TECHNICAL PAPER

Migration of 250 Virtual Machines from Block Storage to Hitachi NAS Platform with Server Workload Performance Comparison

Tech Note

By Daniel Worden and Jeff Chen

March 9, 2015

Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@hds.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Table of Contents

Migration of 250 Virtual Machines from Block Storage to Hitachi NAS Platform with Server Workload Performance Comparison

Tech Note

Hitachi NAS Platform (HNAS) for VMware vSphere is positioned and deployed as an enterprise NFS solution that delivers a level of scalable and predictable performance (IOPS and latency), capacity, efficiency, and data protection to meet the needs of VMware cloud environments in addition to other benefits. This paper is meant to address the following areas:

- Provide recommended configurations for certain virtual machine (VM) counts and expected performance (IOPS and latency) using current best practice deployments. During this testing, 250 virtual machines per cluster were tested.
- Test the life cycle of virtual machines at scale through. The virtual machine life cycle includes:
  - Starting out on block storage running a variety of applications that achieved their target IOPS with acceptable latency.
  - The virtual machines were migrated to HNAS NFS datastores.
  - The virtual machines ran a variety of applications that achieved their target IOPS with acceptable latency.
- Show the HNAS Platform's ability to migrate virtual machines with VMDKs of 1TB and greater.

This document shows the following benefits of Hitachi NAS Platform during this testing:

- Superflush reduced the workload on the backend block storage array, allowing more resources for other workloads on Hitachi Unified Storage VM. Superflush is described in more detail later in this paper.
- All workloads were able to maintain the target I/O during svMotion operations.
- Application latency was lower during HNAS testing compared to block testing.

Table 1. Verified Test Cases

Test Case	*Pass/Fail Criteria*	*Result*
Hitachi Unified Storage VM Fibre Channel storage baseline.	Virtual machines should be able to reach and maintain the I/O target configured for them.	Pass
Virtual Machine migration from Fibre Channel storage to HNAS Platform NFS datastores	All virtual machines will be successfully migrated and the workloads will successfully run after the migration, while other tiles are being migrated.	Pass
Run workload on all virtual machines on HNAS NFS Datastores for 12 hours	Virtual machines should be able to reach and maintain the I/O target configured for them.	Pass
Migrate a virtual machine with a 1 TB+ VMDK from the Fibre Channel storage to the HNAS Platform.	The virtual machine should be successfully migrated from its current storage to the HNAS platform.	Pass

The following tests were run for this phase of the project:

- Test execution against 250 virtual machines using block storage from an Hitachi Unified Storage VM for a 12-hour test to establish a baseline.
- Migrate the 250 virtual machines from the block storage to HNAS NFS datastores. The 250 virtual machines were migrated in four phases:
  - The workloads were stopped on 1 tile (25 virtual machines), and migrated to the Hitachi NAS Platform NFS datastores. A tile is defined in Table 1.
  - The workloads were restarted on the first tile, and with the workloads stopped on tiles 2-4 they were migrated to the Hitachi NAS Platform NFS datastores.
  - The workloads were restarted on tiles 2-4, and with the workloads stopped on tiles 5-7 they were migrated to the Hitachi NAS Platform NFS datastores.
  - The workloads were restarted on tiles 5-7, and with the workloads stopped on tiles 8-10 they were migrated to the Hitachi NAS Platform NFS datastores.
- After all 10 tiles (250 virtual machines) were migrated to the Hitachi NAS Platform NFS datastores, a 12-hour test was run for validation of the workload and configuration, as well as comparison of the performance of Hitachi NAS Platform compared to the block baseline.

This document provides the following:

- Validation of 250 virtual machines migrating using Storage vMotion (SvMotion) from block to Hitachi NAS Platform
- Validation of 250 virtual machines running a variety of workloads on Hitachi NAS Platform
- High-level overview of this solution's compute, network, and storage configurations

Note — Testing of this configuration was in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow the recommended practice of conducting proof-of-concept testing for acceptable results in a non-production, isolated test environment that otherwise matches your production environment before your production implementation of this solution.

Use Case Overview

During this phase of testing 250 virtual machines were used, and they ran a variety of workload profiles. The workloads were divided into tiles of 25 virtual machines each. Each workload profile had a light, medium, and heavy component. The virtual machines in each tile were configured for an average of 25 IOPS. The workload profile within each tile is described in Table 2.

Table 2. Per Tile Virtual Machine Workload Profile

Workload	Weight	IOPS	VMDK Size	Number of Virtual Machines	Total IOPS	Total VMDK Size
SQL Server	Light	10	60 GB	3	30	180 GB
Web Server	Light	10	50 GB	6	60	300 GB
Exchange	Light	10	60 GB	3	30	180 GB
OLTP	Light	10	60 GB	3	30	180 GB
SQL Server	Medium	25	200 GB	1	25	200 GB
Web Server	Medium	25	80 GB	2	50	160 GB
Exchange	Medium	25	100 GB	1	25	100 GB
OLTP	Medium	25	100 GB	1	25	100 GB
SQL Server	Heavy	70	40 GB	1	70	40 GB
Web Server	Heavy	70	40 GB	2	140	80 GB
Exchange	Heavy	70	40 GB	1	70	40 GB
OLTP	Heavy	70	40 GB	1	70	40 GB
Total		25 per virtual machine		25	625	1600 GB

The total VMDK size includes a 20 GB VMDK for the operating system. The remainder of the VMDK size was used as a VMDK for a data volume. Each virtual machine had the following resources assigned to it:

- 1 vCPU
- 4 GB of RAM

Table 3 shows the I/O profile for each workload.

Table 3. Workload I/O Profiles

Workload	I/O Size	Percent Random	Percent Read
Microsoft® SQL Server®	64 KB	100%	66%
Web Server	8 KB	75%	95%
Microsoft Exchange	8 KB	80%	55%
OLTP	8 KB	100%	70%

Tested Components

Table 4 lists the specific hardware components used during testing.

Table 4. Tested Hardware Components

Hardware	Description	Version	Quantity
Hitachi Unified Storage VM	▪ Dual controllers ▪ 16 × 8 Gb/sec Fibre Channel ports ▪ 23 GB cache memory ▪ 208 × 600 GB 10k RPM SAS disks, 2.5 inch SFF ▪ 4 × FMD	73-03-06-00/00	1
Hitachi NAS Platform 4100	▪ 2 × 10 Gb/sec Cluster ports ▪ 4 × 10 Gb/sec Ethernet ports ▪ 4 × 8 Gb/sec Fibre Channel ports	12.2.3719.02	1
Hitachi Compute Blade 500 Chassis	▪ 8-blade chassis ▪ 2 Brocade 5460 Fibre Channel switch modules, each with 6 × 8 Gb/sec uplink ports ▪ 2 Brocade VDX 6746 Ethernet switch modules, each with 8 × 10 Gb/sec uplink ports ▪ 2 management modules ▪ 6 cooling fan modules ▪ 4 power supply modules	SVP: A0170-D-8920 5460: FOS 7.0.2C VDX6746: NOS 4.1.2	1
520HB2 server blade	▪ Half blade ▪ 2 × 12-core Intel Xeon E5-2697 processors, 2.70 GHz ▪ 256 GB RAM ▪ 16 × 16 GB DIMMs	BMC/EFI: 01-27	2
Brocade 6510	▪ SAN switch 48 × 8 Gb Fibre Channel ports	FOS 7.0.1a	2
Brocade VDX 6740	▪ Ethernet switch with 40 × 10 Gb/sec ports	NOS 4.1.2	2

Table 5 lists the specific software components used during testing.

