VMware VAAI with HA3969 and HA3969U

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Date: August, 2014

Agenda

- Introduction to VAAI
- VAAI supported
- Physical environment
- Testing and results
- Summary

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Introduction to VAAI

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Date: August, 2014

VAAI introduction

- vStorage APIs for Array Integration (VAAI) is a feature introduced ESX/ESXi 4.1 that provides hardware acceleration functionality.

- VAAI enables key data operations to be executed at the storage level rather than at the ESX server layer.

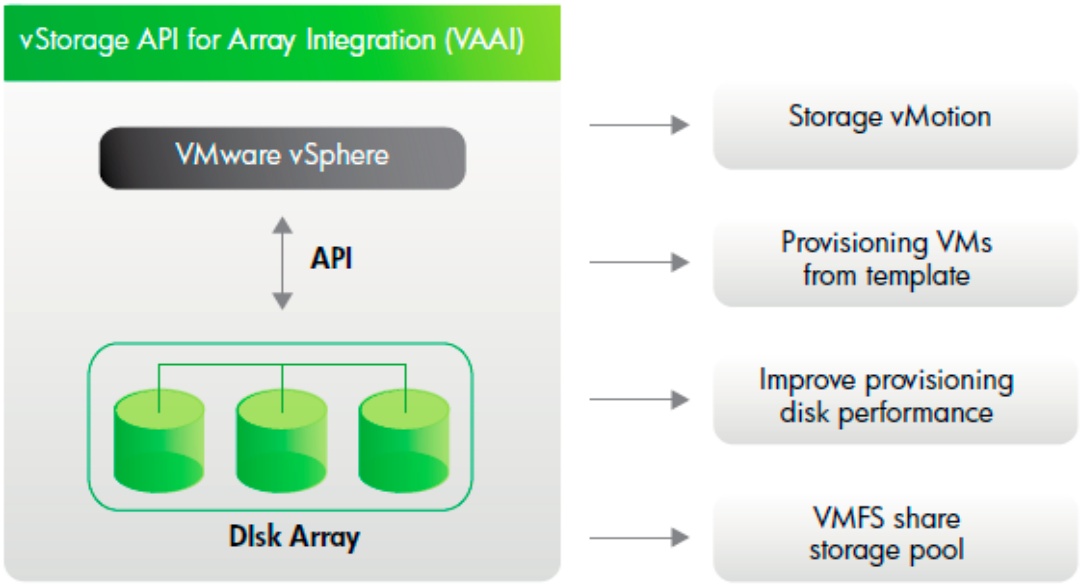

Storageflex
NETWORK DATA STORAGE SOLUTIONS

VAAI Primitives

VAAI uses these primitives:

- **Full Copy:** Clone Blocks/Full Copy/XCOPY, which is used to copy or migrate data within the same physical array
- **Block Zeroing:** Zero Blocks/Write Same, which is used to zero-out disk regions
- **Hardware Assisted Locking:** Atomic Test & Set (ATS), which is used during creation and locking of files on the VMFS volume

Full Copy (Hardware-Accelerated Copy)

- **Leverage native array copy capability to move blocks within the array**

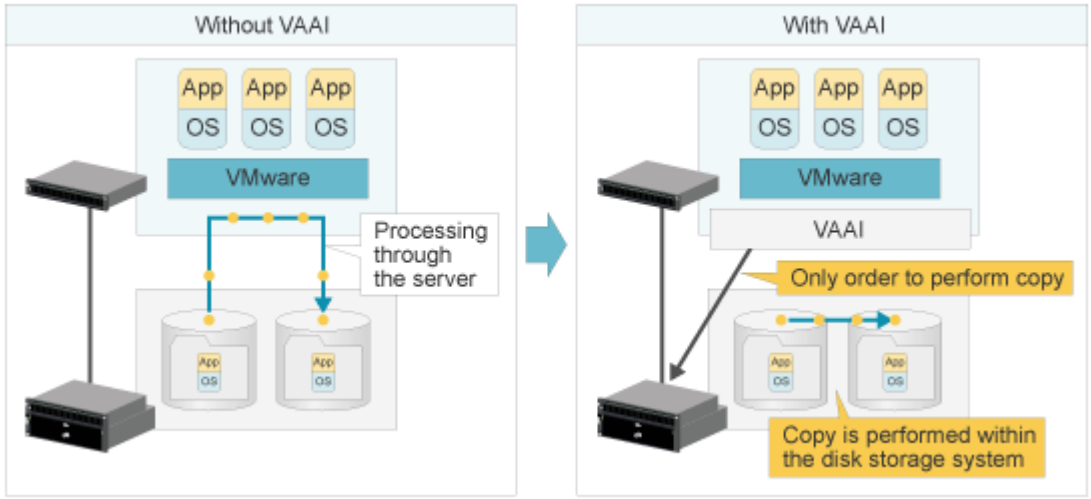

Without VAAI

- **SCSI Read** (Data moved from array to host)
- **SCSI Write** (Data moved from host to array)
- Repeat
- Huge periods of large VMFS level IO, done via millions of small block operations

With VAAI

- **SCSI Extended Copy** (XCOPY, Data moved within array)
- Repeat
- Order of magnitude reduction in host/array I/O operations

Block Zeroing (Hardware-Accelerated Zero)

- **Eliminate redundant and repetitive host-based write commands with optimized internal array commands**

