the Hitachi NAS Platform 4100 8-node cluster, using VSP G1000 all flash storage system. Hitachi NAS Platform 4100 8-node cluster delivered an overall performance of 1,222,089 SPECsfs2008\_nfs.v3 operations per second with an overall response time of 0.75 ms. (<http://www.spec.org/sfs2008/results/res2014q2/sfs2008-20140331-00245.html>). When compared to a similar submission in its class<sup>1</sup>, even when using fewer number of flash drives this HNAS 4100 8-node solution delivered 110% more IOPS and 4% better overall response time.

This report demonstrates the dedication Hitachi maintains to the SPEC.org benchmark standards, and is a continuation of our past 2-node and 4-node submissions. When compared to all the submissions, as of this writing<sup>2</sup> the HNAS 4100 2-node configuration still delivers the industry's best overall response time.

---

<sup>1</sup> Results as of April 23, 2014. Refers to this submission: <http://www.spec.org/sfs2008/results/res2013q3/sfs2008-20130819-00226.html>.

<sup>2</sup> Results as of April 23<sup>rd</sup> 2014. For more information, go to <http://www.spec.org/sfs2008/results/sfs2008.html>.

## About the SPECsfs2008 Benchmark

SPECsfs2008 is the current version of the Standard Performance Evaluation Corporation benchmark suite measuring file-server throughput and response time; it provides a standardized method for comparing performance across different vendor platforms. SPECsfs2008 results summarize the server's capabilities regarding the number of operations that can be handled per second, as well as the overall latency of the operations. The suite is a follow-on to the **SFS97\_R1** benchmark, with an updated NFSv3 workload, support for additional client platforms, and a new test harness and reporting or submission framework. SPEC and the benchmark name SPEC SFS are registered trademarks of the Standard Performance Evaluation Corporation. For the latest SPEC SFS2008 benchmark results, visit <http://www.spec.org/sfs2008/>.

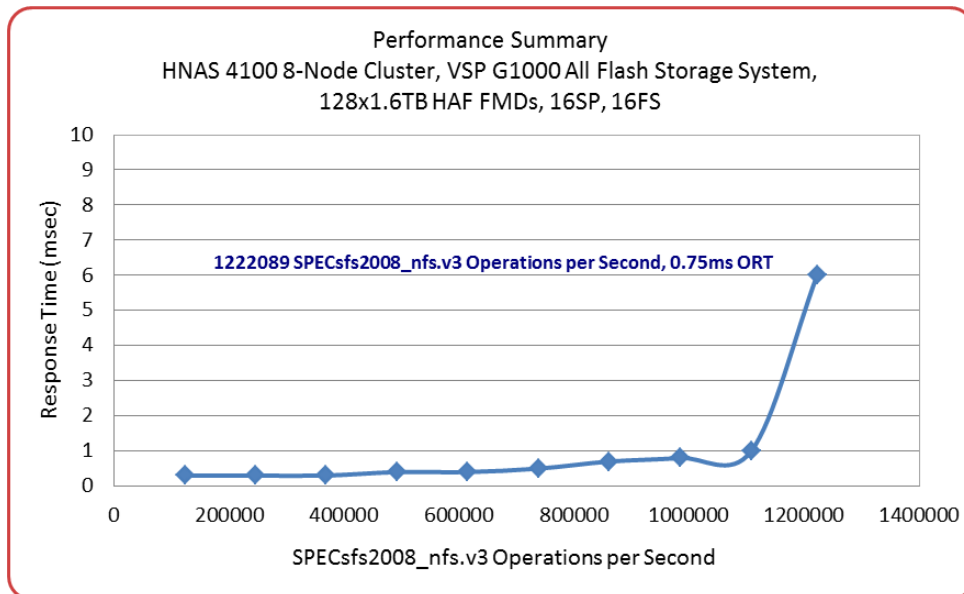


## Performance Summary

The tests focused on measuring NFS performance of the Hitachi NAS Platform 4100 8-node cluster (see Figure 1). NFS is one of the most commonly used distributed file systems in NAS solutions and is the standard file-sharing mechanism used in UNIX and Linux environments. Sixteen load-generating clients were used for the tests, each driving up to 448 threads on the HNAS file module cluster at a time.

As mentioned in the executive summary, the HNAS 4100 8-node cluster with VSP G1000 storage system delivered an overall performance of 1,222,089 SPECsfs2008\_nfs.v3 operations per second with an overall response time of 0.75 ms. *When compared to a similar submission in its class, even when using fewer number of flash drives this Hitachi solution delivered 110% more operations per second and 4% better overall response time.* These results show that the HNAS file modules offer capacity efficient, highly optimized, enterprise-class performance NAS services in NFSv3 environments. Figure 1 summarizes the performance of the Hitachi NAS Platform 4100 8-node cluster.