Table 5. Tested Software Components

Software	Version
Hitachi Storage Navigator Modular 2	Microcode Dependent
VMware vCenter server	5.5.0 U2
VMware Virtual Infrastructure client	5.5.0 U2
VMware ESXi	5.5.0 U2
VDbench	5.04
Microsoft Windows Server 2008 R2 (Microsoft SQL Server virtual machine operating system)	
Microsoft Windows Server 2012 R2 (Microsoft Exchange server virtual machine operating system)	
SUSE Linux Enterprise Server 11 SP2 (OLTP and web server virtual machine operating system)	

High Level Test Infrastructure

Figure 1 illustrates the Hitachi Unified Storage VM storage configuration for the Fibre Channel baseline testing.

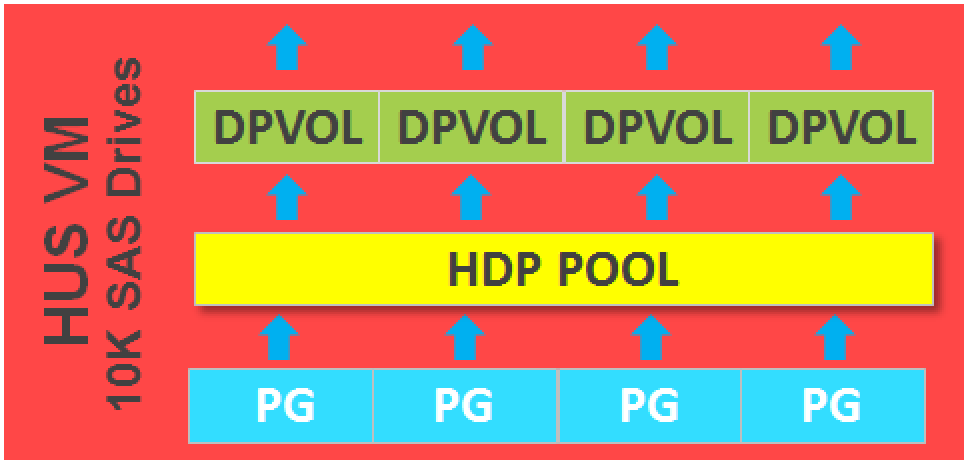

Figure 1

Hitachi Unified Storage VM was configured with 10K SAS drives grouped and presented as:

- 13 × RAID-6 (6D+2P) parity groups (PG)
- 1 dynamic provisioning pool was created from the 13 parity groups
- 4 DP-VOLs were created from the single dynamic provisioning pool and presented as datastores

Figure 2 illustrates the storage configuration for the Hitachi NAS Platform NFS storage testing.

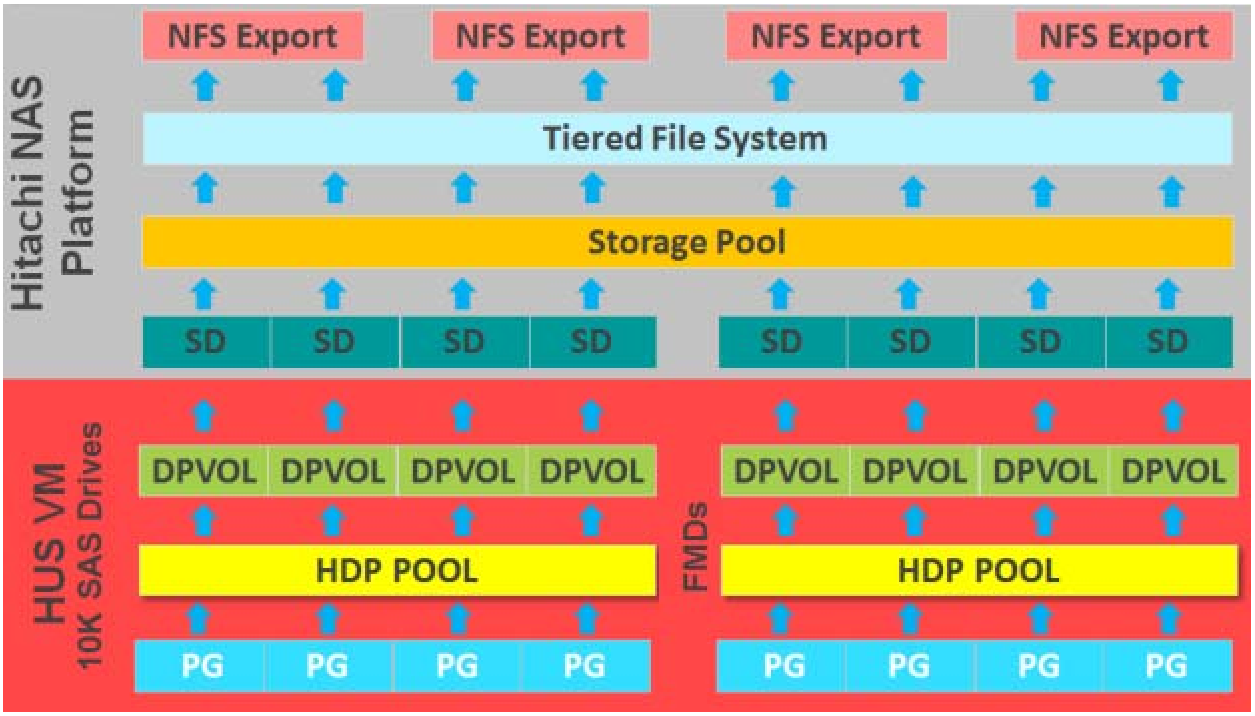

Figure 2

Hitachi Unified Storage VM was configured with 10K SAS drives grouped and presented as the following:

- 13 RAID-6 (6D+2P) parity groups (PG)
- 1 dynamic provisioning pool was created from the 13 parity groups
- 16 × 8 TB thin provisioned DP-VOLs were created from the single dynamic provisioning pool

Hitachi Unified Storage VM was configured with FMDs grouped and presented as the following:

- 1 RAID-5 (3D+1P) parity group (PG)
- 1 dynamic provisioning pool was created from the single parity group
- 8 × 4 TB thin provisioned DP-VOLs were created from the single dynamic provisioning pool

The following configuration was used for Hitachi NAS Platform:

- 16 system drives (SD) were created from the 16 DP VOLs that were provisioned from the 10K SAS drives
- 8 system drives (SD) were created from the 8 DP VOLs from the FMDs
- One Storage Pool with a Tiered Filesystem (TFS) was created from the 24 System Drives as listed below. This Storage Pool creation pattern takes future filesystem expansion into consideration. The filesystem should be expanded with eight SDs at a time.
  - Stripeset 0 (Tier 0 - metadata) 8 × FMD 4 TB SDs
  - Stripeset 1 (Tier 1) 8 × 10K SAS 8 TB SDs
  - Stripeset 2 (Tier 1) 8 × 10K SAS 8 TB SDs
- One filesystem was created from the Storage Pool
- 4 NFS Exports were created as mount points for the ESXi datastore

Thin provisioned DP-VOLs as Hitachi NAS Platform system drives were introduced in version 12.1.3613.10. Since this version, thin provisioned DP-VOLs are recommended for the following reasons:

- It is quicker and easier to expand an Hitachi NAS Platform file system
- Data rebalancing is automatically performed by the dynamic provisioning pool more efficiently
- Prior to the support of thin provisioned SDs, new SDs had to be presented to Hitachi NAS Platform, and the storage pool had to be expanded by a stripeset. Then the file system had to be rebalanced.

When using DP thin provisioned volumes with Hitachi NAS Platform, it is a best practice to not over-provision more than three times the actual available space.

Eight DP-VOLs were created from a parity group of FMD because FMDs are fast enough to handle the additional I/O threads compared to mechanical spindle-based disks.

Superflush was configured on all Hitachi NAS Platform system drives with the setting of 3 wide by 128 Kb (3 × 128). These settings are the best practice when using Superflush with Hitachi Unified Storage VM.

All best practices were followed when configuring Hitachi Unified Storage VM and Hitachi NAS Platform.

For information on Ethernet networking configuration recommendations please see the reference architecture guide Deploy Hitachi Unified Compute Platform Select for VMware vSphere using Hitachi NAS Platform with Hitachi Unified Storage VM.

Test Results

These are the test results.

Hitachi Unified Storage VM Fibre Channel Storage Baseline

The first test run was a baseline on Hitachi Unified Storage VM using the Fibre Channel protocol.

- 10 tiles (250 virtual machines) were distributed evenly across four datastores.
- The 10 tiles were provisioned as Eager Zero Thick VMDKs which used 17 TB of disk space.

The 10 tiles ran for 12 hours to establish a baseline for comparison when the workloads were moved to Hitachi NAS Platform.

Table 6. First Test Case

Test Case	Description	Result
Hitachi Unified Storage VM Fibre Channel storage baseline.	Run workload on 10 tiles or 250 virtual machines for 12 hours	Passed
	Expected result: Virtual machines should be able to reach and maintain the I/O target configured for them.	

During baseline testing, the 10 tiles were able to reach greater than 99% of the target I/O, as shown in Figure 3.

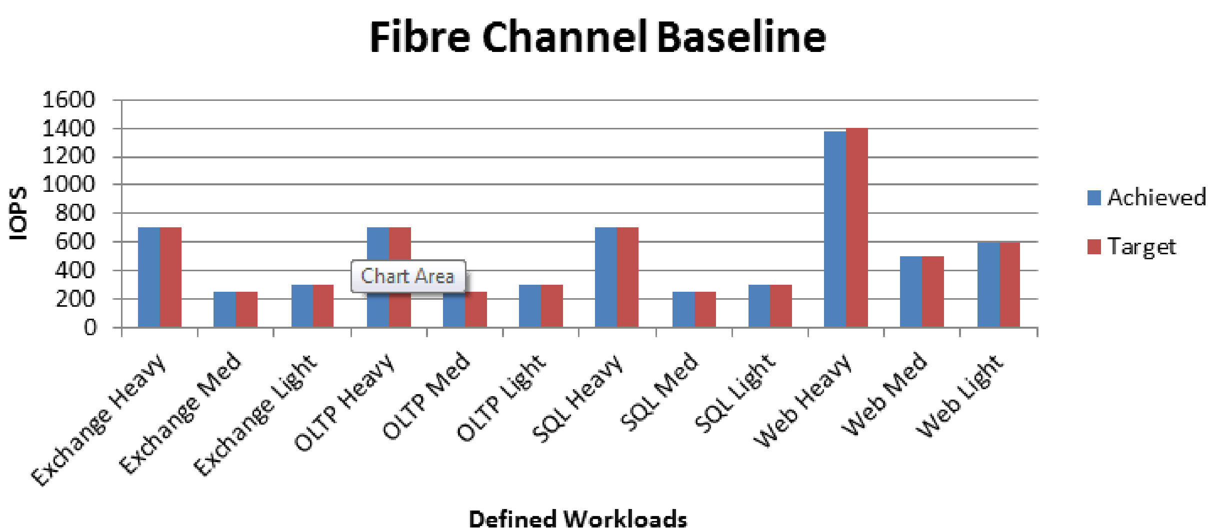

Figure 3

Figure 4 shows the average busy rate for the 13 SAS parity groups during the 12-hour test. During the Fibre Channel baseline testing, the average percent busy for the 13 parity groups was between 77and 80 percent. This indicates that the number of SAS disks were sized to be fully utilized for these workloads.

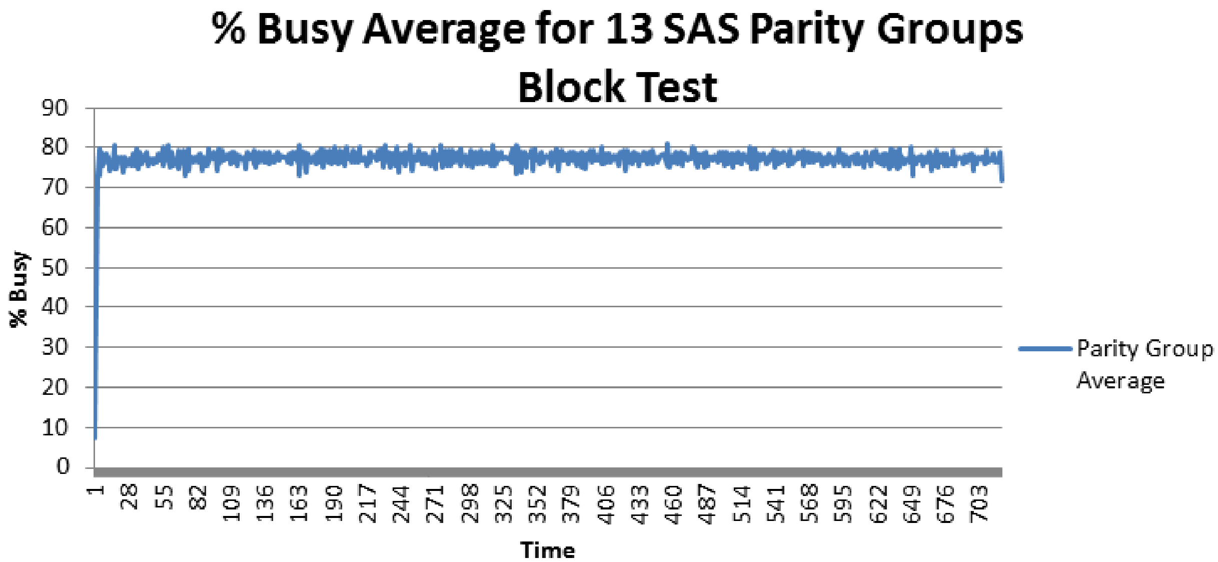

Figure 4

The average percent busy was between 7 and 8 percent for the controllers on Hitachi Unified Storage VM, as shown in Figure 5.