Without Block Zeroing with VAAI

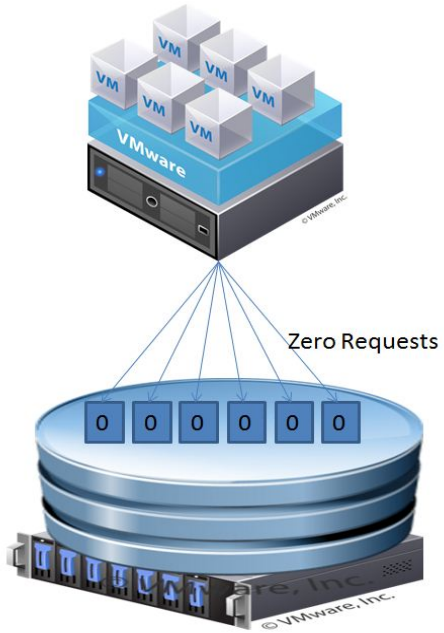

Block Zeroing with VAAI

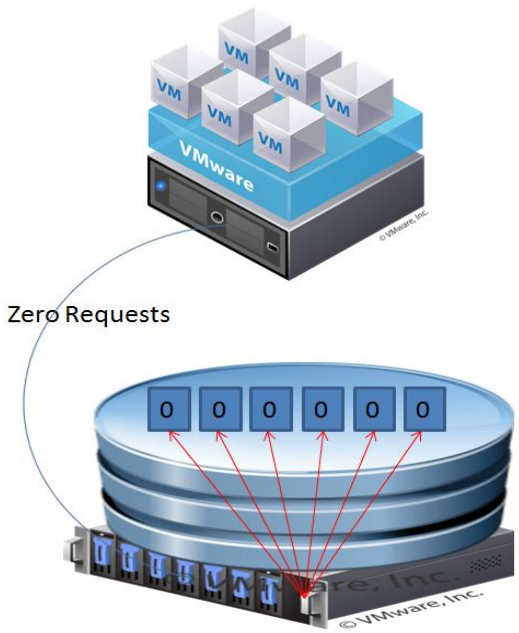

Without VAAI

- **SCSI Write** - Many identical small blocks of zeroes moved from host to array for MANY VMware IO operations

With VAAI

- **SCSI Write Same** - One giant block of zeroes moved from host to array and repeatedly written

ATS; Atomic Test & Set (Hardware-Accelerated Locking)

- **Provides 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**

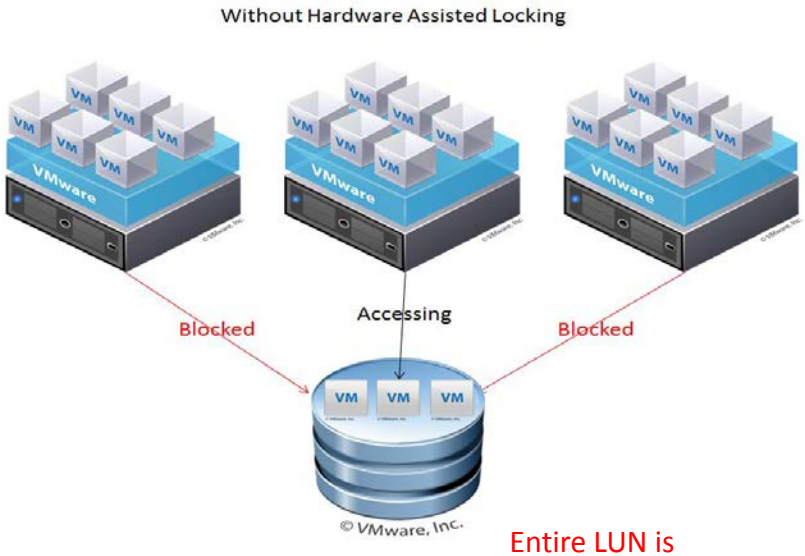

Entire LUN is locked

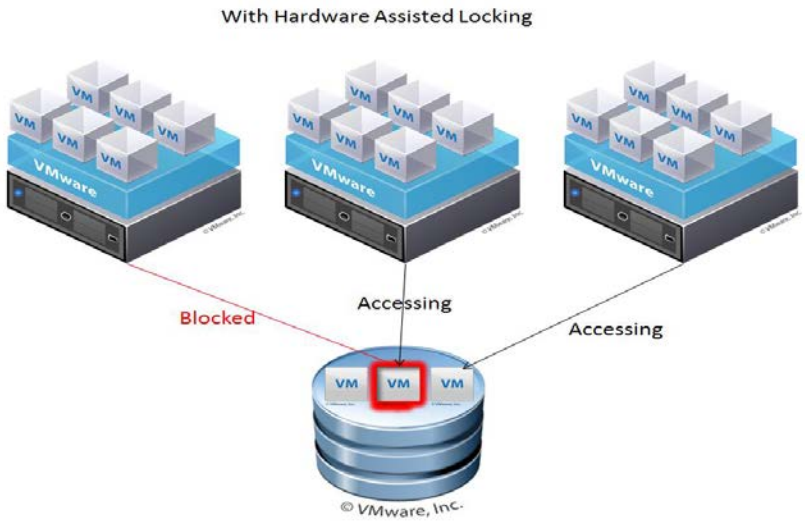

Only the VMDK files are locked

Without VAAI

- Reserves the complete LUN so that it could obtain a lock and required several SCSI commands
- LUN level locks affect adjacent

With VAAI

- Locks occur at a block level with one efficient SCSI command: **SCSI Compare and Swap (CAS)**
- Block level locks have no effect on adjacent hosts

VAAI supported

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Date: August, 2014

Hardware acceleration status

- **In vShaper client, the hardware acceleration status can be verified for the storage arrays.**

