PERFORMANCE BRIEF

HITACHI
Inspire the Next

Hitachi NAS Platform 3080 Cluster Using the Hitachi Adaptable Modular Storage 2500: SPEC SFS2008 Performance Analysis

Performance Evaluation Using Network File System (NFS) v3 Protocol

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

July 2010

Table of Contents

Executive Summary

Unstructured data such as files and objects are the fastest growing data types in the enterprises 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.

The Hitachi NAS Platform 3080 and 3090, powered by BlueArc®, offer best-in-class performance and scalability, native intelligent file tiering, clustering up to four nodes, single namespace, large 256TB volumes, data migrator software and integration with Hitachi high-quality block storage and management products. The hardware-based Hitachi NAS Platform architecture leverages Field Programmable Gate Arrays to perform operations in parallel. This facilitates high throughput between servers and Hitachi storage systems. The Hitachi NAS Platform 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 power, cooling and space savings in the data center as well. With Hitachi NAS Platform organizations are assured of saving up to 40% in storage costs and up to 65% in backup costs.

Hitachi Data Systems tested the Hitachi NAS Platform 3080 server using storage from a Hitachi Adaptable Modular Storage 2500. Industry leading results were experienced for its two node cluster configuration on the SPEC SFS2008_nfs.v3 tests for network storage server performance. These results were generated utilizing a true market-ready system configuration.

Hitachi NAS Platform 3080 dual-node cluster delivered performance of 79,058 SPEC SFS2008_nfs.v3 operations per second, with an Overall Response Time of 3.29 msec. This result outperforms competing solutions in the market while using fewer disk drives for greater spindle efficiency. SPEC SFS2008 test results show that the Hitachi NAS Platform 3080 dual-node cluster using Adaptable Modular Storage 2500 can offer optimized, high-performance NAS services for NFS v3 environments.

Performance Summary

The tests focused on measuring NFS performance of the Hitachi NAS Platform 3080 as a dual-node cluster (see Figures 1 and 2). NFS is one of the most commonly used file systems in NAS solutions, and it is the standard file sharing mechanism used in UNIX and Linux environments. Twenty load generating clients were used for the tests, each driving 32 threads with 640 threads active overall on the Hitachi NAS Platform 3080 cluster configuration.

Figure 1. SPEC SFS2008 NFS Performance Summary: Hitachi NAS Platform 3080 Cluster and Hitachi Adaptable Modular Storage 2500 with 160 x 450GB SAS Hard Disk Drives