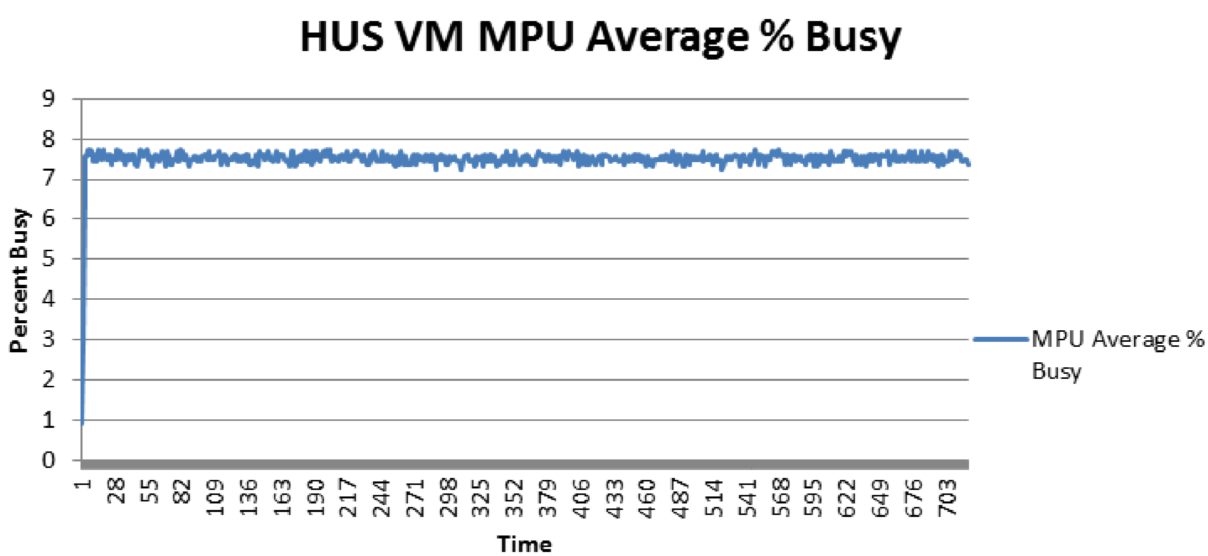

Figure 5

Virtual Machine Migration From Fibre Channel Datastores to NFS Datastores on Hitachi NAS Platform

The next test case was Storage vMotion of the 10 tiles, or 250 virtual machines from the block storage, to Hitachi NAS Platform. Four NFS exports were created on Hitachi NAS Platform and mounted on each of the two ESXi hosts used during testing. All of the workloads were stopped on the tiles during the SvMotion operations.

- Tile 1 was migrated from the Fibre Channel datastore to an NFS datastore.
- When migration completed for the first tile, the workload on tile 1 was restarted and tiles 2-4 were migrated from the Fibre Channel datastores to the NFS datastores.
- When migration completed for tiles 2-4, the workloads were started on them. Tiles 5-7 were then migrated from the Fibre Channel datastores to the NFS datastores
- After tiles 5-7 completed migration, the workloads were started on them, and tiles 8-10 were migrated to the NFS datastores.

Table 7. Second Test Case

Test Case	Description	Result
Virtual machine migration from Fibre Channel datastores to HNAS NFS datastores	Migrate 10 tiles in phases, from Fibre Channel datastores to NFS datastores. Upon the completion of each phase the workload for the migrated tiles was started before the next phase was migrated. Table 8 below shows the tiles that were migrated together and the time to completion.	Passed
	Expected result: All tiles will be successfully migrated and the workloads will successfully run after the migration, while other tiles are being migrated.	

Table 8 shows the time to completion when migrating the tiles during each part of this test.

Table 8. Tile Migration Time to Completion

Tiles Migrated	Tiles Running Workloads	Time to Migration Completion
Tile 1	None	53 Minutes
Tiles 2,3,4	Tile 1	1 Hour 56 Minutes
Tiles 5,6,7	Tiles 1,2,3,4	1 hour 43 Minutes
Tiles 8,9,10	Tiles 1,2,3,4,5,6,7	1 hour 22 Minutes

Note in Table 8 that the time to completion decreased as more tiles were migrated, even though the workloads were increasing. The first migration took 53 minutes. The second migration of three tiles with one tile's workload running took only a little more than twice as long as the single tile.

Figure 6 shows the application latency on 10 tiles running simultaneously, without any migrations, compared to tiles 1-7 running simultaneously while tiles 8-10 are being migrated from Fibre Channel to NFS datastores on Hitachi NAS Platform. As Figure 6 shows, the application latency increased by 1 millisecond on the 7 tile workload with the 3 tiles being migrated. This 1 millisecond increase was acceptable.