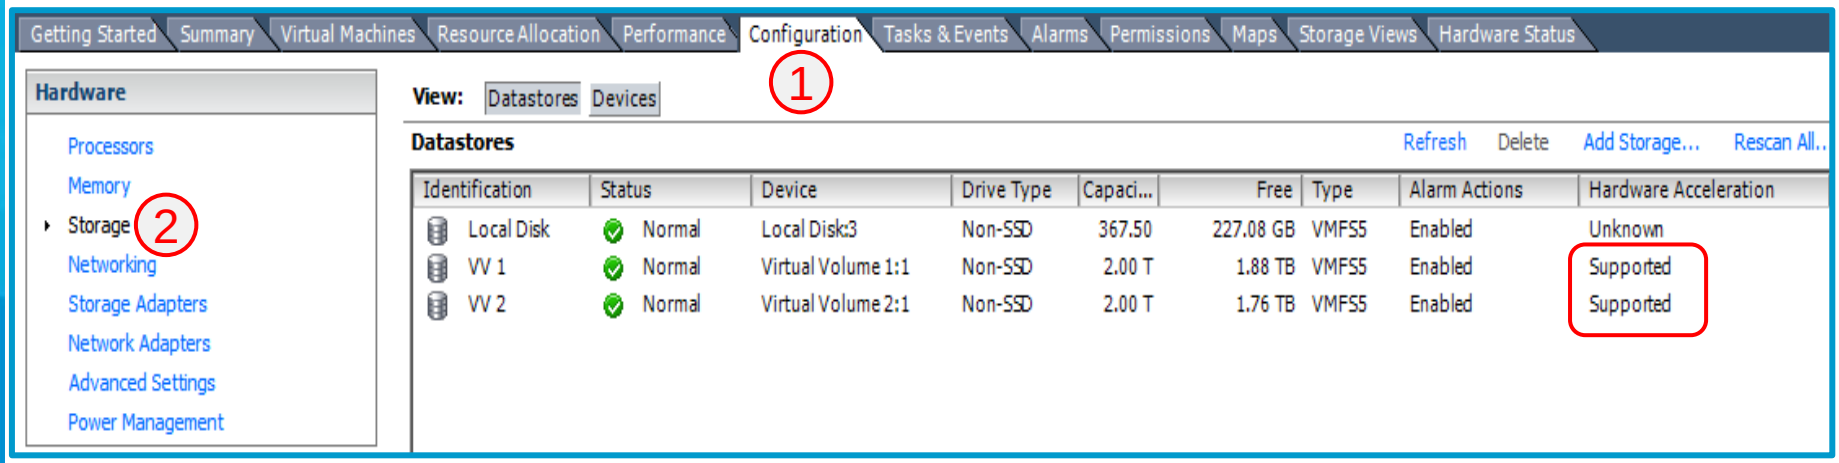

- **There are three kind of hardware acceleration status:**
  - Unknowns: The initial value is Unknown.
  - Supported: The status changes to Supported after the host successfully performs the offload basic operations.
  - No supported: If the offload operation fails, the status changes to Not Supported.

Storageflex
NETWORK DATA STORAGE SOLUTIONS

VAAI enabled

To determine if VAAI is enabled using the vSphere Client:

- In the vSphere Client inventory panel, select the host.
- Click the Configuration tab, and click Advanced Settings under Software.
- Check that these options are set to 1 (enabled):

  DataMover -> HardwareAcceleratedMove -> 1
  DataMover -> HardwareAcceleratedInit -> 1
  VMFS3 -> HardwareAcceleratedLocking -> 1

  Note: These options are enabled by default.

Enabling full copy and block zero.

