



VMware Storage Offload

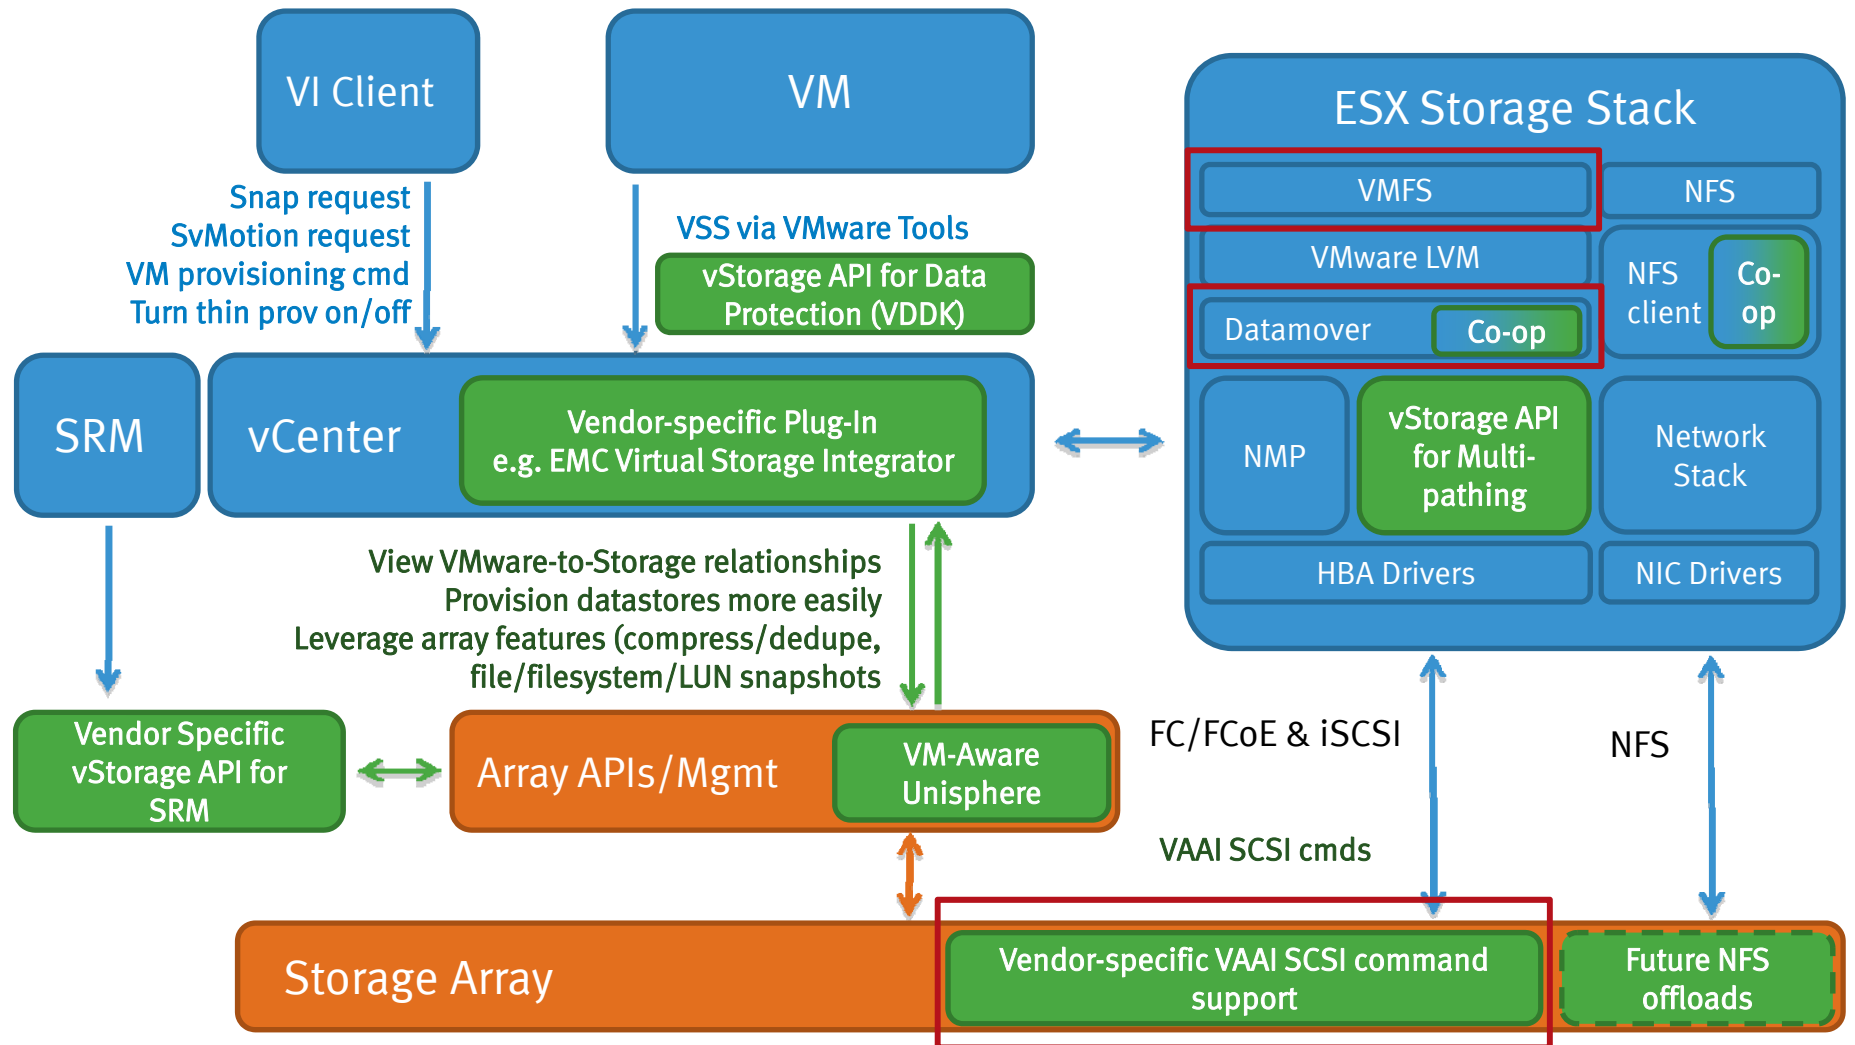
vStorage APIs for Array Integration (VAAI) –
what they are, and what to expect

