



# Hitachi Virtual Storage Platform Integration with VMware® VAAI

Lab Validation Report

*By*  
*Henry Chu, Hitachi Data Systems*  
*Roger Clark, Hitachi Data Systems*  
*Erika Nishimoto, Hitachi, Ltd.*

*August 2011*

## Feedback

Hitachi Data Systems welcomes your feedback. Please share your thoughts by sending an email message to [SolutionLab@hds.com](mailto:SolutionLab@hds.com). Be sure to include the title of this white paper in your email message.

# Table of Contents

<b>VMware VAAI Overview .....</b>	<b>3</b>
Full Copy .....	4
Block Zeroing .....	5
Hardware-assisted Locking .....	6
<b>Engineering Validation.....</b>	<b>7</b>
Test Environment .....	7
Test Methodology .....	8
Test Results.....	9
<b>Conclusion .....</b>	<b>20</b>

# Hitachi Virtual Storage Platform Integration with VMware® VAAI

## Lab Validation Report

Data center administrators look for ways to improve scalability, performance, and efficiency to reduce administrative overhead and costs. One way to do this in a VMware environment is through the Hitachi Virtual Storage Platforms integration with VMware vStorage APIs for Array Integration (VAAI).

VAAI is a set of APIs, or primitives, that allow ESX hosts to offload processing for certain data-related services to a Hitachi Virtual Storage Platform. This can enable significant improvements in virtual machine performance, virtual machine density, and availability in vSphere 4.1 environments. Moving these functions to a storage system offers many benefits, but requires the use of a highly available, scalable, high performance storage system like a Hitachi Virtual Storage Platform.

A Hitachi Virtual Storage Platform works seamlessly with VMware's ESX 4.1 and ESXi 4.1 virtualization software. This enhances vSphere tasks such as Storage vMotion, cloning, and provisioning new virtual machines. It also reduces SCSI reservation locking overhead that can lead to decreased performance.

Hitachi Dynamic Provisioning enables the creation of a storage pool from which capacity can be used as needed to further improve performance, scalability, and utilization. In addition, the Hitachi Virtual Storage Platform can do this utilizing supported external storage, protecting your current storage investment. The external storage used for this paper was a Hitachi Adaptable Modular Storage 1000.

This white paper validates the benefits of using VAAI with a Hitachi Virtual Storage Platform with internal and external storage. It is written for storage administrators, vSphere administrators, and application administrators who manage large, dynamic environments. It assumes familiarity with SAN-based storage systems, VMware vSphere, and general IT storage practices.

## VMware VAAI Overview

VMware VAAI enables execution of key data operations at the storage level such as a Hitachi Virtual Storage Platform, rather than at the ESX server layer. This reduces resource utilization and potential bottlenecks on physical servers. It also enables more consistent server performance and higher virtual machine density.

When used with vSphere 4.1, a Hitachi Virtual Storage Platform supports the following API primitives:

- **Full Copy**—Enables the storage system to make full copies of data within the storage system without having the ESX host read and write the data. Read more about [Full Copy](#).
- **Block Zeroing**—Enables storage systems to zero out a large number of blocks to speed provisioning of virtual machines. Read more about [Block Zeroing](#).
- **Hardware-assisted Locking**—Provides an alternative means to protect the metadata for VMFS cluster file systems, thereby improving the scalability of large ESX host farms sharing a datastore. Read more about [Hardware-assisted Locking](#).

## Full Copy

For common ESX administration tasks, use the full copy primitive on the ESX host to offload the actual data copy to a Hitachi Virtual Service Platform. For example, the full copy primitive helps with tasks such as provisioning virtual machines or migrating VMDK file between datastores within a storage system using Storage vMotion. The following operations are some examples of when the full copy primitive is used:

- **Virtual Machine Provisioning**—The source and destination locations are within the same volume. Hitachi integrates with the full copy API to clone virtual machines or datastores from a golden image. This process dramatically reduces I/O between the ESX nodes and Hitachi storage.
- **Storage vMotion**—The source and destination locations are different volumes within the same storage system. This feature enables VMDK files to be relocated between datastores within a storage system. Virtual machines can be migrated to facilitate load-balancing or planned maintenance without service interruption. By integrating with full copy, host I/O offload for VMware Storage vMotion operations accelerates virtual machine migration times considerably.

Figure 1 compares copy functions with and without VMware VAAI. As shown, the full copy primitive removes the ESX host from the data path of the VMDK cloning operation. This reduces the number of disk I/Os from the ESX host, saving host-side I/O bandwidth while copying virtual machines.