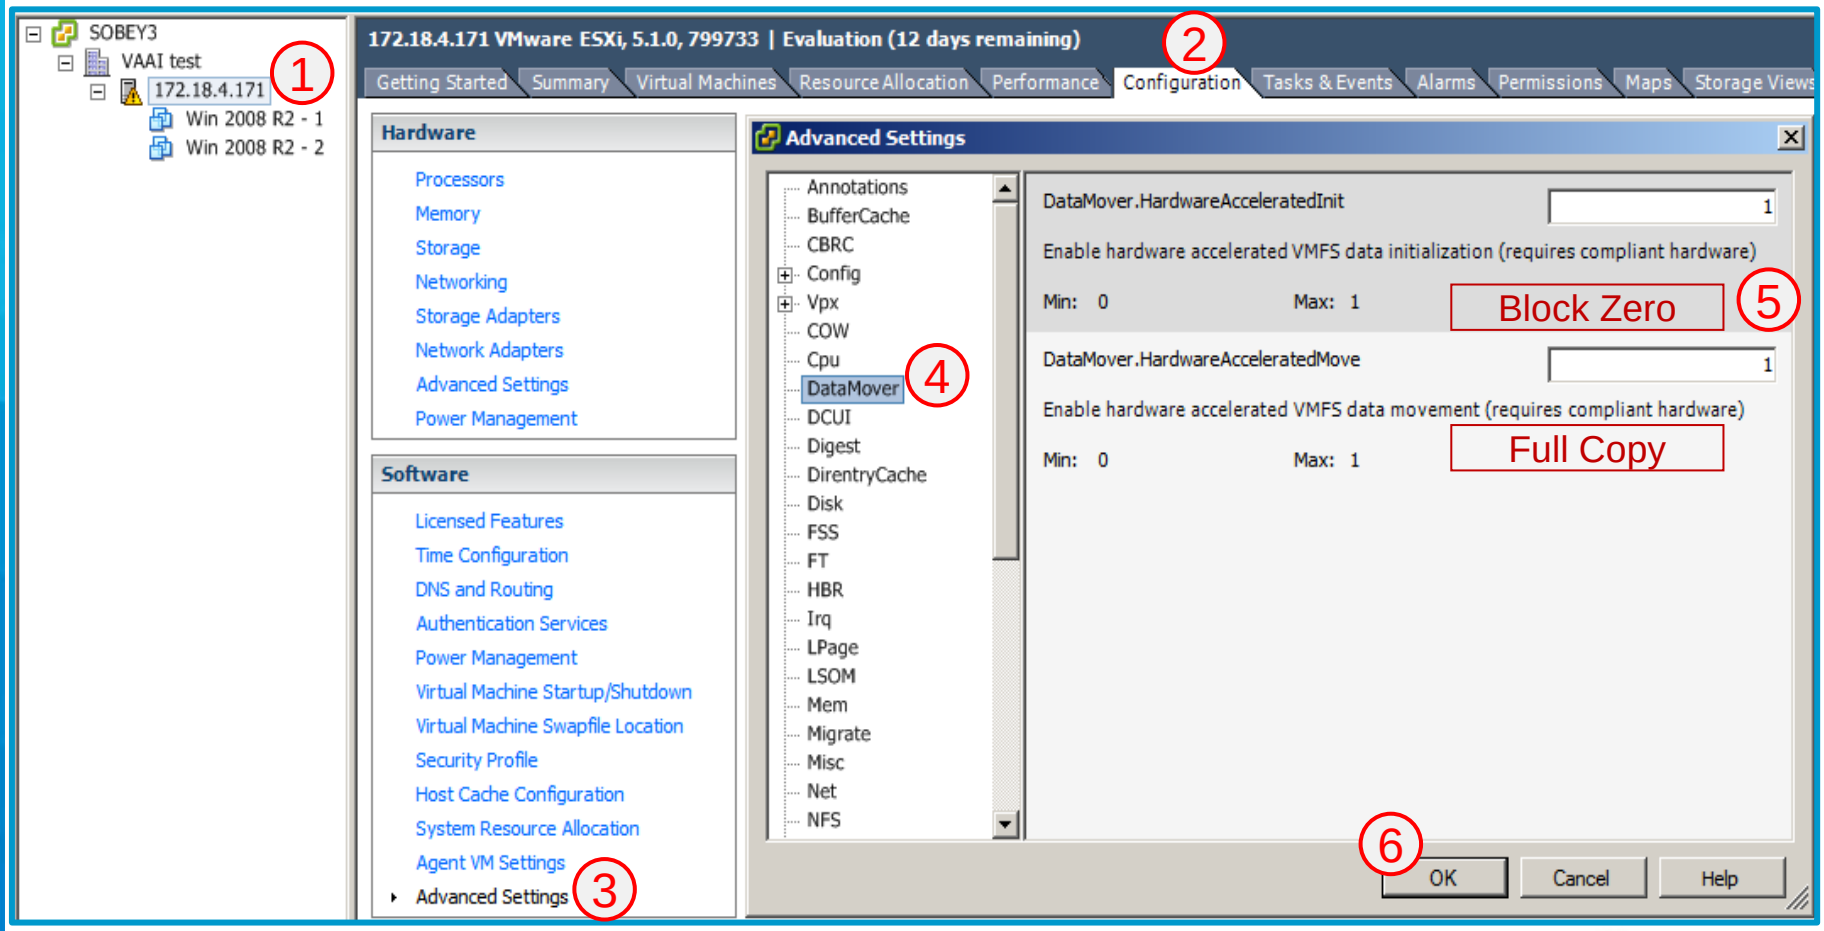

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Enabling hardware-assisted locking.