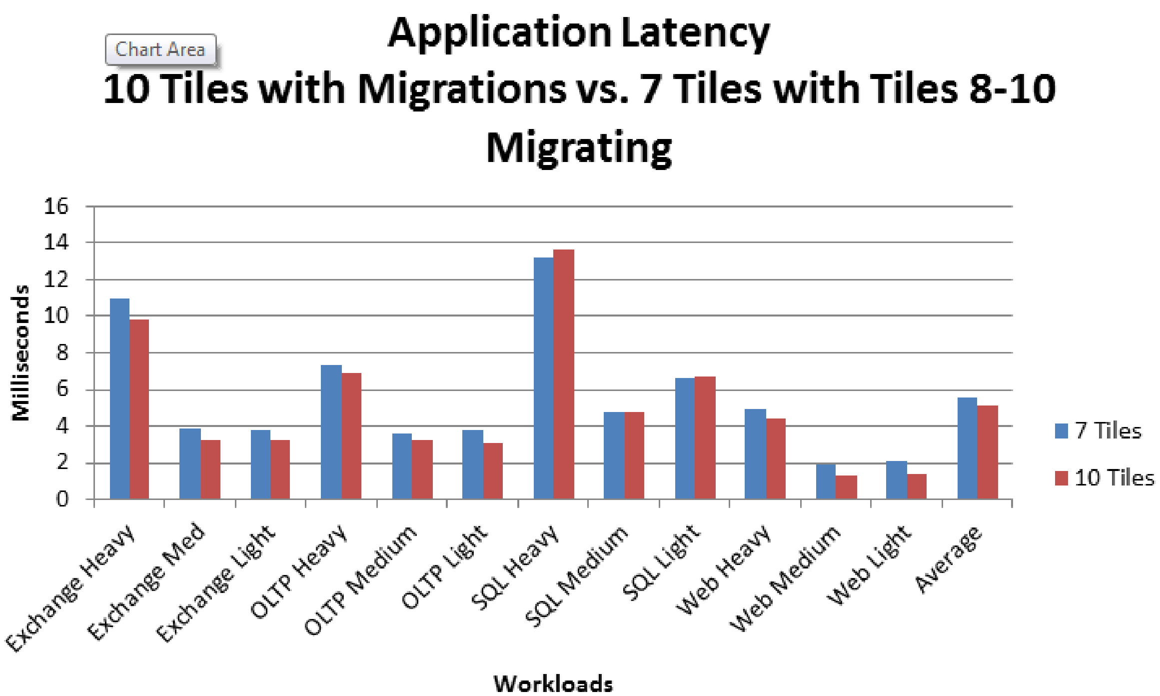

Figure 6

NFS Datastores on Hitachi NAS Platform Baseline

The third test case ran all 10 workloads for 12 hours on NFS datastores on Hitachi NAS Platform.

Table 9. Third Test Case

Test Case	*Description*	*Result*
Run all 10 tiles on NFS Datastores for 12 hours	Run workload on 10 tiles (250 virtual machines) for 12 hours	Passed
	Expected result: Virtual machines should be able to reach and maintain the I/O target configured for them.	

The tiles were distributed across the ESXi host as shown in Table 10.

Table 10. Tile Distribution Across Hosts

Host	*Tiles*
Host1	1,2,3,4,5
Host2	6,7,8,9,10

The tiles were distributed across the four NFS datastores as shown in Table 11.

Table 11. Tile Distriution Across Datastores

NFS Datastore	*Tiles*
NFS_Datastore_001	1,5,9
NFS_Datastore_002	2,6,10
NFS_Datastore_003	3,7
NFS_Datastore_004	4,8

VMDKs were converted from Eager Zeroed Thick to Thin Provisioned as they were migrated from the Fibre Channel datastores to the NFS datastores on Hitachi NAS Platform. Thin provisioned VMDKs are the best practice when using Hitachi NAS Platform with ESXi hosts, and are the default VMDK type when a VMDK is migrated to, or created on, an NFS datastore. Once migrated to Hitachi NAS Platform NFS datastores, the 10 tiles used 5.21 TB of File System space.

During testing of 10 tiles on NFS datastores on Hitachi NAS Platform, 100 percent of the target I/O was achieved, as shown in Table 7.

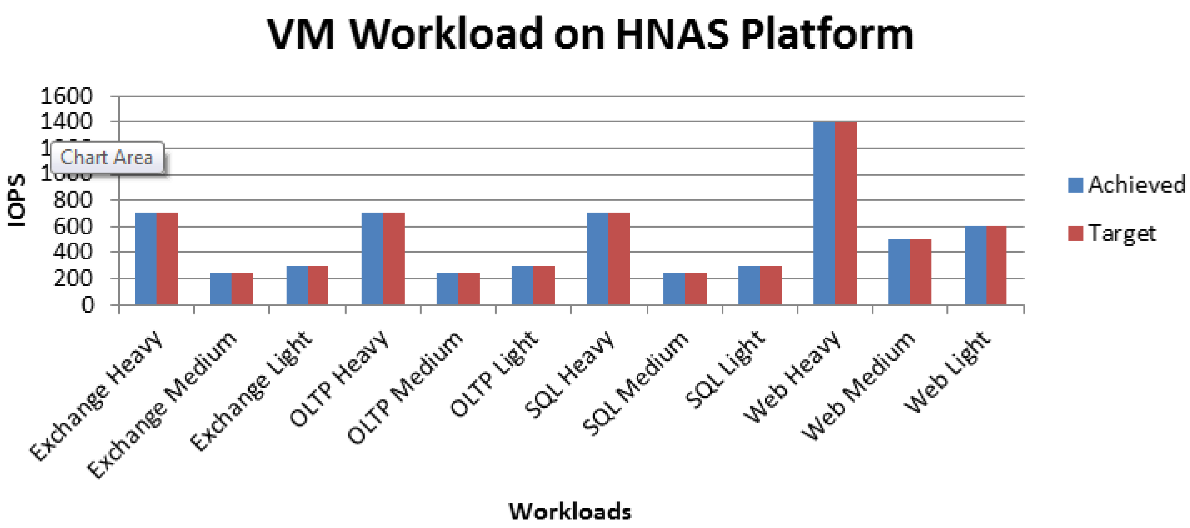

Figure 7

During the migration testing of tiles 8-10, with tile 1-7 workloads running, CPU utilization on Hitachi NAS Platform was 4 percent higher than it was while running 10 tiles workload and no migrations occurring. Table 8 illustrates these findings. Hitachi NAS Platform was working harder to process the workloads and migrations, but the applications were still able to reach their target workloads.

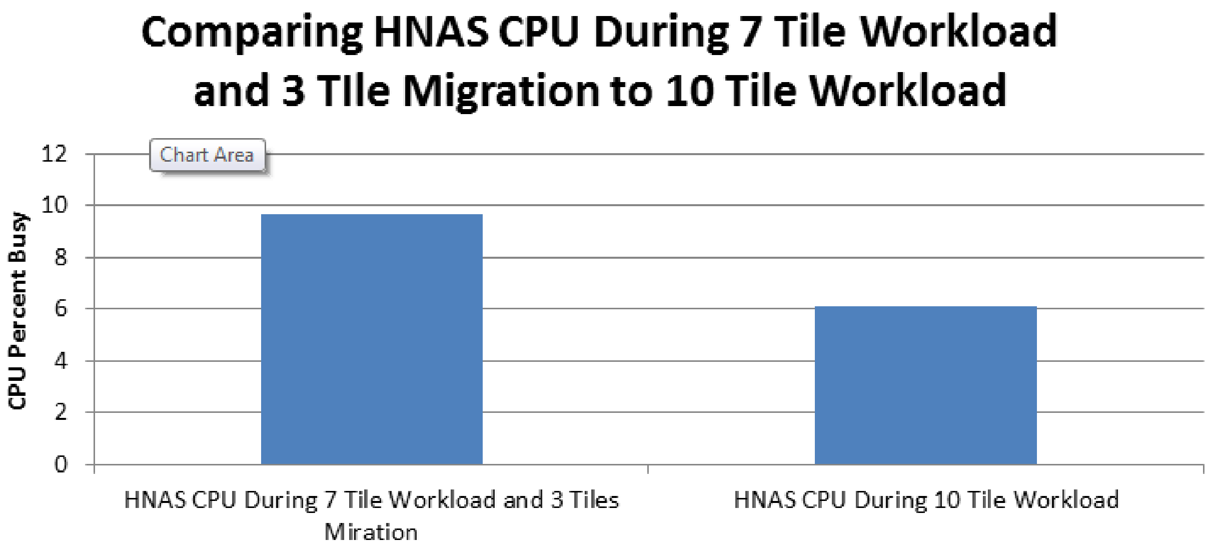

Figure 8

During the migration testing of tiles 8-10, with tile 1-7 workloads running, the application latency was an average of 6 milliseconds. While running 10 tiles workload and no migrations occurring, application latency was 5 milliseconds. Table 9 illustrates the comparison for each application's workloads and the average for all workloads.