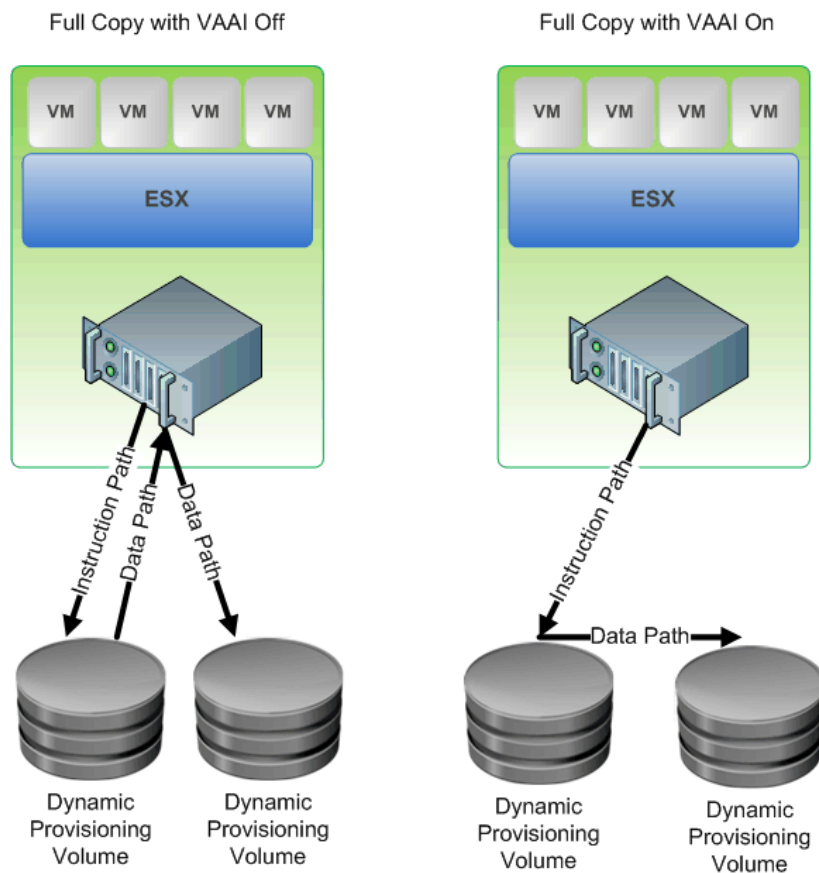


Figure 1

## Block Zeroing

ESX supports different space allocation options when creating new virtual machines or virtual disks. When using the zeroedthick format, the virtual disk's space is pre-allocated but not all pre-zeroed. Instead, the space is zeroed when the guest OS first writes to the virtual disk.

When using the eagerzeroedthick format, the virtual disk's space is pre-allocated and pre-zeroed. This means that it can take much longer to provision eagerzeroedthick virtual disks. The block zeroing primitive offloads these zeroing operations to the storage system without the host having to issue multiple commands.

Figure 2 compares block zeroing with and without VAAI. The block zeroing primitive allows zeroedthick or eagerzeroedthick to quickly provision VMDKs by writing zeros across hundreds or thousands of blocks on the VMFS datastores. This primitive is particularly useful when provisioning eagerzeroedthick VMDKs for VMware fault tolerant virtual machines due to the large number of blocks that need to be zeroed.




Figure 2

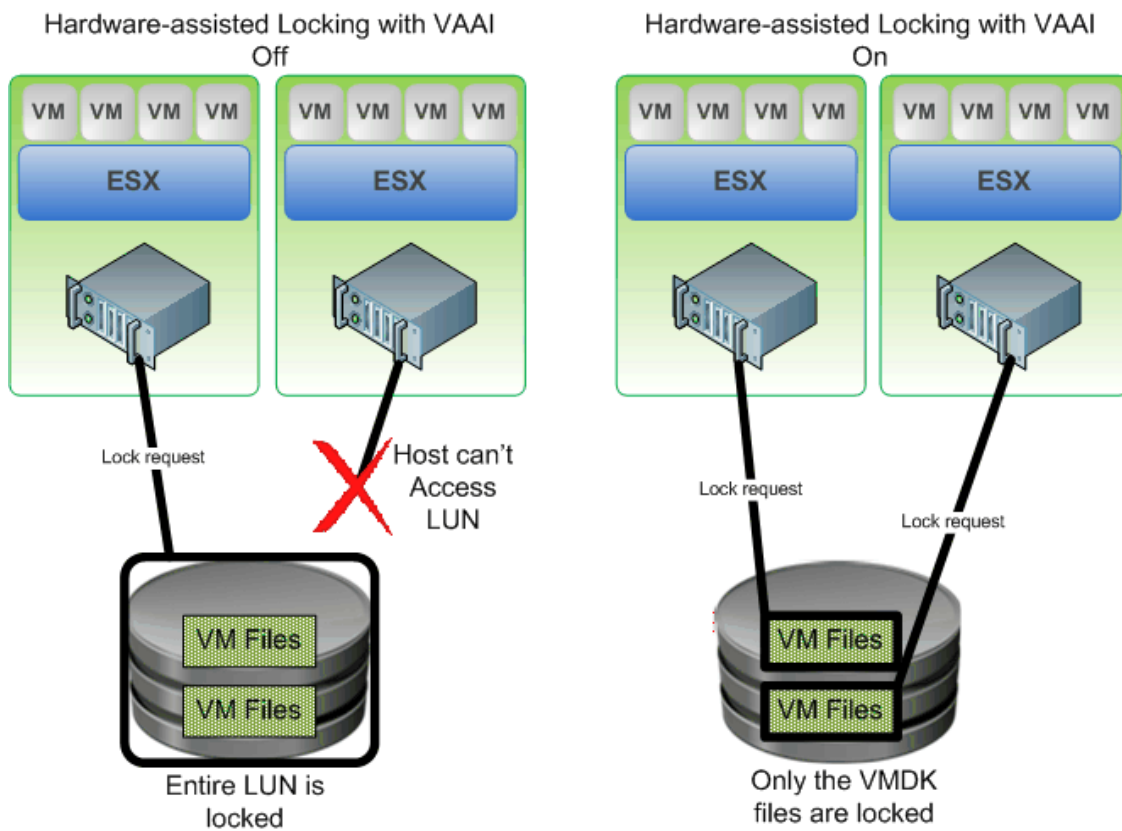
## Hardware-assisted Locking

ESX 4.1 environments rely on locking mechanisms to protect VMFS metadata, particularly in clustered environments where multiple ESX hosts access the same LUN. Hardware-assisted locking provides a granular LUN locking method to allow locking at the logical block address level without the use of SCSI reservations or the need to lock the entire LUN from other hosts.

