

# Hitachi NAS Platform 3090 Two- node Cluster Using Hitachi Virtual Storage Platform: SPECsfs 2008 Performance Analysis

Performance Evaluation Using Network File System (NFS)  
v3 Protocol

*By File Services Competency Center, Technical Operations,  
Hitachi Data Systems*

August 2011

# Table of Contents

<b>Executive Summary</b>	<b>3</b>
<b>Test Setup and Methodology</b>	<b>4</b>
<b>Performance Summary</b>	<b>4</b>
<b>About the SPEC SFS2008 Benchmark</b>	<b>5</b>
<b>Detailed Component Summary of the Test Environment</b>	<b>6</b>
<b>Appendix A: Contributors</b>	<b>8</b>

---

## Executive Summary

Unstructured data, such as files and objects, is the fastest growing data type in the enterprise today. This has become a tremendous data management challenge for IT managers, as past solutions have created silos of data and nonintegrated management interfaces. It is a chaotic scenario that only consolidation and data migration can cleanly solve. Conventional NAS systems are not optimized for performance-oriented file system processing and disk access to help with consolidation.

Hitachi NAS Platform models 3080-G2 and 3090-G2, powered by BlueArc®, offer best-in-class enterprise midrange performance and scalability. They provide native intelligent file tiering, clustering up to 4 nodes, single namespace, large 256TB volumes, data migrator software and integration with high-quality Hitachi storage systems and management products. The hardware-based Hitachi NAS Platform (HNAS) architecture leverages field programmable gate arrays (FPGAs) to perform operations in parallel. This facilitates high throughput between servers and Hitachi storage systems. HNAS provides the capabilities and feature set to help IT managers consolidate many, if not all, of their server-based filers and NAS appliances into just a few Hitachi systems. This leads to not only immediate capital expenditure (CAPEX) and operating expense (OPEX) savings, but also power, cooling and space savings in the data center. With HNAS, organizations are assured of saving up to 40% in storage costs, and up to 65% in backup costs.

Hitachi Data Systems tested Hitachi NAS Platform 3090-G2 using storage from a Hitachi Virtual Storage Platform (VSP). Industry-leading results were experienced for its 2-node cluster configuration on the SPEC SFS2008\_nfs.v3 tests for network storage server performance. These results were generated utilizing a truly market-ready system configuration.

Hitachi NAS Platform 3090-G2 2-node cluster delivered an overall performance of **148,439 SPEC SFS2008\_nfs.v3** operations per second, with an overall response time of **2.89** msec. (Visit <http://www.spec.org/sfs2008/results/res2011q2/sfs2008-20110613-00195.html> for detailed results). This result outperforms competing solutions in the market while using fewer disk drives for greater spindle efficiency. SPEC SFS2008 test results show that HNAS 3090-G2 using storage from a Hitachi Virtual Storage Platform can offer optimized, high-performance NAS services for NFS v3 environments.

## Test Setup and Methodology

The test bed consisted of multiple components, including a Hitachi NAS Platform 3090-G2 2-node cluster, 2 Brocade 5320 Fibre Channel Switches, a Hitachi Virtual Storage Platform, 16 NFS clients and a Brocade Turbolron 24x Ethernet Switch. VSP was configured with 64GB cache memory and 320 x 146GB 15K RPM SAS disk drives. Eighty LUNs were created using RAID-5 (3D+1P). The LUNs were distributed across the sixteen 4Gb/sec Fibre Channel front-end ports. Each LUN was presented to 2 front-end ports, thereby ensuring multipathing and failover capabilities.

The HNAS 3090-G2 cluster was connected to VSP via a redundant pair of Brocade 5320 Fibre Channel Switches. A zone was created on each Fibre Channel switch. Each of the HNAS 3090-G2 nodes were connected to each zone via two 4Gb/sec Fibre Channel ports. VSP was connected to each zone via 8 x 4Gb/sec Fibre Channel ports, providing the I/O paths from the server to the storage.

Eight storage pools were created on the HNAS 3090-G2 cluster, with 10 LUNs 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. In total, there were 8 file systems and 8 NFS exports on the NAS Platform 3090-G2 cluster.

Only one 10Gbit NIC interface integrated into the HNAS 3090-G2 node was used per node for the tests. Sixteen Oracle Sun Fire x2200 M2 servers were used as load-generating clients for driving the client workloads. The assignment of processes to file systems was done in such a way that they were uniformly divided across all the file systems.

For network connectivity among the clients and the HNAS 3090-G2 2-node cluster, a Brocade Turbolron 24x Switch was used, all in a 10Gb/sec network environment. The standard Ethernet MTU size (1500 bytes) was used for the tests. Each client had an Intel XF SR10 gigabit Ethernet (GbE) single port PCIe network interface.