Ed Walsh
vSpecialist
Walsh_ed@emc.com
518-928-0694

A couple of things to set the stage...

- vStorage APIs are a “family”
- The “family” as of vSphere 4.0 has 3 “siblings”:
 1. vStorage API for Site Recovery Manager
 2. vStorage API for Data Protection
 3. vStorage API for Multipathing
- Future VMware technologies expand the “family” of vStorage APIs
 4. vStorage API for Array Integration (VAAI)
- VMware is not stopping anytime soon....

“Where Does Storage Integration Happen?”



5 key things to know about VAAI

1. A set of APIs that allows future vSphere releases to offload specific operations to block array targets
2. supported on VMware File Systems (VMFS) and Raw Device Mappings (RDM) volumes & backwards compatible
3. Improves many things by having the vmkernel offload operations to more efficient array-based hardware
4. Implemented via standard T10 SCSI commands – open standard
5. Supported with EMC Unified and EMC VMAX arrays

3 Simple Pre-Requisites

- vSphere 4.1 release
 - Configured “on” by default
 - If array errors on advanced command, falls back to “heavy lifting” legacy datamover software codepath
- VAAI-compliant array firmware
 - EMC Unified – FLARE 30
 - VMAX – Enginuity 5875
- Specific configuration considerations
 - On EMC Unified – ALUA mode (failover=4) must be configured

Q: Why no support for older platforms?