Without hardware-assisted locking, ESX uses SCSI reservations to prevent hosts from activating or sharing virtual disk content on more than one host at the same time. These SCSI locking algorithms lock an entire LUN and do not provide the granularity to lock a particular block on the LUN. In addition, this algorithm requires four separate commands to acquire a lock (simplified as *reserve*, *read*, *write*, and *release*).

In addition to lack of granularity, locking the entire LUN introduces SCSI reservation contention. Both of these affect scalability.

Figure 3 shows a comparison of hardware-assisted locking with and without VAAI.



**Figure 3**

Transferring the LUN locking process to a Hitachi Virtual Storage Platform reduces the number of commands required to access a lock and allows more granular locking. This leads to better overall performance and increases the number of virtual machines per datastore and the number of hosts accessing the datastore.

The following are example use cases for hardware-assisted locking:

- Migrating a virtual machine with vMotion
- Creating a new virtual machine or template
- Deploying a virtual machine from a template
- Powering a virtual machine on or off
- Creating, deleting, or growing a file
- Creating, deleting, or growing a snapshot

## Engineering Validation

To demonstrate the VMware VAAI capabilities of the Hitachi Virtual Storage Platform, Hitachi Data Systems configured a vSphere environment and tested it using the following use cases:

- **Provisioning VMDK time Improvement** — Creating eagerzeroedthick VMDK files on fresh and fragmented datastores using block zeroing.
- **Disk Type conversation with Fault Tolerance** — Recover space from converted Fault Tolerance Virtual Machine using Zero Page reclaim
- **Warm-up time improvement** — Improve warm-up time of thin provisioned Virtual Machine using Block Zeroing
- **VM Cloning time improvement** — Cloning Virtual Machines using Full Copy
- **Boot Storm** — Improving Virtual Machine boot times using Hardware-assisted Locking
- **Large scale simultaneous vMotion** — Reducing SCSI locks using Hardware-assisted Locking

The goal of these tests was to compare times and I/O performance with VAAI on and off. The test results report IOPS on each Fibre Channel port, response time, and total completion times.

---

**Note**—All testing was done in a lab environment. In production environments, results can be affected by many factors that cannot be predicted or duplicated in a lab. Conduct proof-of-concept testing using your target applications in a non-production, isolated test environment that is identical to your production environment. Following this recommended practice allows you to obtain results closest to what you can expect to experience in your deployment. The test results included in this document are not intended to demonstrate actual performance capability of the Hitachi Virtual Storage Platform.

---

## Test Environment

The environment for these tests consisted of one to four VMware ESX 4.1 hosts attached to a Hitachi Virtual Storage Platform using internal and external storage. All of the ESX hosts used redundant paths for both the HBAs and the NICs. The host configuration followed VMware recommended practices.

For more information, see [“Optimizing the Hitachi Virtual Storage Platform in vSphere 4 Environments.”](#)

Table 1 lists the hardware used in the Hitachi Data Systems lab.

**Table 1. Hardware Resources**

<i>Hardware</i>	<i>Description</i>	<i>Version</i>
Hitachi Virtual Storage Platform	4 × 8GB Fibre Channel ports 42GB cache memory 38 × 300GB 10k SAS 26 × 146GB 15K SAS	Microcode: 70-02-03-00/00 SVP: 70-02-03/00 RIM server: 06-02-00/00
Hitachi Adaptable Modular Storage 1000	Dual Controllers 4 × 4GB Fiber Channel ports per controller 8GB cache memory, 2GB per controller	Microcode: 0786/A-H
Brocade 48000 director	Director-class SAN switch with 4Gb/sec Fibre Channel ports	FOS 5.3.1a
Hitachi Compute Blade 2000 chassis	2 × 8Gb/sec Fibre Channel switch modules 4 × 1Gb/sec Ethernet switch modules	A0160-G-5666
Hitachi X55A2 server blades	2 × six core Intel CPU 72GB RAM 2 × 300GB 10k SAS	03-79
Dell R710	1 × quad-core Intel CPU 16GB RAM 2 × 72GB SAS	A.50, A15

Microsoft Windows 2008 R2 Enterprise was installed on the virtual machines used in the Hitachi Data Systems lab for this test environment. Each virtual machine was configured with one virtual CPU and 1GB of RAM.

A standalone server running Windows 2008 R2 Enterprise with quad-core CPUs and 2GB RAM was used to host the VMware vCenter server and VMware vCenter client.

## Test Methodology

To measure the duration of each test case, the VMware vSphere client captured the requested start time and the end time of each task to determine how long each task required to complete. Each test was performed four times to validate the results, with the average value reported.

The tests that required disk I/O to be measured used a custom script to simultaneously launch VDbench 5.02 across all of the virtual machines. VDbench repeatedly created and deleted a 500GB file on each virtual machine to generate a consistent level of VMFS metadata operations and disk I/O traffic.

To measure the number of host I/Os and SCSI conflicts being generated on each host, `esxtop` was used. At the end of each test, `esxtop` logs were collected from all the ESX hosts and then averaged for each host.

## Test Results

The following test results show the benefits of using the built in VAAI support in the Hitachi Virtual Storage Platform. The tests show how the Hitachi Virtual Storage Platform improves performance and scalability when using the VAAI primitives. Each test consists of everyday administrative tasks that can put extra stress on a VMware vSphere environment that result in performance slow downs.