The Brocade Turbolron 24X Switch is a 24-port 10GbE solution with a compact, high-performance, high-availability and high-density design that meets mission-critical data center and high-performance computing (HPC) requirements. With an ultra-low-latency, cut-through, non-blocking architecture, the Brocade Turbolron 24X provides a cost-effective connectivity solution for Hitachi NAS Platform clustering, helping demanding applications reduce response time and increase application availability. The Turbolron 24X is designed to save valuable rack space, power and cooling in the data center while delivering 24/7 service through its high-availability design.

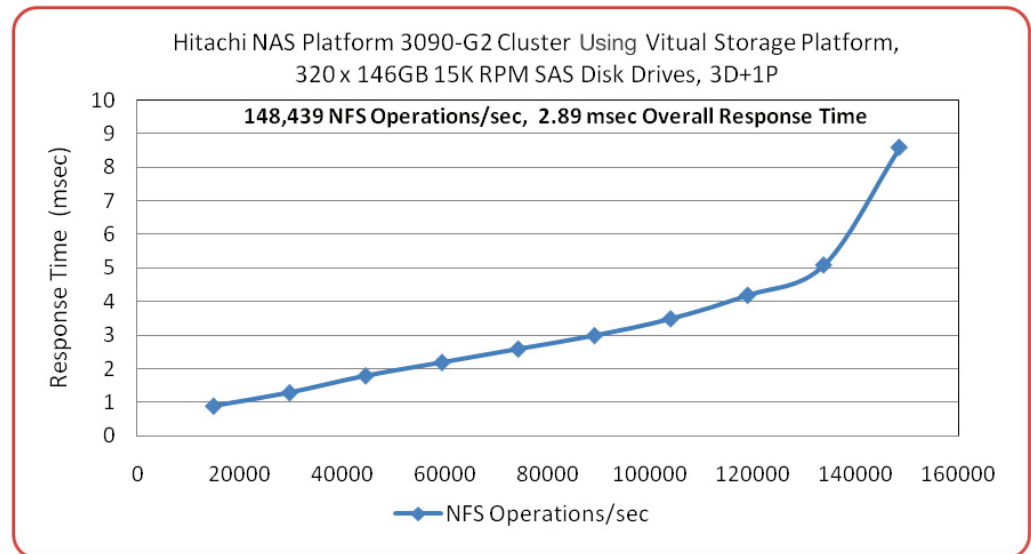
The Brocade 5300 Switch is designed to consolidate connectivity in rapidly growing mission-critical storage environments. It delivers full 8Gb/sec performance in configurations of 48, 64 or 80 ports in an efficiently designed 2U package. The combination of density and performance enables a highly scalable storage pool for Hitachi NAS clusters.

## Performance Summary

The tests focused on measuring the NFS performance of the Hitachi NAS Platform 3090-G2 2-node cluster (see Figures 1 and 2). NFS is one of the most commonly used file systems in NAS solutions, and the standard file sharing mechanism used in UNIX and Linux environments. Sixteen load-

generating clients were used for the tests, each driving 72 threads, with 1152 threads active overall on the HNAS 3090-G2 cluster at a time.

**Figure 1. SPEC SFS2008 NFS Performance Summary**



## About the SPEC SFS2008 Benchmark

SPEC SFS2008 is the latest 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. SPEC SFS2008 results summarize the server's capabilities with respect to the number of operations that can be handled per second, as well as the overall latency of the operations. The suite is 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/>.