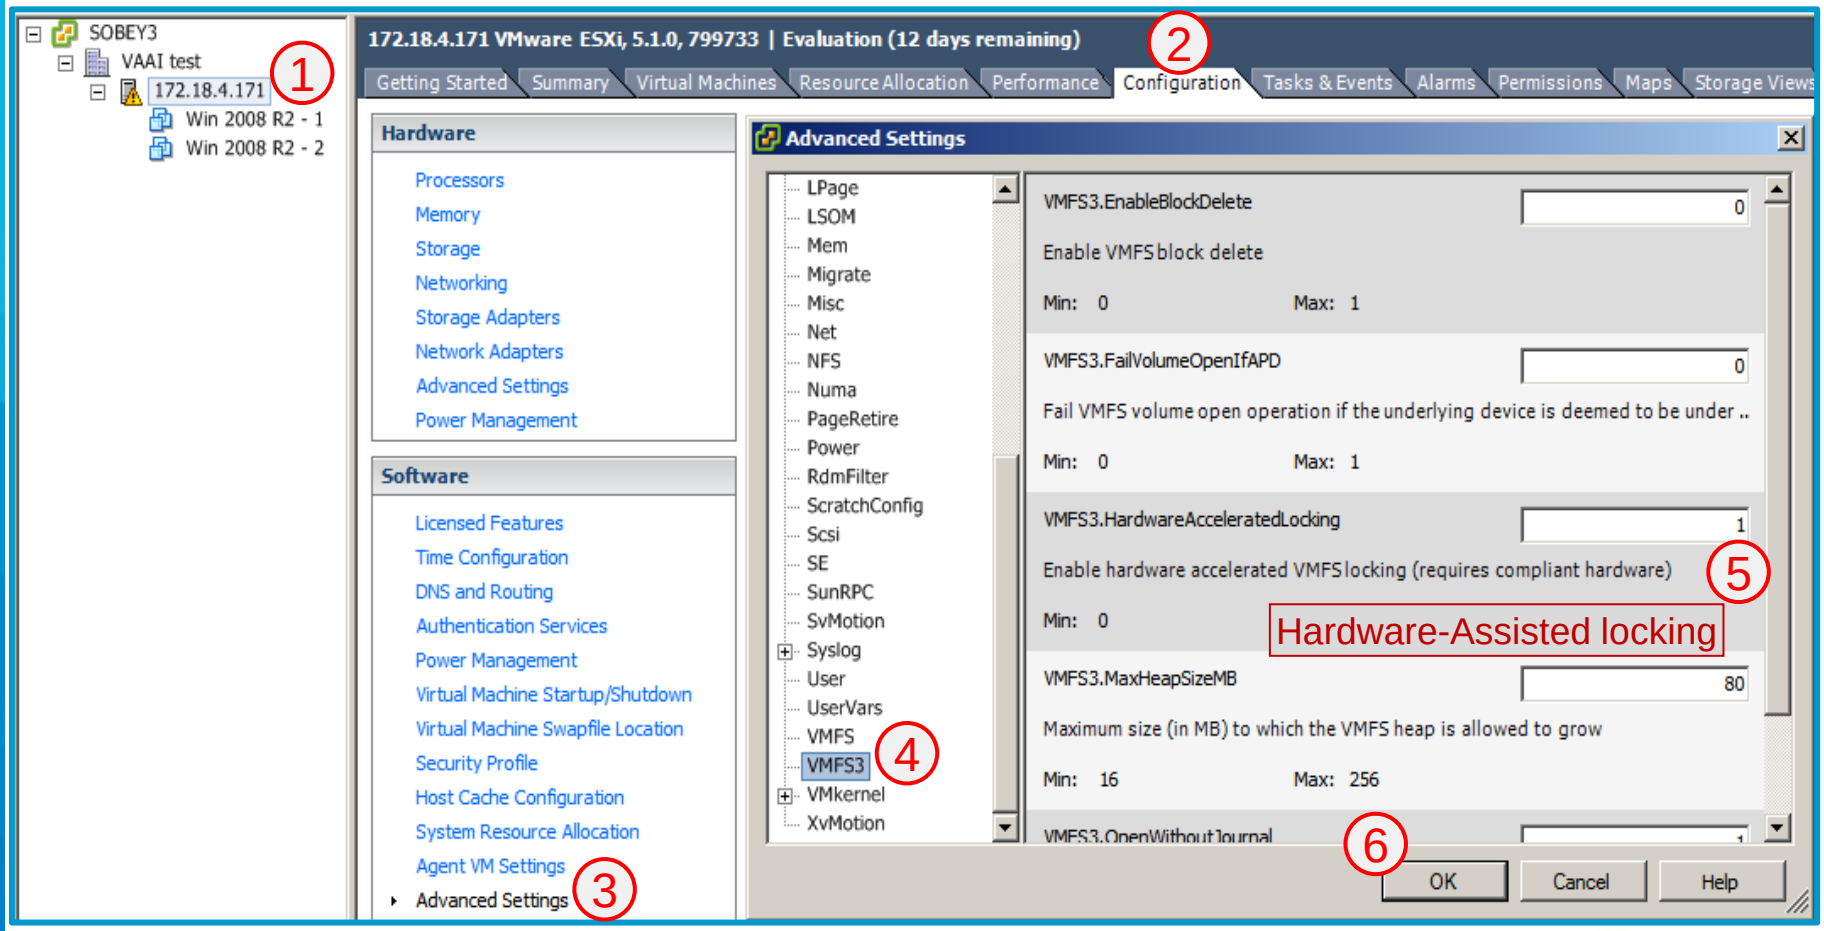

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Test physical environment

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Date: August, 2014

Network architecture

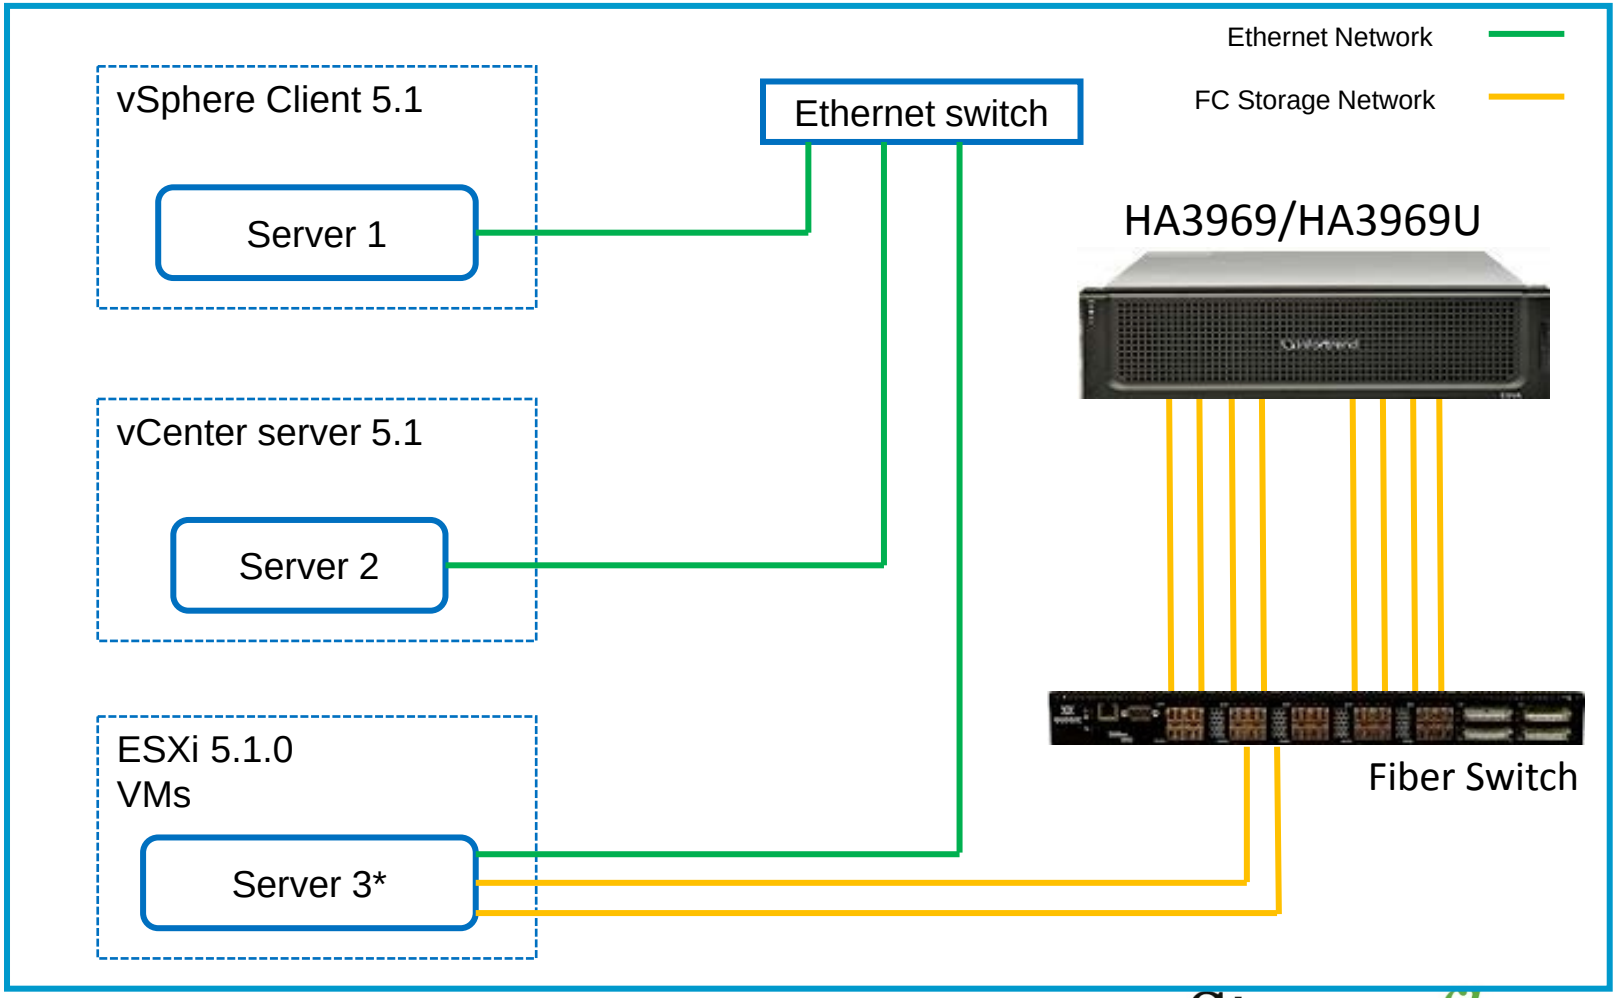

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Hardware and software resources

Item	Resource	Detail	Configuration
1	Hardware	HA3969	HDD SAS, 450 GB, 1500 RPM x 16
	Software	NA	Firmware: 3.91C0.8
2	Hardware	Server 1	Memory: 12 GB of RAM CPU: Intel Core i7 3770K CPU 3.50GHz (8 CPUs), ~ 3.9GHz
	Software	Windows server 2008 R2	vSphere Client 5.1, SANWatch 2.4.c.01
3	Hardware	Server 2	Memory: 8 GB of RAM CPU: Intel Pentium III, Xenon (4 CPUs), ~ 2.3GHz
	Software	Window server 2003	vCenter Server 5.1
4	Hardware	Server 3 - Supermicro X7DB8	Memory: 16 GB of RAM CPU: 2 x Intel Xeon E5405 CPU 2.00GHz
	Software	ESXi 5.1.0	Virtual Machines (Window Server 2008 R2 64-bits)