### *Provisioning VMDK Time Improvement*

This test is designed to test the block zeroing primitive of VAAI. For this test five new 159GB eagerzeroedthick virtual machines are created on fresh and dirty datastores. The time it took to complete the provisioning was measured. Internal and virtualized external storage are tested.

A fresh datastore is a newly formatted datastore that has not been used. A dirty datastore is a datastore that has had VMDK files created and deleted from it without enabling the Write Same or Zero Page Reclaim options.

To simulate a dirty datastore, a 797GB eagerzeroedthick virtual machine was created on a fresh 799GB datastore with VAAI disabled. Then, the virtual machine was deleted without running a zero page reclaim on the DP-VOL hosting the datastore.

Table 2 shows the internal storage configuration used on the Hitachi Virtual Storage Platform for this test. DP-VOL 0 holds the virtual machine image used as the source of the virtual machines creation. DP-VOL 1 hosts the destination datastore.

**Table 2. Internal Hitachi Virtual Storage Platform Configuration for Provisioning VMDK Time Improvement**

DP-VOL (LU)		Pool-VOL (LDEV)		PG (Parity Group)					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1950	0	801	0	805	SAS	10k	300GB	RAID-5 (3D+1P)
		1	805	1	805	SAS	10k		
1	799	2	397	2	401	SAS	15k	146GB	RAID-5 (3D+1P)
		3	401	3	401	SAS	15k		

Table 3 shows the external storage configuration used on the Hitachi Virtual Storage Platform for this test. DP-VOL 0 holds the virtual machine image used as the source of the virtual machines creation. DP-VOL 1 hosts the destination datastore.

**Table 3. External Hitachi Virtual Storage Platform Configuration for Provisioning VMDK Time Improvement**

DP-VOL (LU)		Pool-VOL (LDEV)		PG (Parity Group)					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1950	0	801	0	805	SAS	10k	300GB	RAID-5 (3D+1P)
		1	805	1	805	SAS	10k		
1	799	E1	1067	E1	1071.50	SAS	10k	300GB	RAID-5 (4D+1P)
		E2	1071	E2	1071.50	SAS	10k		

Creation time of virtual machines was improved 96 percent for internal storage and 98 percent for external storage by using block zeroing on a fresh datastore. Creation time was improved 2 percent for internal storage and 5 percent for external storage by using block zeroing on a dirty datastore.

Table 4 shows the time required to complete each test.

**Table 4. Time required to complete the test**

<i>Cloning with Full Copy</i>	<i>Status</i>	<i>Internal or External</i>	<i>VAAI Off (Min:Sec)</i>	<i>VAAI On (Min:Sec)</i>	<i>Improvement using VAAI</i>
159GB Eagerzeroedthick	Fresh	Internal	5:32	0:11	96%
159GB Eagerzeroedthick	Dirty	Internal	5:56	5:50	2%
159GB Eagerzeroedthick	Fresh	External	10:20	0:13	98%
159GB Eagerzeroedthick	Dirty	External	18:40	17:50	5%

---

**Key Finding**—Because the full copy primitive uses the storage system to execute commands, the ESX hosts are less burdened with respective copy commands. This allows the host more cycles for processing other tasks.

---

### *Disk Type Conversion with Fault Tolerance*

This tested the block zeroing primitive of VAAI. A 100GB zeroedthick virtual machine was converted to eagerzeroedthick by enabling the fault tolerant option for a virtual machine in vCenter. After completing the conversion, a zero page reclaim was run on the dynamic pool hosting the datastore. The storage capacity recovered was measured. For this test, only internal storage was tested.

Block zeroing enables the formatting and provisioning of both zeroedthick and eagerzeroedthick VMDKs and in VMFS datastores to be handled by the storage system rather than the ESX hosts.

Table 5 shows the internal storage configuration used on the Hitachi Virtual Storage Platform for this test.

**Table 5. Internal Hitachi Virtual Storage Platform Configuration for Disk Type Conversion Test**

<i>DP-VOL (LU)</i>		<i>Pool-VOL (LDEV)</i>		<i>PG (Parity Group)</i>					
<i>ID</i>	<i>Capacity (GB)</i>	<i>ID</i>	<i>Capacity (GB)</i>	<i>ID</i>	<i>Capacity (GB)</i>	<i>Drive Type</i>	<i>Rotation</i>	<i>Drive Size</i>	<i>RAID Level</i>
0	1945GB	0	397GB	0	401GB	SAS	15k	146GB	RAID-5 (3D+1P)
		1	401GB	1	401GB	SAS	15k		

With zero page reclaim, any pages that are zeroed are freed to be reallocated. To demonstrate this functionality, VMware VAAI was enabled and a 100GB zeroedthick VMDK file with 8GB of used space within the VMDK file was converted to an eagerzeroedthick VMDK file.

Figure 4 shows an eagerzeroedthick VMDK file on a dynamic provisioning volume before and after running a zero page reclaim operation. Table 6 lists the test results.