A: Older gen (DMX, CX3) was a major architectural transition (from 32-bit to 64-bit in CLARiiON, from PowerPC to Intel on Symmetrix)

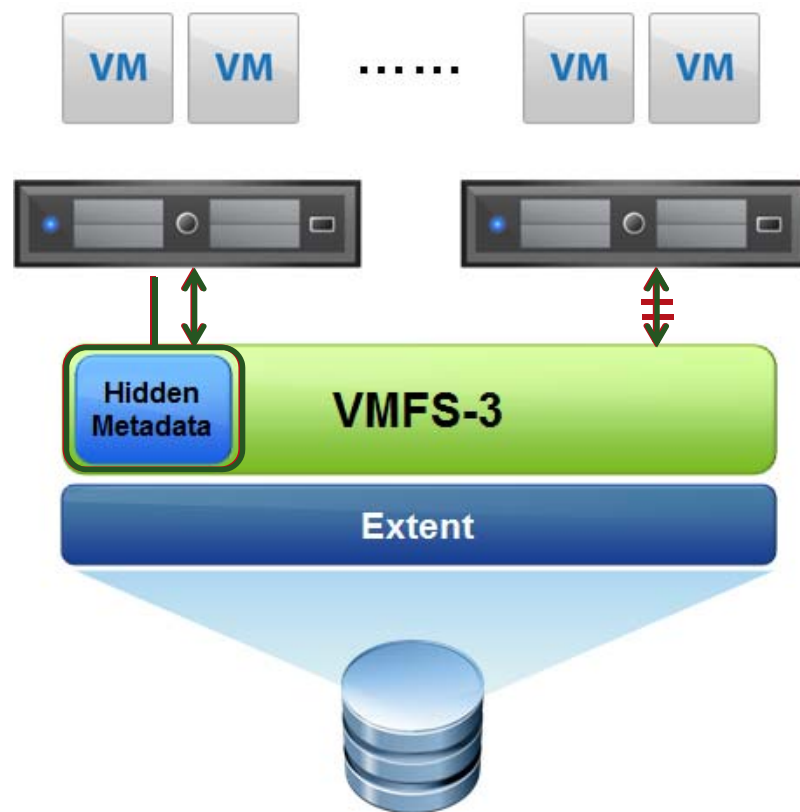
On the whole, expect new EMC coolness to *mostly* be on the platforms released in the last 2 years.