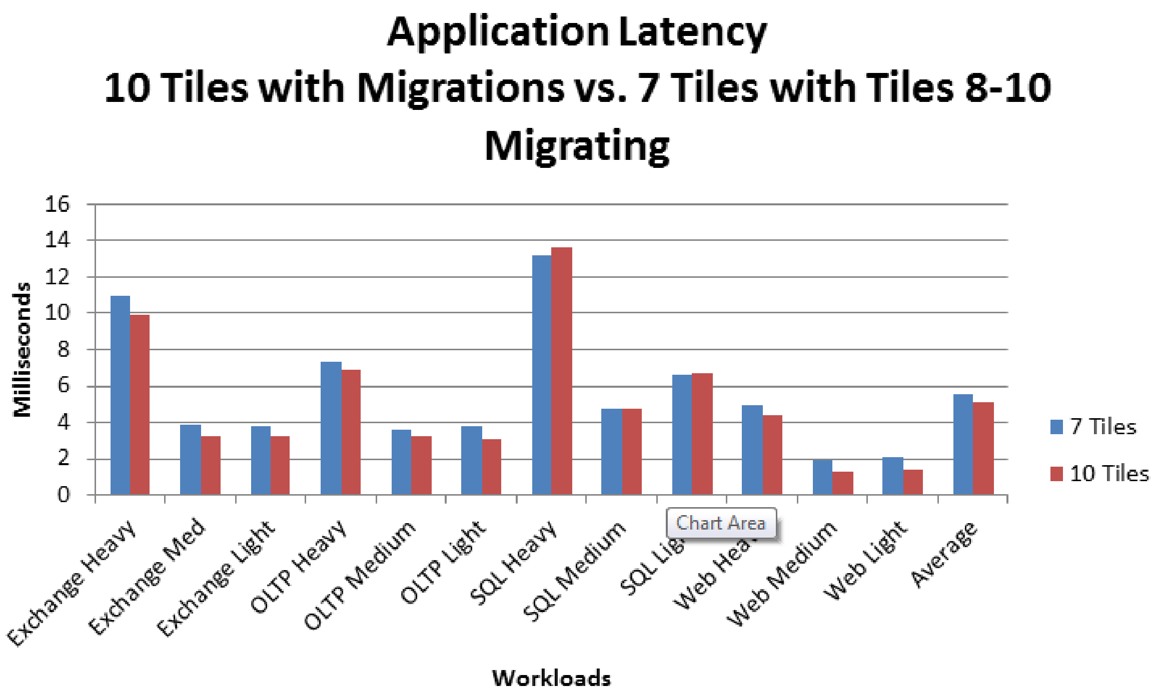

Figure 9

During the Fibre Channel testing, parity group percent busy was between 77 and 80 percent. During Hitachi NAS Platform NFS datastore testing, parity group percent busy was between 20 and 24 percent, as shown in Figure 10. This decrease in the parity group percent busy enabled the increase in application-based I/O. The drop in parity group percent busy was due to the Superflush feature. Superflush coalesces write I/O from the applications in sector cache, if blocks are contiguous, providing larger sequential I/O to increase throughput per IOP.

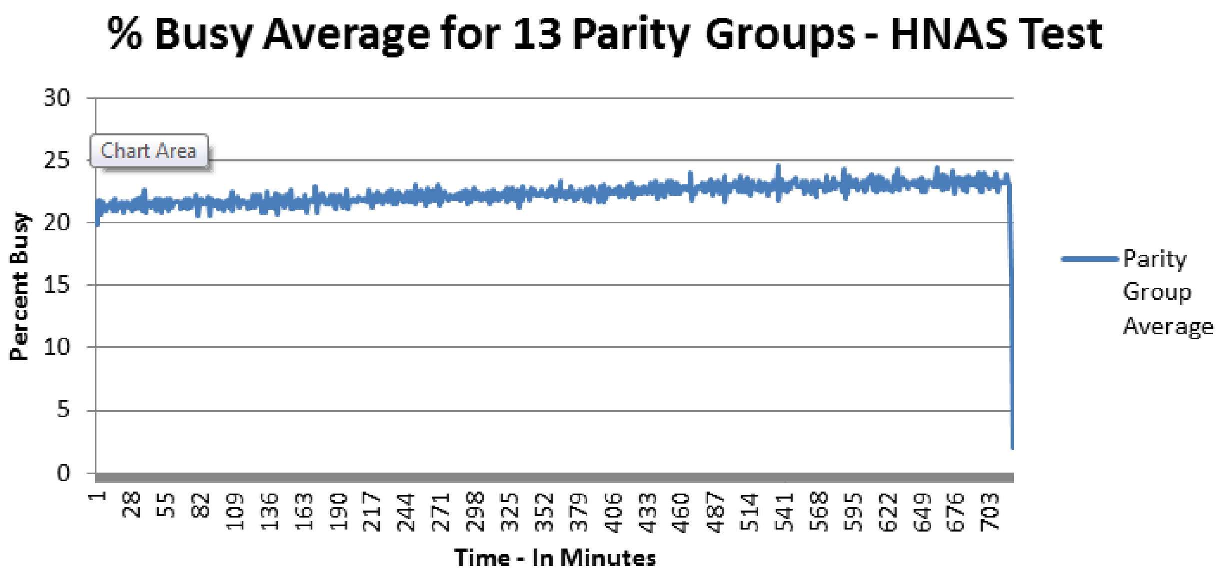

Figure 10

The Hitachi Unified Storage VM controllers were busy between 5 and 8 percent during testing, as shown in Figure 11. This leaves plenty of room for an increase in workload on the Hitachi NAS Platform node.