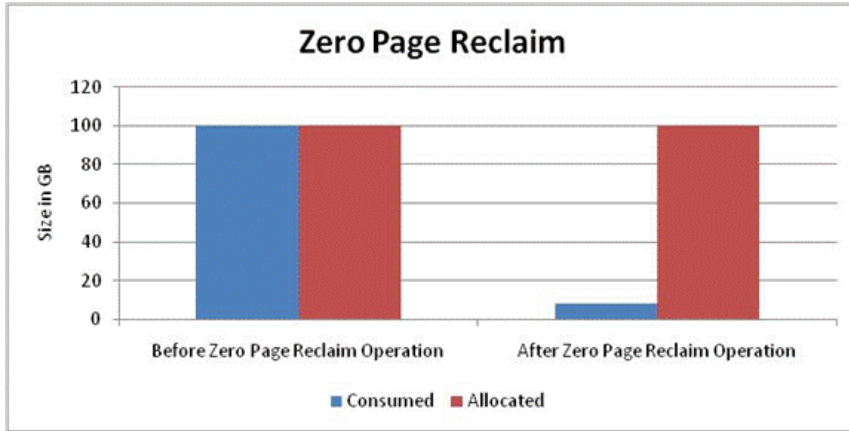


Figure 4

Table 6. Zero Page Reclaim During Disk Type Conversion Test

<i>Zero Page Reclaim</i>	<i>Before Zero Page Reclaim</i>	<i>After Zero Page Reclaim</i>	<i>Storage Reclaimed</i>
Consumed	100GB	8GB	92.0%
Allocated	100GB	100GB	0.0%

As expected after the conversion, the entire 100GB that was allocated to the converted VMDK file was zeroed and committed to the virtual machine. After the zero page reclaim operation ran, the virtual machine was still allocated the original 100GB but only 8GB of space was actually being consumed on the storage system by the virtual machine.

---

**Key Finding**—The block zeroing primitive in VAAI used with Hitachi Dynamic Provisioning allows for faster provisioning of eagerzeroedthick VMDK files. The acceleration results from the Hitachi Virtual Storage Platform handling the zeroing internally rather than ESX host handling the zeroing using multiple write requests.

---

### *Warm-up Time Improvement*

This tests the effect of thin provisioning eagerzeroedthick VMDK with VAAI using Hitachi Dynamic Provisioning. A 100GB eagerzeroedthick virtual machine and a zeroedthick virtual machine were booted. The time it took for the virtual machine to warm-up was captured.

A virtual machine was considered warmed-up when the IOPS reached a steady state. Because both zeroedthick and eagerzeroedthick virtual machines is provisioned using the same storage configuration, the IOPS will be nearly identical when warm-up is complete. Tests included internal and external storage.

Table 7 shows the internal storage configuration used on the Hitachi Virtual Storage Platform for this test. DP-VOL 0 holds the virtual machine image used as the source of the virtual machines creation. DP-VOL 1 hosts the destination datastore.

**Table 7. Internal Hitachi Virtual Storage Platform Storage Configuration for the Warm-up Time Improvement Test**


DP-VOL (LU)		Pool-VOL (LDEV)		PG (Parity Group)					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1945.6	0	801	0	805	SAS	10k	300GB	RAID-5 (3D+1P)
		1	805	1	805	SAS	10k		
1	1945.6	2	397	2	401	SAS	15k	146GB	RAID-5 (3D+1P)
		3	401	3	401	SAS	15k		

Table 8 shows the external storage configuration used on the Hitachi Virtual Storage Platform for this test. For this test DP-VOL 0 holds the virtual machine image used as the source of the virtual machines creation. DP-VOL 1 hosts the destination datastore.

**Table 8. External Hitachi Virtual Storage Platform Storage Configuration for the Warm-up Time Improvement Test**

DP-VOL (LU)		Pool-VOL (LDEV)		PG (Parity Group)					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1945.6	0	801	0	805	SAS	10k	300GB	RAID-5 (3D+1P)
		1	805	1	805	SAS	10k		
1	1945.6	E1	1067	2	1071.50	SAS	10k	300GB	RAID-5 (4D+1P)
		E2	1071	3	1071.50	SAS	10K		

Figure 5 shows the warm-up test times for eagerzeroedthick compared to zeroedthick in a dynamic provisioning volume on internal storage.



**Figure 5**

Figure 6 shows the warm-up test times for eagerzeroedthick versus zeroedthick in a dynamic provisioning volume on external storage.




Figure 6

Table 9. Time to Complete the Warm-up Time Improvement Test

<i>Internal/External</i>	<i>Disk Format</i>	<i>Improvement</i>
Internal	Zeroedthick	76%
	Eagerzeroedthick	
External	Zeroedthick	61%
	Eagerzeroedthick	

With thin provisioned eagerzeroedthick VMDK, the warm-up time decreased by 76% for internal storage and 61% for external storage. The overall IOPS during warm-up for eagerzeroedthick is substantially higher than zeroedthick VMDK.

---

**Key Finding**—Thin provisioning eagerzeroedthick VMDK with VAAI in a Hitachi Dynamic Provisioning pool results in reduced warm-up time.

---

### *VM Cloning Time Improvement*

This test is designed to test the block zeroing and full copy primitives of VAAI. For this test a 100GB eagerzeroedthick and zeroedthick virtual machines were cloned. During the testing IOPS and the time it took to complete the cloning was captured. Internal and external storage was tested.

Table 10 shows the internal storage configuration used on the Hitachi Virtual Storage Platform for this test. DP-VOL 0 holds the virtual machine image used as the source of the virtual machines creation. DP-VOL 1 hosts the destination datastore.

**Table 10. Internal Hitachi Virtual Storage Platform Configuration for the VM Cloning Time Improvement Test**

<i>DP-VOL (LU)</i>		<i>Pool-VOL (LDEV)</i>		<i>PG (Parity Group)</i>					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1945.6	0	397	0	401	SAS	15k	146GB	RAID-5 (3D+1P)
		1	401	1	401	SAS	15k		
1	1945.6	2	397	2	401	SAS	15k	146GB	RAID-5 (3D+1P)
		3	401	3	401	SAS	15k		

Table 11 shows the external storage configuration used on the Hitachi Virtual Storage Platform for this test. For this test DP-VOL 0 holds the virtual machine image used as the source of the virtual machines creation. DP-VOL 1 hosts the destination datastore.