3 APIs being Implemented

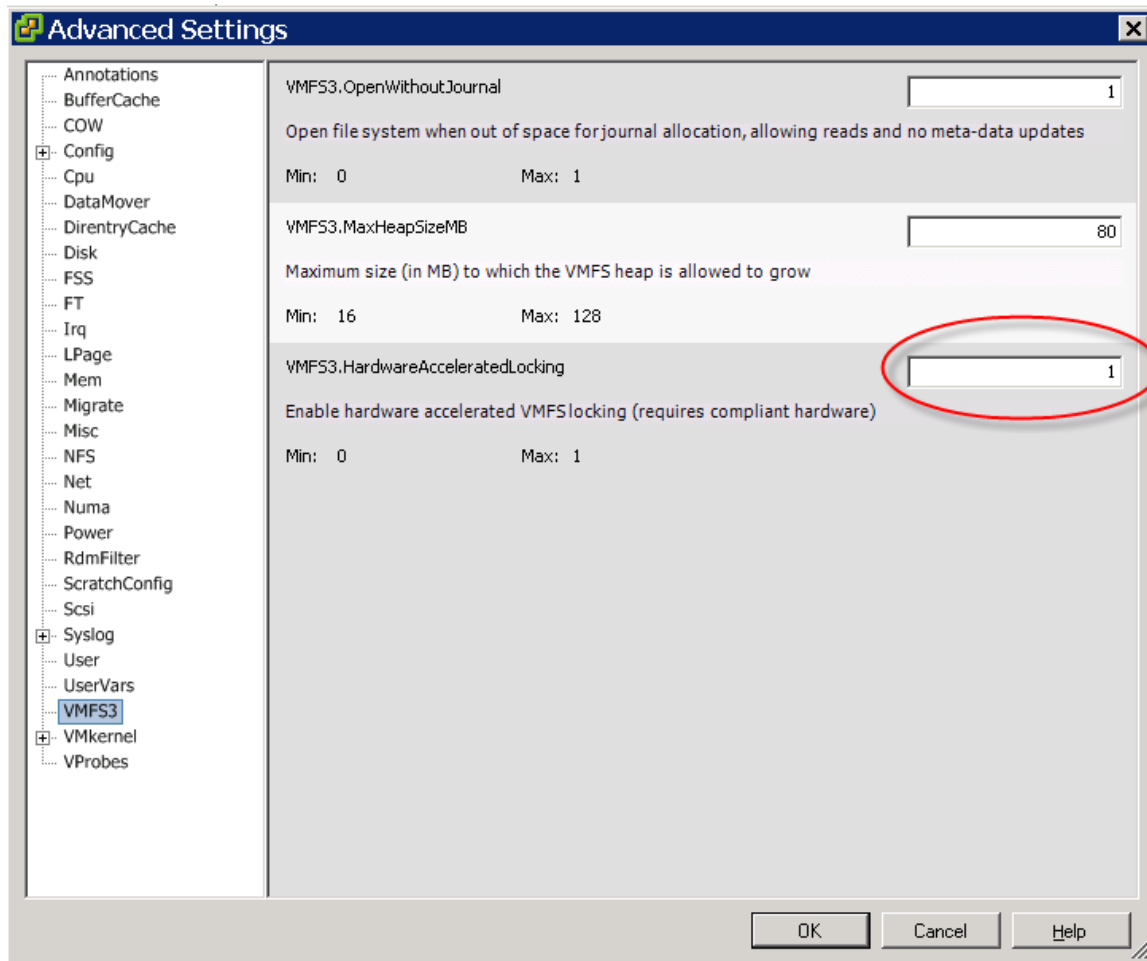
- *Hardware-Accelerated Locking = 10-100x better metadata scaling*
 - Replace LUN locking with extent-based locks for better granularity
 - Reduce number of “lock” operations required by using one efficient SCSI command to perform pre-lock, lock, and post-lock operations
 - Increases locking efficiency by an order of magnitude
- *Hardware-Accelerated Zero = 2-10x lower number of IO operations*
 - Eliminate redundant and repetitive host-based write commands with optimized internal array commands
- *Hardware-Accelerated Copy = 2-10x better data movement*
 - Leverage native array Copy capability to move blocks

Hardware-Accelerated Locking

- Without API
 - Reserves the complete LUN so that it could obtain a lock
 - Required several SCSI commands
 - LUN level locks affect adjacent hosts
- With API
 - Locks occur at a block level
 - One efficient SCSI command - SCSI Compare and Swap (CAS)
 - Block level locks have no effect on adjacent hosts.
- Use Cases
 - Bigger clusters with more VMs
 - View, Lab Manager, Project Redwood
 - More & Faster VM Snapshotting



Future ESX Parameter for Hardware-Accelerated Locking

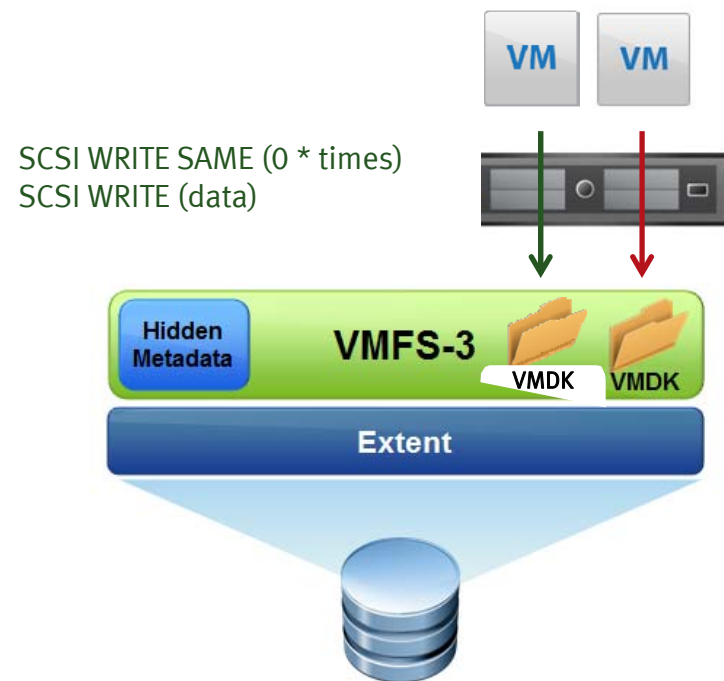


Q: Can this be changed without a reboot?