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Storage architecture

HA3969/HA3969U

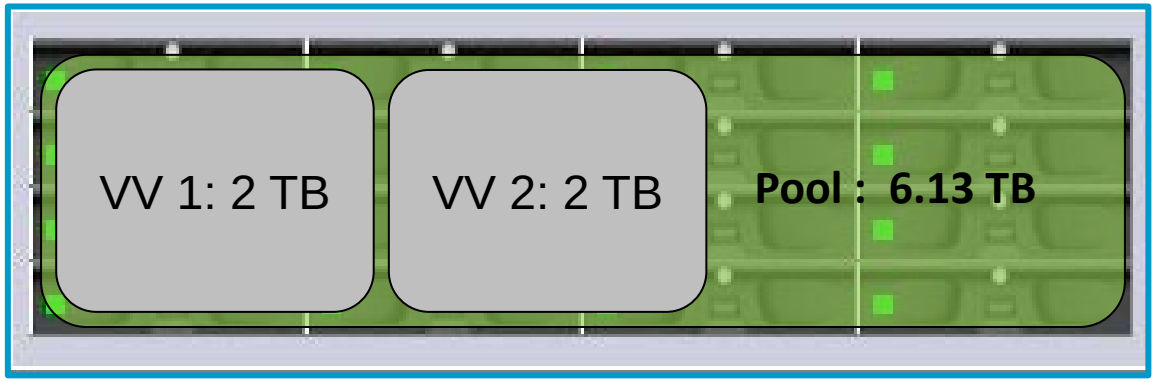

JBOD

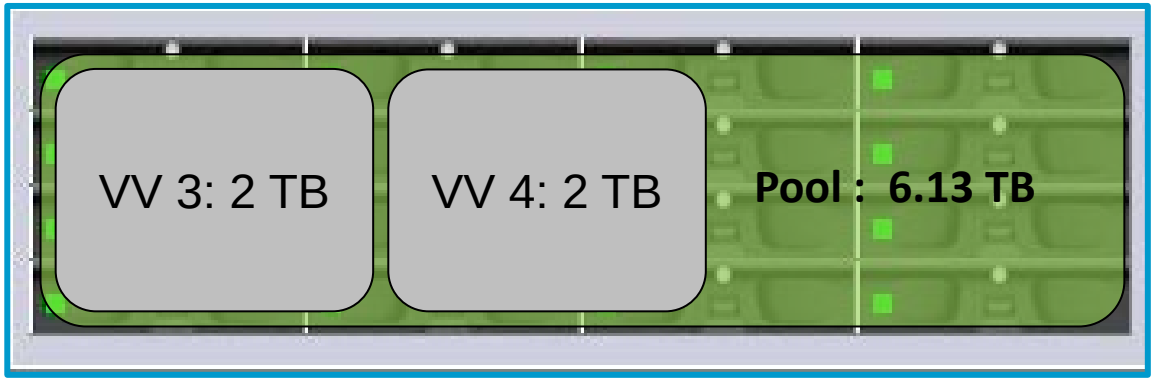

HDD SAS, 450 GB, 1500 RPM **x 32**

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Testing and results

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Date: August, 2014

Full Copy

Steps

1. Create a 150 GB virtual Machine in VV 1 (Think Provision Lazy Zeroed) , the virtual machine was filled with data equal to 63.9 GB.
2. Clone the virtual machine from VV 1 to VV 3. (Different pools)
3. Measure Time, percent of CPU Utilization and Network traffic.
4. Test with VAAI Disabled and VAAI Enabled.

Item	VAAI Disabled	VAAI Enabled
Time	96 Sec	78 Sec
CPU %	5.09 %	0.45 %
Commands / Sec	11651	305

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Full Copy - results

Key finding:

- Cloning a Virtual Machine is 18.75 % faster using VAAI.
- The percent of CPU utilization decrease in 91.16 %.
- The decrease of network traffic was 97.38 % reduced.

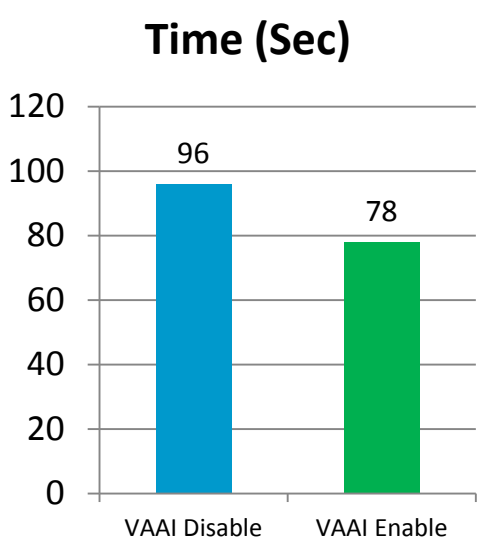

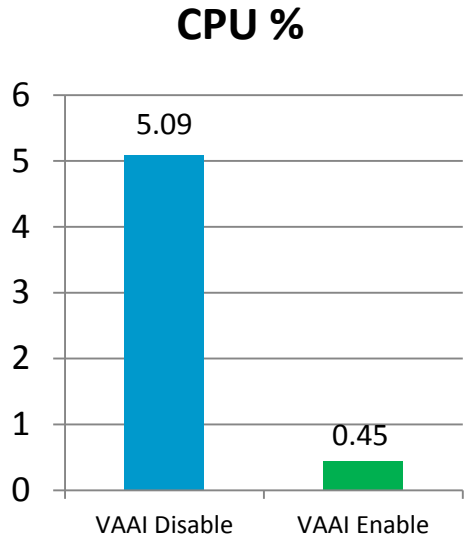

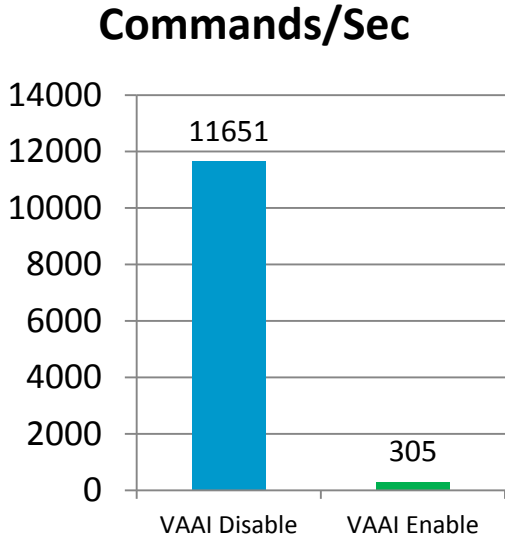

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Full Copy – Bonus results

Bonus Key finding:

- Cloning 10 VM Virtual Machine at the same time is 40.24 % faster using VAAI.