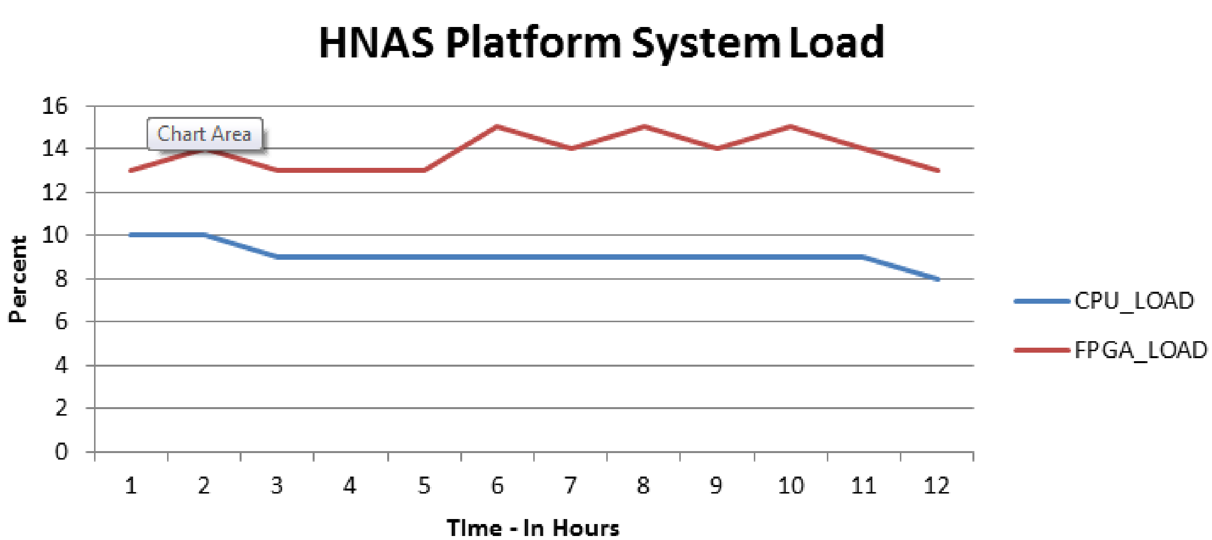

Figure 11

Figure 12 shows the received and transmitted Ethernet traffic on Hitachi NAS Platform 10 GbE interfaces.

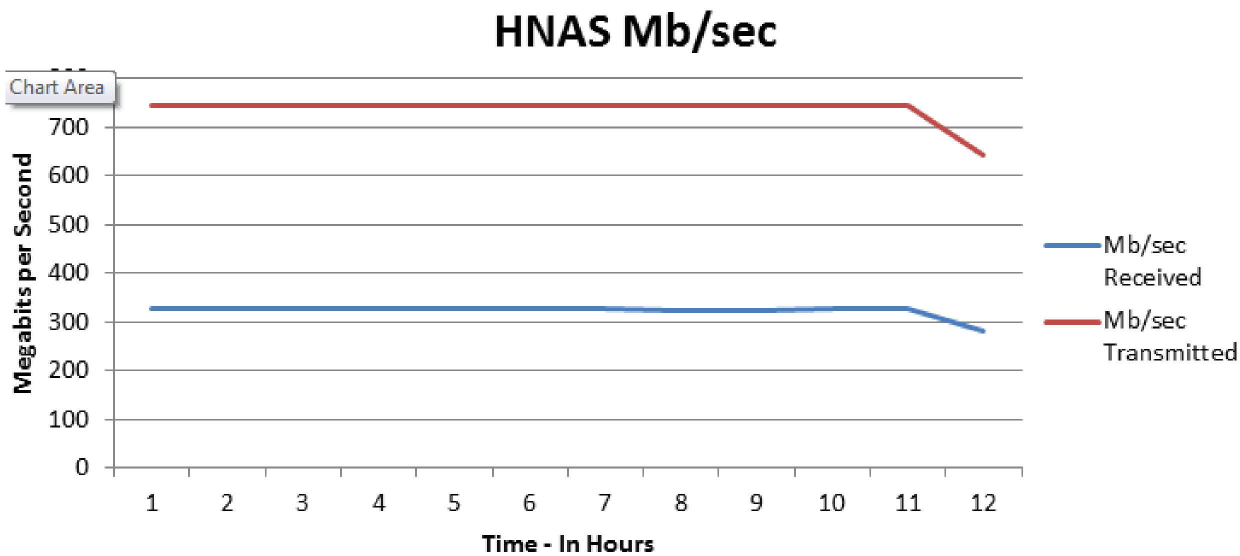

Figure 12

Migration of Very Large Virtual Machines (1 TB+) from Fibre Channel Storage to Hitachi NAS Platform

This test was performed outside of, and separate from, the other tests in this paper. This test was to validate that very large virtual machines (VMDKs of 1 TB and greater) could be migrated to Hitachi NAS Platform.

Table 12. Fourth Test Case

Test Case	*Description*	*Result*
Migrate a virtual machine with a 1.2 TB VMDK to Hitachi NAS Platform	Perform an svMotion on a virtual machine with a 1.2 TB VMDK with an Hitachi NAS Platform NFS datastore as the target.	Passed
	Expected result: The virtual machine should migrate successfully and be operational when completed.	

A virtual machine with a 20 GB operating system VMDK, and a 1.2 TB data VMDK was created. A workload ran against the data VMDK to fill the file with data so when using svMotion on Hitachi NAS Platform as thin provisioned, the resulting file would still be over 1 TB.

After the virtual machine was created and the data VMDK had data written to it, svMotion was started against the very large virtual machine with the Hitachi NAS platform as the target. The migration was completed in 37 minutes.

Conclusion

When Fibre Channel performance is compared to Hitachi NAS Platform during this testing, an advantage is demonstrated using Hitachi NAS Platform in terms of latency. Figure 13 shows that the latency for all but one workload was decreased during testing, and that the average overall average latency was decreased by six milliseconds. Lower latency time indicates higher I/O performance.