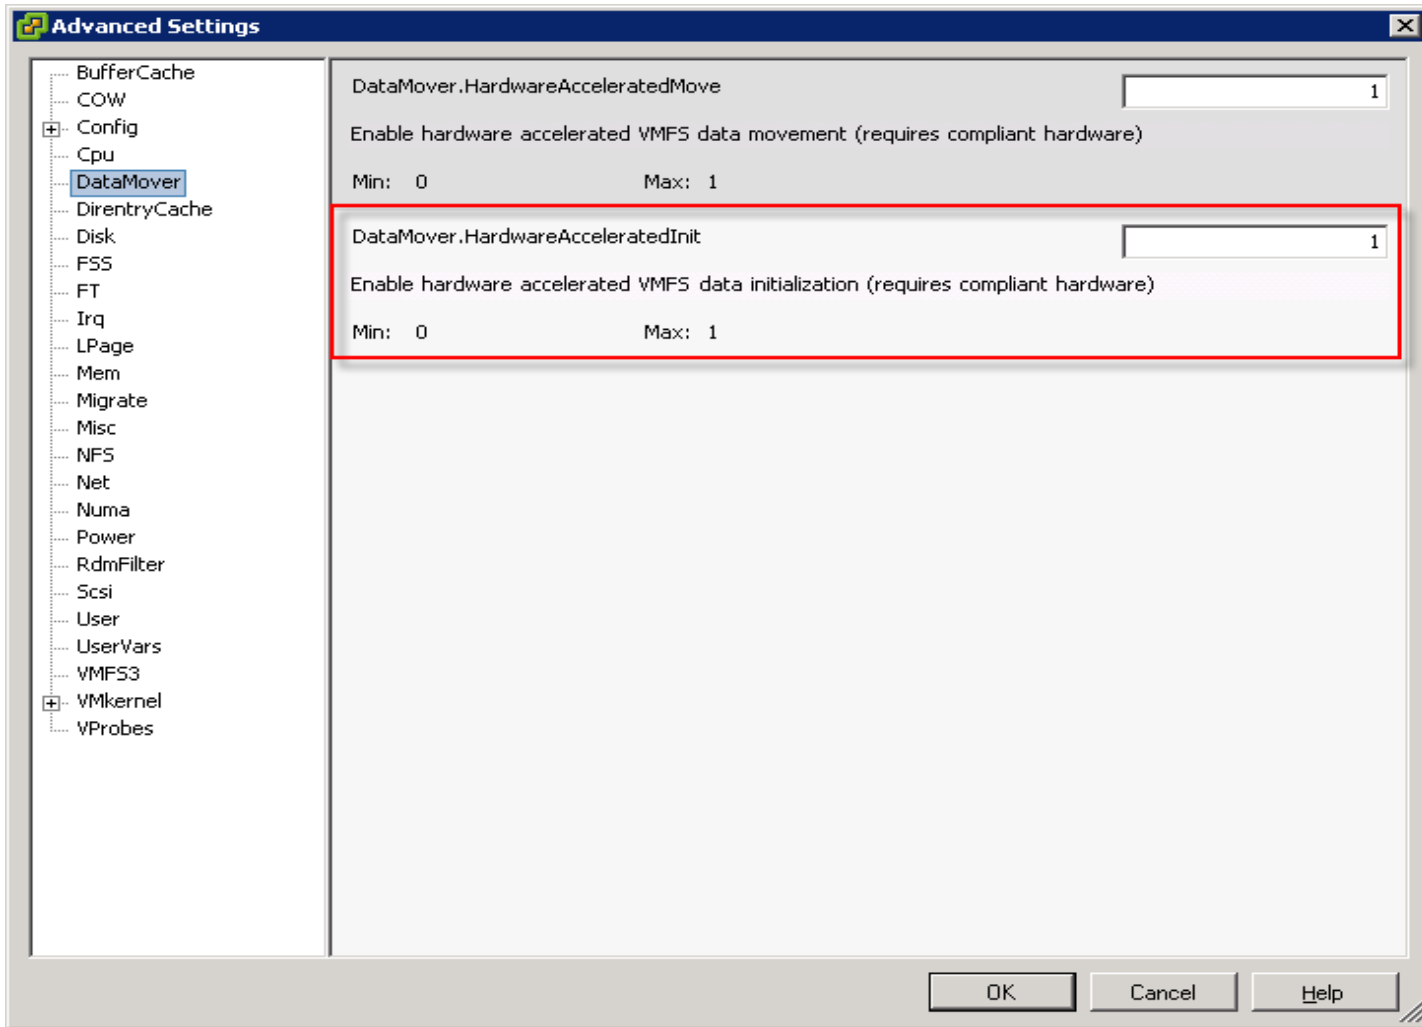
A: Yes!