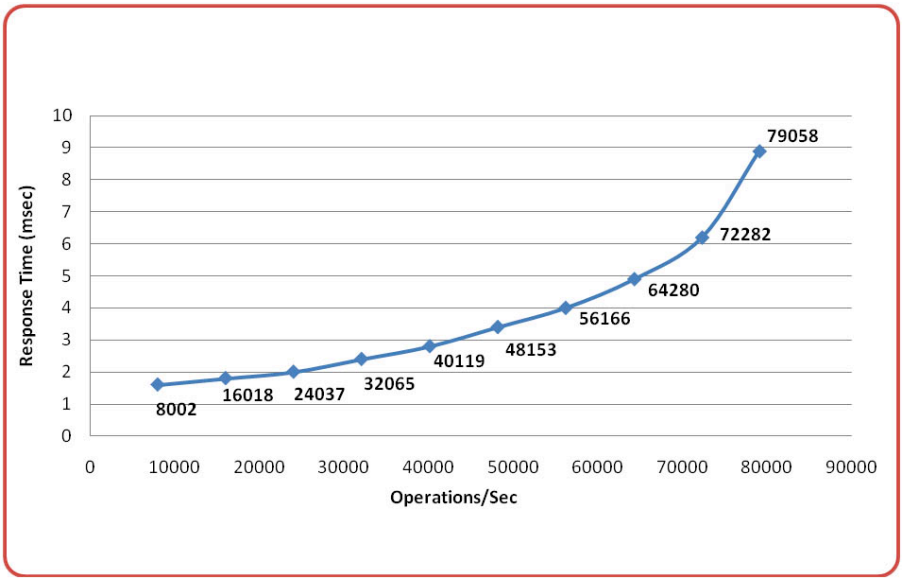

Figure 2. High-level Configuration Overview

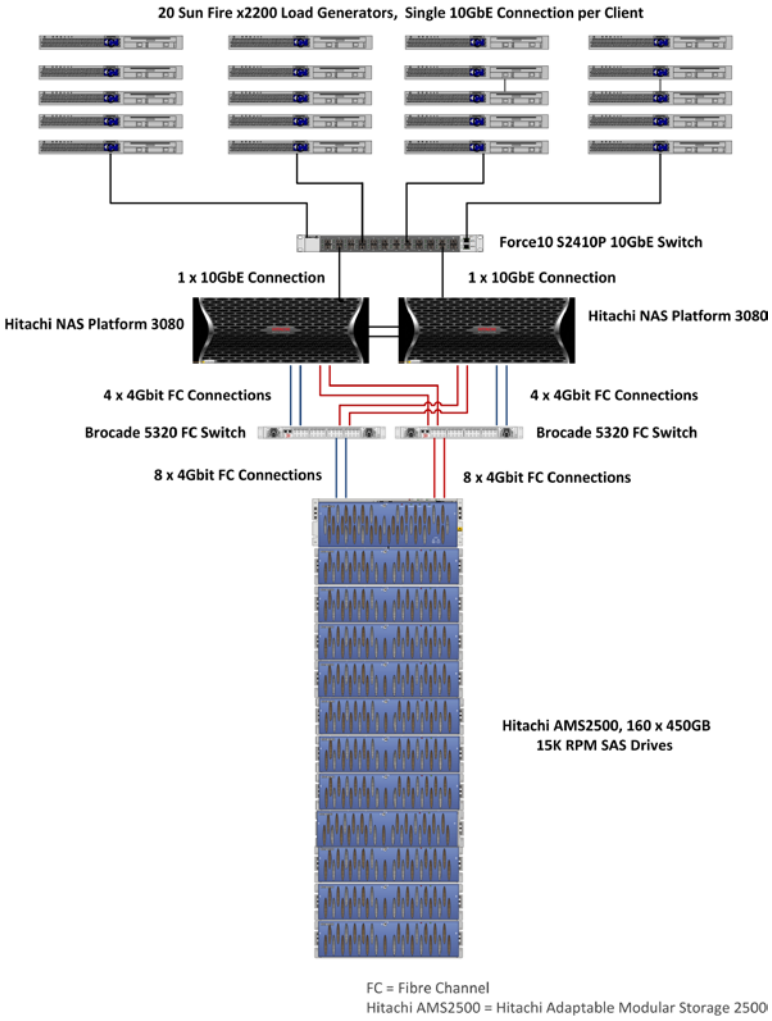

Test Setup and Methodology

The test bed consisted of multiple components, including a Hitachi NAS Platform 3080 dual-node cluster, two Brocade 5320 Fibre Channel Switches, a Hitachi Adaptable Modular Storage 2500 storage system, 16 clients and a Force10 Ethernet switch. The Hitachi Adaptable Modular Storage 2500 storage system was equipped with 32GB cache memory and (160) 450GB 15K RPM SAS disk drives. There were 32 LUNs created using RAID-5 (4D+1P) and formatted using the optional 64KB stripe size. These LUNs were distributed across the the sixteen 4Gb/sec Fibre Channel front-end ports. Each LUN was presented to two front-end ports, thereby ensuring multipathing.

The Hitachi NAS Platform 3080 cluster was connected to the Adaptable Modular Storage 2500 via a redundant pair of Brocade 5320 Fibre Channel switches. A zone was created on each Fibre Channel switch. Port 1 and 3 of both Hitachi NAS Platform nodes were connected to the first switch and Port 2 and 4 were connected to the second switch. The Adaptable Modular Storage 2500 was connected to each zone via 8 x 4Gb/sec Fibre Channel ports providing the I/O paths from the server

to the storage. Sixteen storage pools were created on the Hitachi NAS Platform 3080 cluster, with two LUNs assigned to each storage pool. One file system was created on each storage pool using 4KB file system block size. In total, there were 16 file systems on the Hitachi NAS Platform 3080 cluster. The 16 file systems were aggregated into a single namespace using Cluster Namespace (CNS) feature. Eight file systems were assigned as the primary file system to each Hitachi NAS Platform 3080 node, although both nodes can access all the 16 file systems through the active-active cluster link, or directly via the Fibre Channel SAN in the event of a node failure. One NFS export was created on the Hitachi NAS Platform cluster.

With this configuration, clients were able to mount the file shares using the NFS protocol. One 10Gbit NIC interface integrated into each Hitachi NAS Platform 3080 node was used for the test, making two Ethernet ports in total for client access. Sixteen Oracle SunFire x2200 M2 servers used as load generating clients for driving the client workloads. Each load generating client hosted 32 processes. 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 Hitachi NAS Platform 3080 cluster, a Force10 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 SR10GbE single port PCIe network interface, which connected to the ports on the Force10 S2410 network switch.

Detailed Component Summary of the Test Environment

The detailed component summary of the Hitachi NAS Platform 3080 cluster test environment for the SPEC SFS2008 test is given in Table 1.

TABLE 1. TEST ENVIRONMENT OVERVIEW

Vendor	Model	Description	Version	Quantity
Hitachi Data Systems	**Hitachi NAS Platform (HNAS 3080)**	Network attached storage (NAS) platform with six 1GbE ports, two 10GbE ports, two 10GbE cluster interconnect ports, four 4Gb/sec Fibre Channel ports and ~32GB distributed memory	6.5.1849	2
Hitachi Data Systems	**Hitachi Adaptable Modular Storage 2500 (AMS2500)**	Medium- to enterprise-level storage system equipped with 32GB cache memory, sixteen 4Gb/sec Fibre Channel ports, 160 x 450GB, 15K RPM SAS disk drives, in RAID-5 (32 x 4D+1P LUNs)	0890/B	1
Oracle	**SunFire x2200 M2 Servers**	Dual AMD Opteron processor 2218 HE 2.6GHz, 2GB RAM, Red Hat Enterprise Linux Server Release 5 64-bit edition (Kernel 2.6.18-8.e15 SMP), Intel XF SR 10Gb/sec NIC card		20
Brocade	**5320**	80 port Fibre Channel switch supporting 1/2/4/8Gb/sec link speeds	6.3.0b	2
Force10	**S2410P**	24 port Ethernet switch with 10GbE ports	2.4.1.11	1

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, providing 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/.

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):

- Fred Oh
- Gokula Rangarajan
- John Dupuis
- Rafnas AK
- Shekhar Berry

Hitachi Data Systems Corporation

Corporate Headquarters
750 Central Expressway
Santa Clara, California 95050-2627 USA
www.hds.com

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