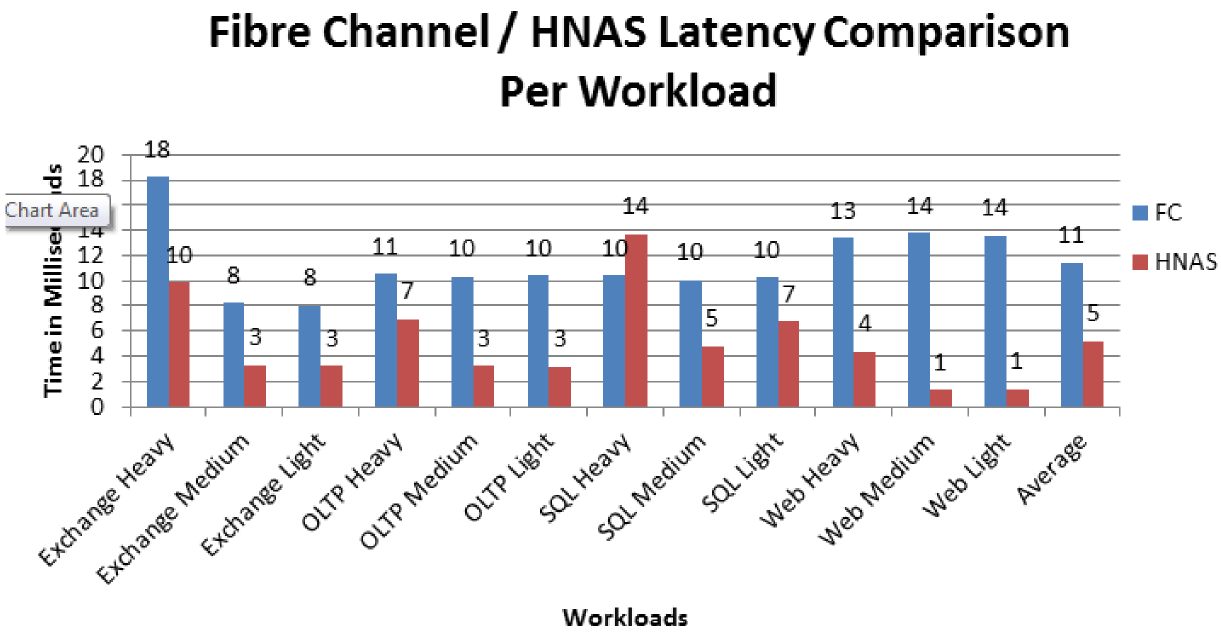

Figure 13

Figure 14 illustrates that MPU utilization for Hitachi Unified Storage VM decreased during Hitachi NAS Platform testing compared to Fibre Channel testing.

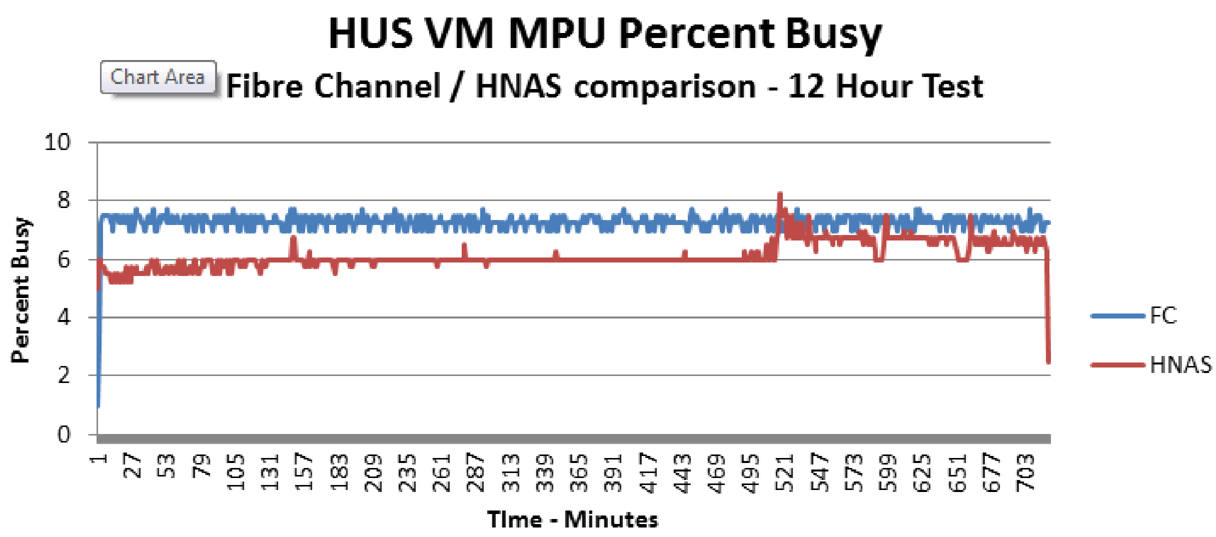

Figure 14

Figure 15 illustrates the sharp decrease in the parity group percent busy when using Hitachi NAS Platform during this testing compared to using Fibre Channel.

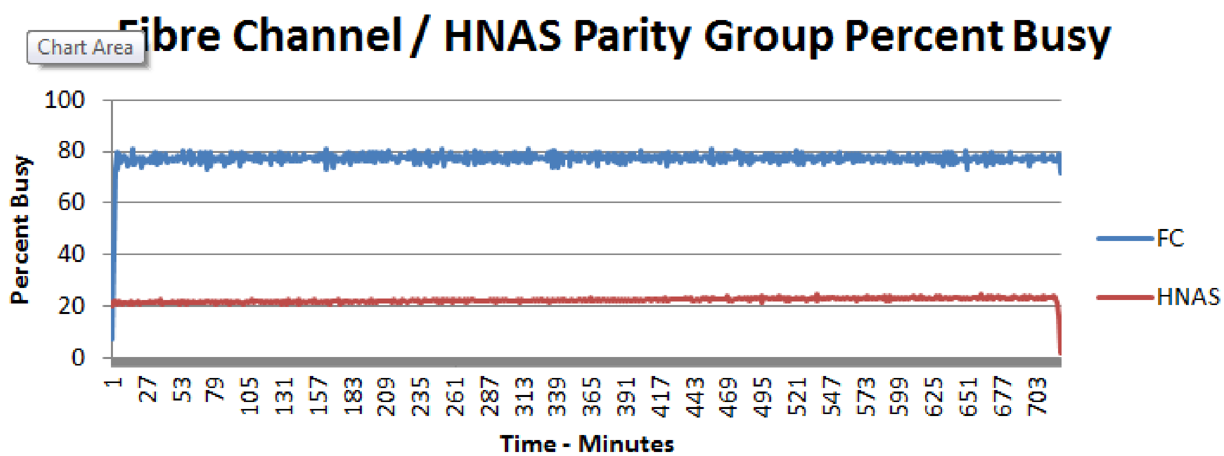

Figure 15

Figure 14 and Figure 15 illustrate a decrease in resource utilization on Hitachi Unified Storage VM when used with Hitachi NAS Platform, which resulted in the lower application latency as shown in Figure 10. This was due to Superflush. Superflush coalesces write I/O from the applications in sector cache, if blocks are contiguous, providing larger sequential I/O to increase throughput per IOP. This resulted in less load on Hitachi Unified Storage VM, and lower latency for the applications.

Superflush is the term for a "full stripe write" and the coalescing of writes it requires.

In order to do this, Hitachi NAS Platform “coalesces” writes, holding them in NVRAM until contiguous free space in the file system is found. Then Hitachi NAS Platform “flushes” the data to disk in a full-stripe-write.

For More Information

Hitachi Data Systems Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Hitachi Data Systems Global Services website.

Live and recorded product demonstrations are available for many Hitachi products. To schedule a live demonstration, contact a sales representative. To view a recorded demonstration, see the Hitachi Data Systems Corporate Resources website. Click the **Product Demos** tab for a list of available recorded demonstrations.

Hitachi Data Systems Academy provides best-in-class training on Hitachi products, technology, solutions and certifications. Hitachi Data Systems Academy delivers on-demand web-based training (WBT), classroom-based instructor-led training (ILT) and virtual instructor-led training (vILT) courses. For more information, see the Hitachi Data Systems Services Education website.

For more information about Hitachi products and services, contact your sales representative or channel partner or visit the Hitachi Data Systems website.

Hitachi Data Systems Corporation

Corporate Headquarters
2845 Lafayette Street, 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

AS-370-00, March 2015