**Table 11. External Hitachi Virtual Storage Platform Storage Configuration for the VM Cloning Time Improvement Test**

<i>DP-VOL (LU)</i>		<i>Pool-VOL (LDEV)</i>		<i>PG (Parity Group)</i>					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1945.6	0	397	0	401	SAS	15k	146GB	RAID-5 (3D+1P)
		1	401	1	401	SAS	15k		
1	1945.6	E3	1067	E3	1071.50	SAS	10k	300GB	RAID-5 (4D+1P)
		E4	1071	3	1071.50	SAS	10k		

Table 12 shows the vMotion Migration times

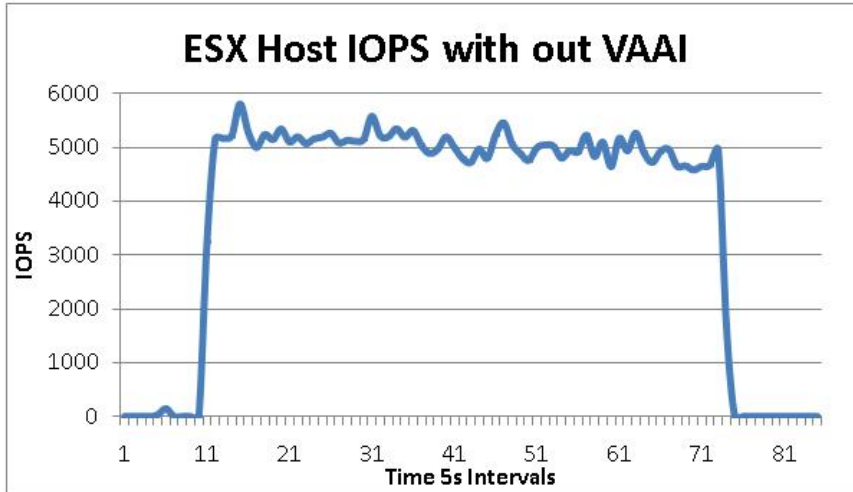
**Table 12. vMotion Migration Times**

<i>VMDK format</i>	<i>Intern/External</i>	<i>Improvement</i>
Eagerzeroedthick	Internal	6%
Zeroedthick	Internal	29%
Eagerzeroedthick	External	0%
Zeroedthick	External	4%

Using internal storage, cloning time was decreased by 6% for eagerzeroedthick virtual machine. The zeroedthick virtual machine cloning time was decreased by 29% on the internal storage.

The zeroedthick virtual machine cloning time was decreased by 4% on the external storage. Although there was no improvement in time on the external eagerzeroedthick virtual machine, there is still the benefit of offloading the IOPS from the ESX hosts to the Hitachi Virtual Storage Platform.

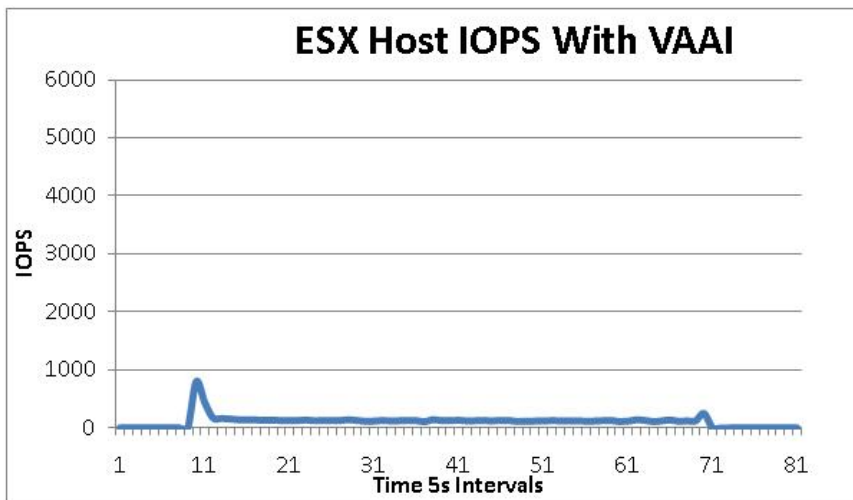
In addition to decreasing cloning times, the number of IOPS is greatly reduced during the cloning process. Figure 7 shows the IOPS for the destination ESX hosts during cloning without VAAI.



**Figure 7**

The IOPS stay near 5000 using the cloning process without VAAI.

Using VAAI, the IOPS are reduced greatly. Figure 8 shows the IOPS for the destination ESX hosts during cloning using VAAI.



**Figure 8**

---

**Key Finding**—Using VAAI improves cloning time for virtual machines. This includes using internal storage or external storage.

---

### *Boot Storm*

This tested the hardware assisted locking primitive of VAAI. In this test case, linked clone virtual machines were evenly distributed across four ESX 4.1 hosts using a single shared datastore.

This was created on an LDEV provisioned from a 12-spindle dynamic provisioning pool for the internal storage. For the external Hitachi Virtual Storage Platform, a 15-spindle dynamic provisioning pool was used. All the virtual machines were powered on simultaneously.

Then, the same test was repeated using a datastore that was created on an LDEV provisioned from a 4-spindle dynamic provisioning pool for the internal storage and a 5-spindle dynamic provisioning pool for the external storage to increase the likelihood of SCSI locking conflicts.

Table 13 shows the internal storage configuration used on the Hitachi Virtual Storage Platform for the test with 4 spindles.