Item	VAAI Disabled	VAAI Enabled
Time	107 Sec	85 Sec

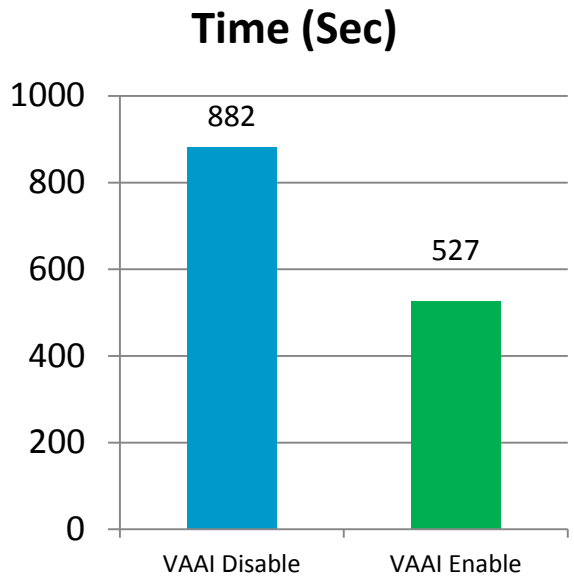

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Block Zeroing

Steps

1. Creating 100 GB virtual machine (Think Provision Eager Zeroed).
2. Measure Time and Network traffic.
3. Test with VAAI Disabled and VAAI Enabled.

Item	VAAI Disabled	VAAI Enabled
Time	107 Sec	85 Sec
Writes / Sec	1069	154

Block Zeroing - results

Key finding:

- Creating a Virtual Machine is 20.56 % faster using VAAI.
- The decrease of network traffic was 85.59 % reduced.

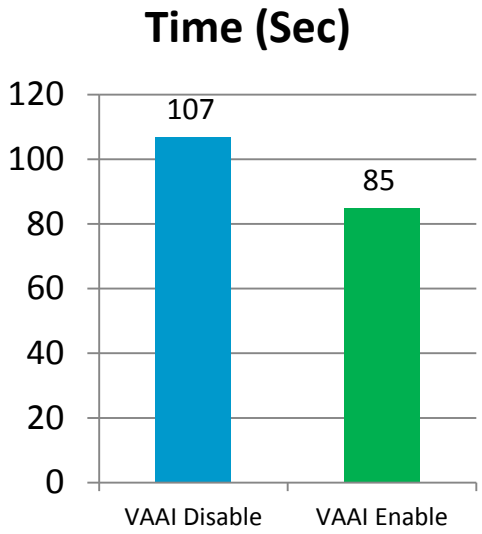

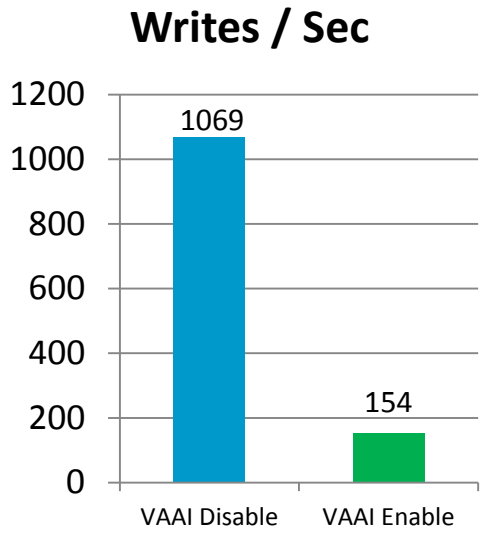

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Hardware Assisted Locking

Steps

1. Create 100 VM over 4 ESXi servers (25 in each one).
2. Measure Time to power on all the 100 VM at the same time.
3. Test with VAAI Disabled and VAAI Enabled.

Item	VAAI Disabled	VAAI Enabled
Time to power ON	97 Sec	97 Sec

Hardware Assisted Locking - results

Key finding:

- Powering on multiple VMs at the same time reduce the bottleneck caused by LUN locking and is 10 % faster using VAAI.

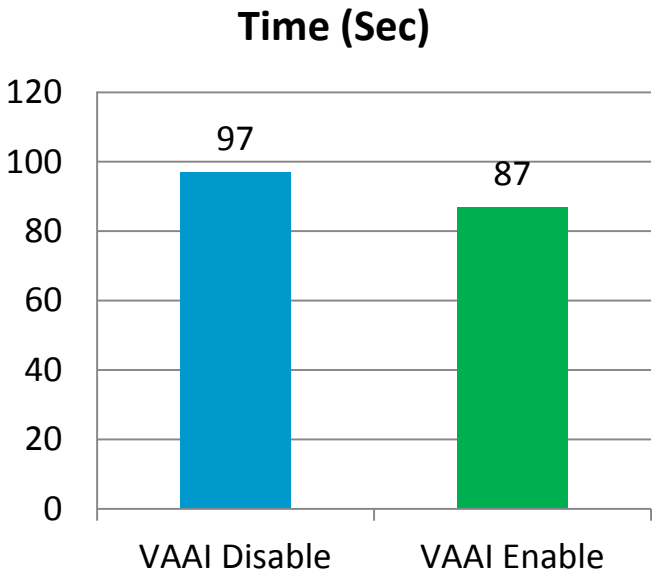

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Summary

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Date: August, 2014

Summary

Combining HA3969/HA3969U and VMware vShpere 5.1 with VAAI achieves an exceptional performance in virtual environments.

The key benefits of using VAAI is notorious in the high speed processing to complete operations and excellent improvement by reducing the server load and the I/O load.

The test results show the perfect coupling and operation of HA3969/HA3969U with VAAI.

- **Full Copy**
  - **19% faster**
  - **91% CPU load reduce**
  - **97% commands traffic reduce**
- **Block Zeroing**
  - **20% faster**
  - **86% Writes commands reduce**
- **Hardware-assisted locking**
  - **10% faster**
  - **Avoid IO bottlenecks caused by LUN locking**

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Thanks !

Storageflex
NETWORK DATA STORAGE SOLUTIONS

Date: August, 2014