Figure 1. Performance Summary



## Test Setup and Methodology

The test bed consisted of a Hitachi NAS Platform 4100 8-node cluster, a Hitachi Virtual Storage Platform G1000 all flash storage system with 4 virtual storage director (VSD) pairs, and a Hitachi Cable Apresia 15000-64XL-PSR Ethernet 10GbE switch. It also included 2 Brocade 6520 Fibre Channel Switches, 16 NFS clients (Hitachi Compute Blade 2000 E55R3 blades) and 2 Brocade VDX-6730 cluster interconnect 10GbE switches. VSP G1000 was configured with 1TB cache memory and up to 128 x 1.6TB Hitachi Accelerated Flash module drives. Each RAID-5 (7D+1P) parity group had 8 LDEVs and a total of 128 LDEVs were created across the 128 drives. The LDEVs were distributed across the 32 VSP G1000 8G/sec Fibre Channel front-end ports. Each LDEV was presented to 2 front-end ports, thereby ensuring multipathing and failover capabilities. The HNAS 4100 cluster was connected to VSP G1000 via a redundant pair of Brocade 6520 Fibre Channel switches. A zone was created on each Fibre Channel switch. Each of the HNAS 4100 file module nodes were connected to each zone via two 8Gb/sec Fibre Channel ports. VSP G1000 was connected to each zone via thirty-two 8Gb/sec Fibre Channel ports, providing the high-speed I/O paths from the server to the storage.

Two storage pools were created on each HNAS 4100 file module node, with 8 LDEVs assigned to each storage pool. One file system was created on each storage pool using the 4KB file system block size along with 1 NFS export on each file system. Only one 10GbE NIC interface integrated into the file module node was used from each node during these tests. Sixteen Hitachi Compute Blade 2000 E55R3 blade servers were used as load-generating NFS clients for driving the client workloads. The assignment of processes to file systems was done in such a way that the processes were uniformly divided across all file systems for optimum load balancing. For network connectivity among the clients and the HNAS 4100 file module clusters, a Hitachi Cable Apresia 15000-64XL-PSR switch was used, all in a 10GbE network environment. Each NFS client had a dual-port Intel X520-SR2 PCIe 10 Gigabit Ethernet (GbE) PCIe network interface.

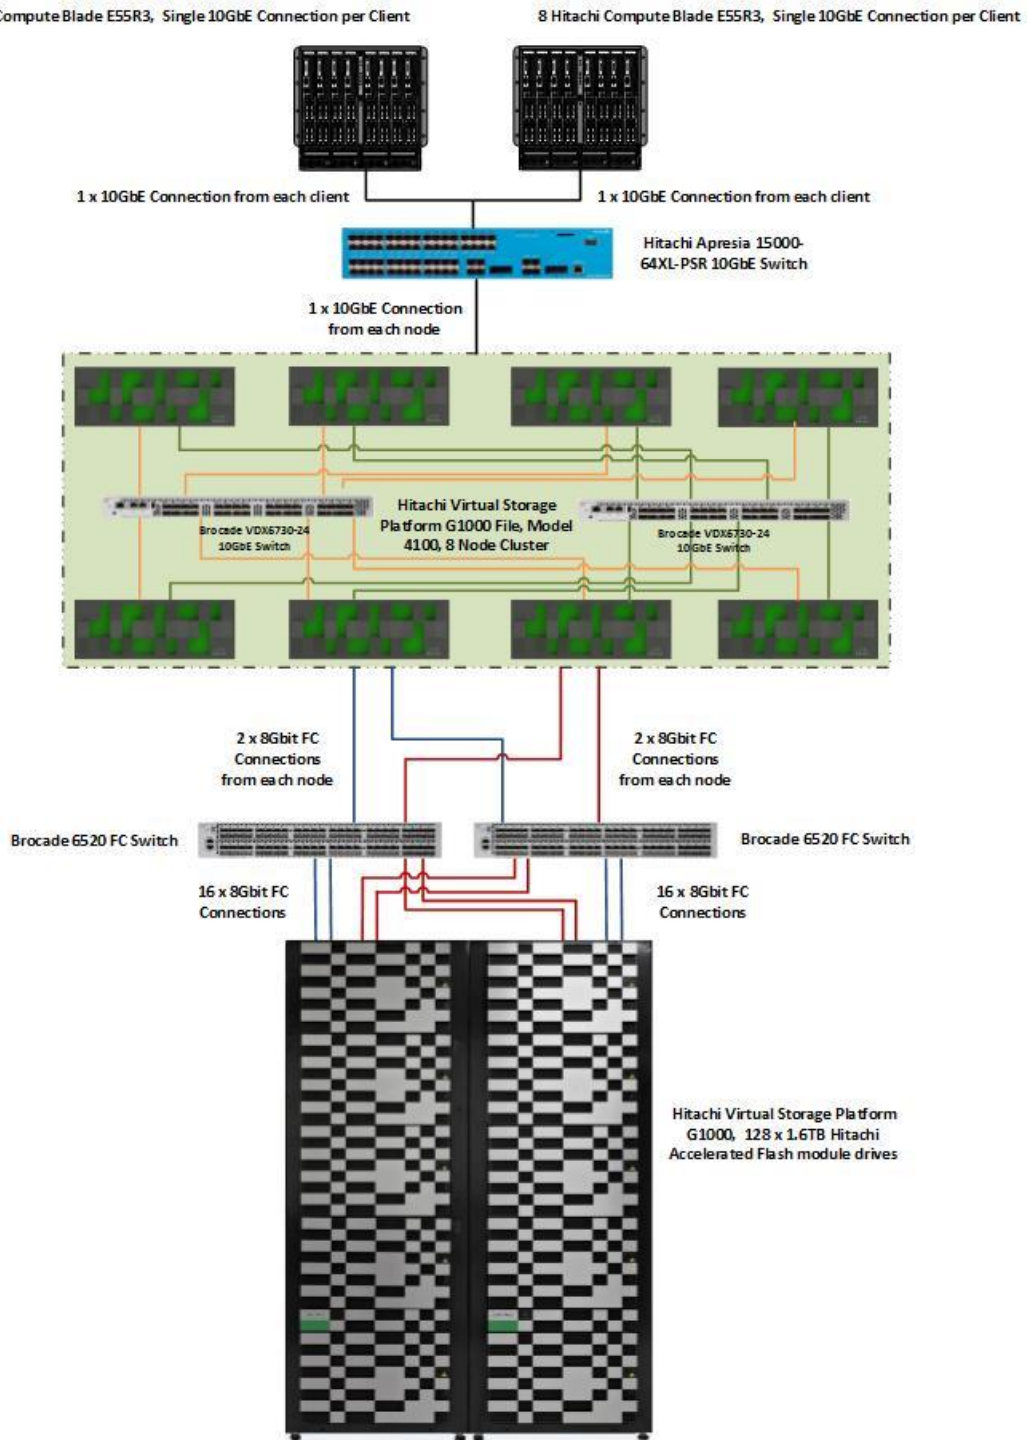
## Detailed Component Summary and High-Level Overview