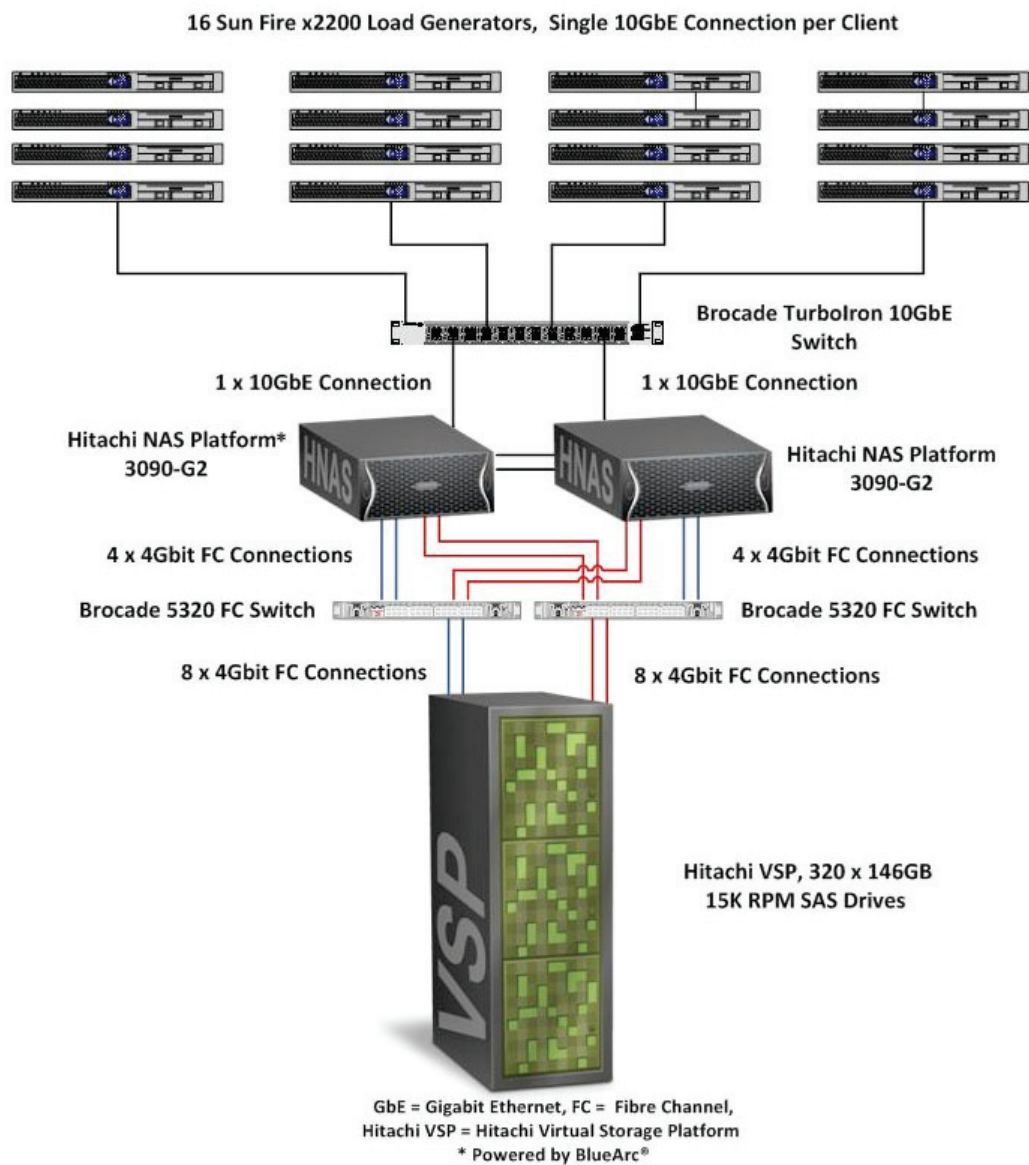
## Detailed Component Summary of the Test Environment

The detailed component summary of the Hitachi NAS Platform 3090-G2 2-node cluster test environment for the SPEC SFS2008 test is presented in Table 1.

**TABLE 1. TEST ENVIRONMENT OVERVIEW**

Vendor	Model	Description	Version	Quantity
<b>Hitachi Data Systems</b>	Hitachi NAS Platform 3090-G2 (HNAS 3090-G2)	Four 4Gbit/sec Fibre Channel front-end ports, two 10Gbit/sec Ethernet paths	8.1.2312.07	2
<b>Hitachi Data Systems</b>	Hitachi Virtual Storage Platform	Enterprise-level storage system configured in dual chassis equipped with 256GB of cache (only 64GB used for tests), 2 front-end director features, sixteen 8Gb/sec Fibre Channel ports (set to 4Gb/sec for tests), 4 back-end director features, 4 grid switch features, 4 virtual storage director features, 320 x 146GB 15K RPM SAS disk drives, RAID-5 (3D+1P), 80 RAID groups, 80 LDEVs	70-02-02-00/00	1
<b>Oracle</b>	Sun Fire x2200 M2 Server	Dual AMD Opteron processor 2218 HE 2.6GHz, 8GB RAM, RHEL5 64-bit, Intel XF SR 10Gb/sec NIC card		16
<b>Brocade Communication Systems</b>	5320 SAN Switch	Fibre Channel switch supporting (80) 1, 2, 4 and 8Gb/sec Fibre Channel ports (port speed hardcoded to 4Gb/sec for this test)	v6.4.1a	1
<b>Brocade Communication Systems</b>	Turbolron 24X	24-port Ethernet switch with 10GbE ports	v4.2.0	1

Figure 2. High-level Configuration Overview



## Appendix A: Contributors

The information included in this document represents the expertise, feedback and suggestions of a number of skilled practitioners. The authors would like to recognize and sincerely thank the following contributors and reviews of this document (listed alphabetically, by first name):

- Gokula Rangarajan
- Jeffrey Blomberg
- Rafnas AK
- Shekhar Berry

## Hitachi Data Systems Corporation

---

### Corporate Headquarters

750 Central Expressway  
Santa Clara, California 95050-2627 USA  
[www.HDS.com](http://www.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 is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names in this document or website 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.

© Hitachi Data Systems Corporation 2011. All Rights Reserved. WP-408-A DG August 2011