**Table 13. Internal Hitachi Virtual Storage Platform Configuration for the Boot Storm Test with 4 Spindles**

<i>DP-VOL (LU)</i>		<i>Pool-VOL (LDEV)</i>		<i>PG (Parity Group)</i>					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1945.6	0	397	0	401	SAS	15k	146GB	RAID-5 (3D+1P)

Table 14 shows the internal storage configuration used on the Hitachi Virtual Storage Platform for the test with 12 spindles.

**Table 14. Internal Hitachi Virtual Storage Platform Configuration for the Boot Storm Test with 12 Spindles**

<i>DP-VOL (LU)</i>		<i>Pool-VOL (LDEV)</i>		<i>PG (Parity Group)</i>					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
0	1945.6	0	397	0	401	SAS	15k	146GB	RAID-5 (3D+1P)
		1	401	1	401	SAS	15k		
		2	401	2	401	SAS	15k		

Table 15 shows the external storage configuration used on the Hitachi Virtual Storage Platform for the test with 5 spindles.

**Table 15. External Hitachi Virtual Storage Platform Configuration for the Boot Storm Test with 5 Spindles**

<i>DP-VOL (LU)</i>		<i>Pool-VOL (LDEV)</i>		<i>PG (Parity Group)</i>					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
2	1945.6	E0	1037.34	E0	1072.5	SAS	10k	300GB	RAID-5 (4D+1P)

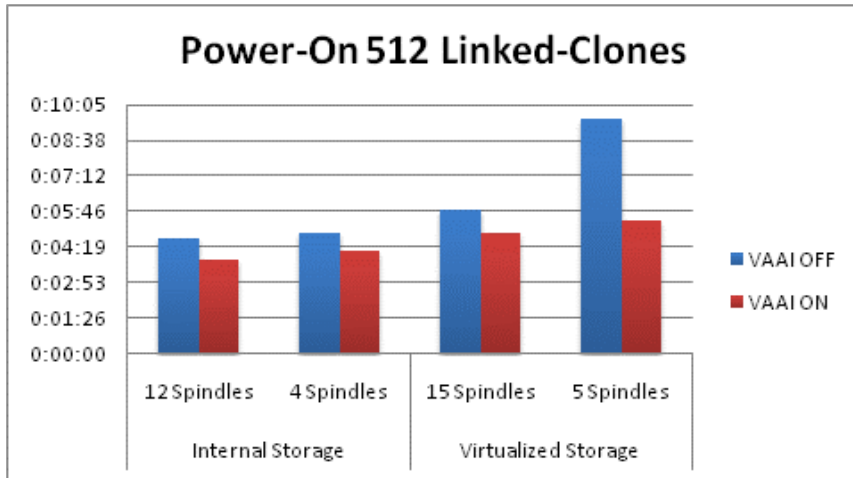
Table 16 shows the external storage configuration used on the Hitachi Virtual Storage Platform for the tests with 15 spindles.

**Table 16. External Hitachi Virtual Storage Platform Configuration for the Boot Storm Test with 15 Spindles**

DP-VOL(LU)		Pool-VOL(LDEV)		PG(Parity Group)					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
3	1945.6	E1	1067.3	E1	1071.5	SAS	10k	300GB	RAID-5 (4D+1P)
		E2	1071.45	E2	1071.5	SAS	10k		
		E3	1071.45	E3	1071.5	SAS	10k		

Figure 9 shows that on internal storage using VAAI, boot times decreased by 20 percent when using the datastore on the 12-spindle dynamic provisioning pool and 15 percent when using the datastore on the 4-spindle dynamic provisioning pool.

With the external storage using VAAI, boot times decreased by 19 percent when using the datastore on the 15-spindle dynamic provisioning pool and 43 percent when using the datastore on the 5-spindle dynamic provisioning pool. Figure 9 shows the boot times for all 512 virtual machines



**Figure 9**

Figure 10 shows the number of SCSI conflicts is reduced greatly during all the tests when using VAAI.

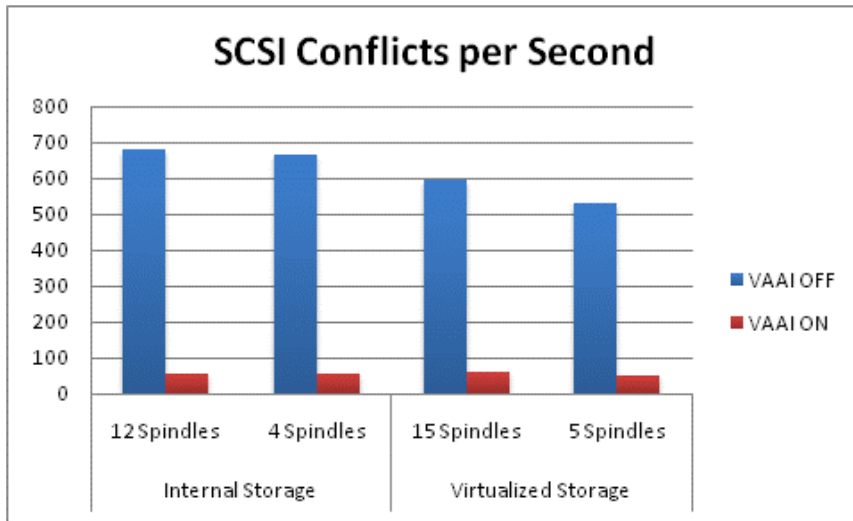


Figure 10

Table 17. SCSI Conflicts for 128 Virtual Machines per Host

<i>Intern/External</i>	<i>Number of Spindles</i>	<i>Improvement</i>
Internal	4	92%
Internal	12	92%
External	5	90%
External	15	90%

Without using VAAI, the datastore was more prone to SCSI reservation locks. This can dramatically reduce performance and the number virtual machines that can be run on a datastore.