Table 1 provides a detailed component summary of the test environment, while Figure 2 provides a high-level overview.

**Table 1. Test Environment Overview**

Vendor	Model	Description	Version	Quantity
Hitachi Data Systems	Hitachi NAS Platform, Model 4100	Up to four 8Gb/sec Fibre Channel ports and four 10GbE Ethernet ports per node	v11.3.3434.01	8
Hitachi Data Systems	Hitachi Virtual Storage Platform G1000 All Flash array	Storage system equipped with 1TB of cache memory, thirty-two 8Gb/sec Fibre Channel ports, 128 x 1.6TB Hitachi Accelerated Flash storage, RAID-5 (7D+1P)	80-01-01-00/00	1
Hitachi Cable	Apresia 15000-64XL-PSR	64-port Ethernet switch with 10GbE ports	v8.19.03	1
Hitachi Data Systems	Hitachi Compute Blade 2000 E55R3	Dual Intel Xeon E5-2690 processors, 8 cores/processor, 2.9GHz, 64GB RAM, Intel X520-SR2 10GbE NIC Card	RHEL 6.4 64-bit	16
Brocade Communication Systems	6520	Fibre Channel switch supporting (96) 2, 4, 8 and 16 Gb/sec Fibre Channel ports	v7.1.1	2
Brocade Communication	VDX6730-24	24-port Ethernet switch with 10GbE ports	v4.0.1	2

Figure 2. High-Level Configuration Overview



## Appendix A: Contributors

The information included in this document by Chayan Sarkar represents the expertise, feedback and suggestions of a number of skilled practitioners. The author would like to recognize and sincerely thank the following contributor and reviewer of this document.

- Fred Oh
- Gokula Rangarajan



 **Hitachi Data Systems**

**Corporate Headquarters**  
2845 Lafayette Street  
Santa Clara, CA 96050-2639 USA  
[www.HDS.com](http://www.HDS.com) [community.HDS.com](http://community.HDS.com)

**Regional Contact Information**  
**Americas:** +1 408 970 1000 or [info@hds.com](mailto:info@hds.com)  
**Europe, Middle East and Africa:** +44 (0) 1753 618000 or [info.emea@hds.com](mailto:info.emea@hds.com)  
**Asia Pacific:** +852 3189 7900 or [hds.marketing.apac@hds.com](mailto:hds.marketing.apac@hds.com)

© Hitachi Data Systems Corporation 2014. All rights reserved. HITACHI is a trademark or registered trademark of Hitachi, Ltd. All other trademarks, service marks, and company names are properties of their respective owners. Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.