However, using the hardware-assisted locking primitive, the likelihood of SCSI reservation locking conflicts during everyday tasks was reduced greatly. These acts include as Storage vMotion, creating or deleting VMDK files, and the powering off or on of virtual machines.

---

**Key Finding**—The use of the hardware-assisted locking primitive greatly improves the scalability of VMware vSphere by allowing more virtual machines per datastore to run concurrently.

---

### *Large Scale Simultaneous vMotion*

This tested the hardware assisted locking primitives of VAAI. For this test, 32 1.9TB linked clone virtual machines were migrated from one datastore to another using storage vMotion. The number of SCSI conflicts and the time it took for the virtual machines to complete the migration were captured. Tests were made on an internal and external Hitachi Virtual Storage Platform.

Table 18 shows the internal storage configuration used on the Hitachi Virtual Storage Platform for this test.

**Table 18. Internal Hitachi Virtual Storage Platform Configuration for the Large Scale Simultaneous vMotion Test**

<i>DP-VOL(LU)</i>		<i>Pool-VOL(LDEV)</i>		<i>PG(Parity Group)</i>					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID Level
1	1945.6	1	397	1	401	SAS	15k	146GB	RAID-5 (3D+1P)
		2	401	2	401	SAS	15k		
		3	401	3	401	SAS	15k		

Table 19 shows the external storage configuration used on the Hitachi Virtual Storage Platform for this test. DP-VOL 0 holds the virtual machine image used as the source of the virtual machines creation. DP-VOL 1 hosts the destination datastore.

**Table 19. External Hitachi Virtual Storage Platform Configuration for the Large Scale Simultaneous vMotion Test**

<i>DP-VOL(LU)</i>		<i>Pool-VOL(LDEV)</i>		<i>PG(Parity Group)</i>					
ID	Capacity (GB)	ID	Capacity (GB)	ID	Capacity (GB)	Drive Type	Rotation	Drive Size	RAID level
3	1945.6	E1	1067.34	E1	1071.5	SAS	10k	300GB	RAID-5 (3D+1P)
		E2	1067.34	E2	1071.5	SAS	10k		
		E3	1067.34	E3	1071.5	SAS	10k		

Table 20 shows the time required complete the test and the number of SCSI conflicts recorded.

**Table 20. Time Required to Complete the Large Scale Simultaneous vMotion Test with the Number of SCSI Conflicts for 32 Virtual Machines**

<i>VMDK format</i>	<i>Intern/External</i>	<i>VAAI</i>	<i>Number of Spindles</i>	<i>Time required</i>	<i>Number of SCSI conflicts</i>
Linked Clone	Internal	Disabled	12	4:58	90
Linked Clone	Internal	Enabled	12	4:36	10
Linked Clone	External	Disabled	15	4:35	84
Linked Clone	External	Enabled	15	5:01	10

---

**Key Finding**—The hardware-assisted locking primitive accelerated large scale migration with vMotion. This benefits maintenance operations, such as host patching and updates, which results in shorter maintenance windows and less downtime.

---

## Conclusion

A Hitachi Virtual Storage Platform, when coupled with VMware VAAI primitives, allows you to build and maintain more scalable and efficient virtual environments. When using the Virtual Storage Platform, the VAAI primitives are extended to external virtualized storage, whether the external storage supports or does not support VAAI. When you use these primitives with the Hitachi Virtual Storage Platform and Hitachi Dynamic Provisioning, you create a robust and highly available infrastructure to support high density virtual machine workloads.

The full copy primitive reduces host-side I/O during common tasks such as the following:

- Moving virtual machines with Storage vMotion
- Deploying a new virtual machine from a template by instructing the storage system to copy data within the Hitachi Virtual Storage Platform rather than sending the traffic back and forth through the ESX hosts

The block zeroing primitive speeds virtual machine deployment by offloading the repetitive zeroing of large numbers of blocks to the Hitachi Virtual Storage Platform. This frees the ESX host resources for other tasks.

The hardware-assisted locking primitive greatly reduces the probability of ESX hosts being locked out when attempting to access files on a VMFS datastore. A lock out can degrade performance and, in some cases, cause tasks to time out or fail completely.

The tight integration of VMware's VAAI and the Hitachi Virtual Storage Platform provides a proven high-performance, highly scalable storage solution for an ESX environment.

Table 21 summarizes the key benefits of using VMware VAAI with the Hitachi Virtual Storage Platform, as supported by the testing performed for this white paper.

**Table 21. Key Findings**

<i>VAAI Primitive</i>	<i>Benefit</i>
Full copy	Speeds virtual machine deployment with reduction of host HBA utilization.
Block zeroing	Enables storage systems to zero out a large number of blocks to speed provisioning of virtual machines. When used with Hitachi Dynamic Provisioning, all virtual disk types are thin provisioned, including eagerzeroedthick.
Hardware-assisted locking	Reduces locking conflicts and accelerates large-scale vMotion. This improves scalability of the vSphere infrastructure.  Allows the creation of large VMFS volumes (up to a 2TB single partition), thus simplifying storage configuration and sizing of LDEVs for VMFS volumes.

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 web site](#).

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 web site](#). 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 Academy web site](#).

For more information about Hitachi products, contact your sales representative or channel partner or visit the [Hitachi Data Systems web site](#).

## Hitachi Data Systems Corporation

---

*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 mentioned in this document 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. AS-101-00 August 2011

### **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)