Hardware-Accelerated Zero

- Without API
 - SCSI Write - Many identical small blocks of zeroes moved from host to array for MANY VMware IO operations
 - Extra zeros can be removed by EMC arrays after the fact by “zero reclaim”
 - New Guest IO to VMDK is “pre-zeroed”
- With API
 - SCSI Write Same - One giant block of zeroes moved from host to array and repeatedly written
 - Thin provisioned array skips zero completely (pre “zero reclaim”)
- Use Cases
 - Reduced IO when writing to new blocks in the VMDK for any VM
 - Time to create VMs (particularly FT-enabled VMs)



Future ESX Parameter for Hardware-Accelerated Zero

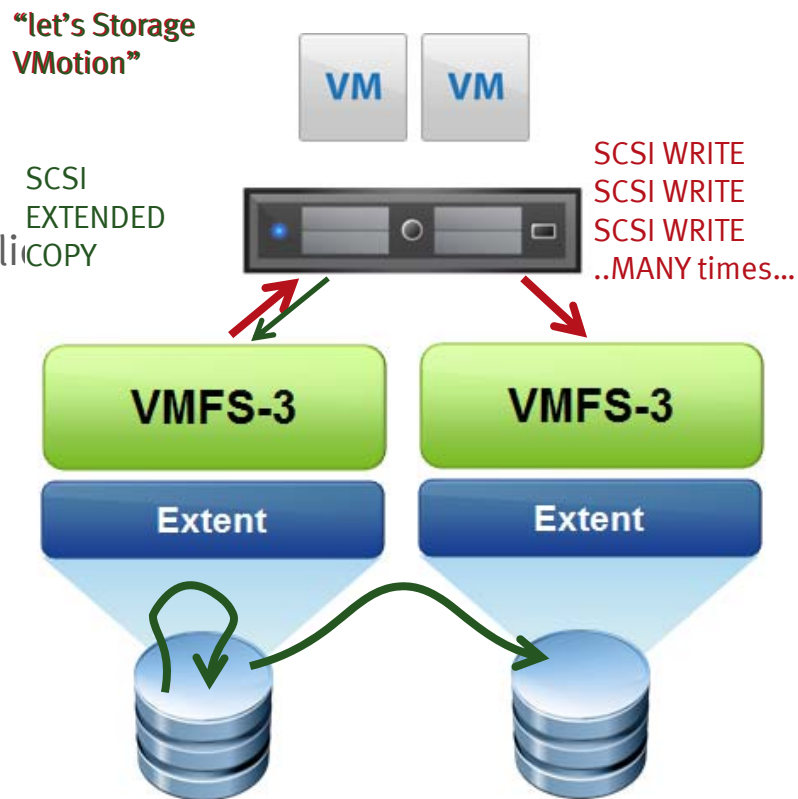


Q: Can this be changed without a reboot?

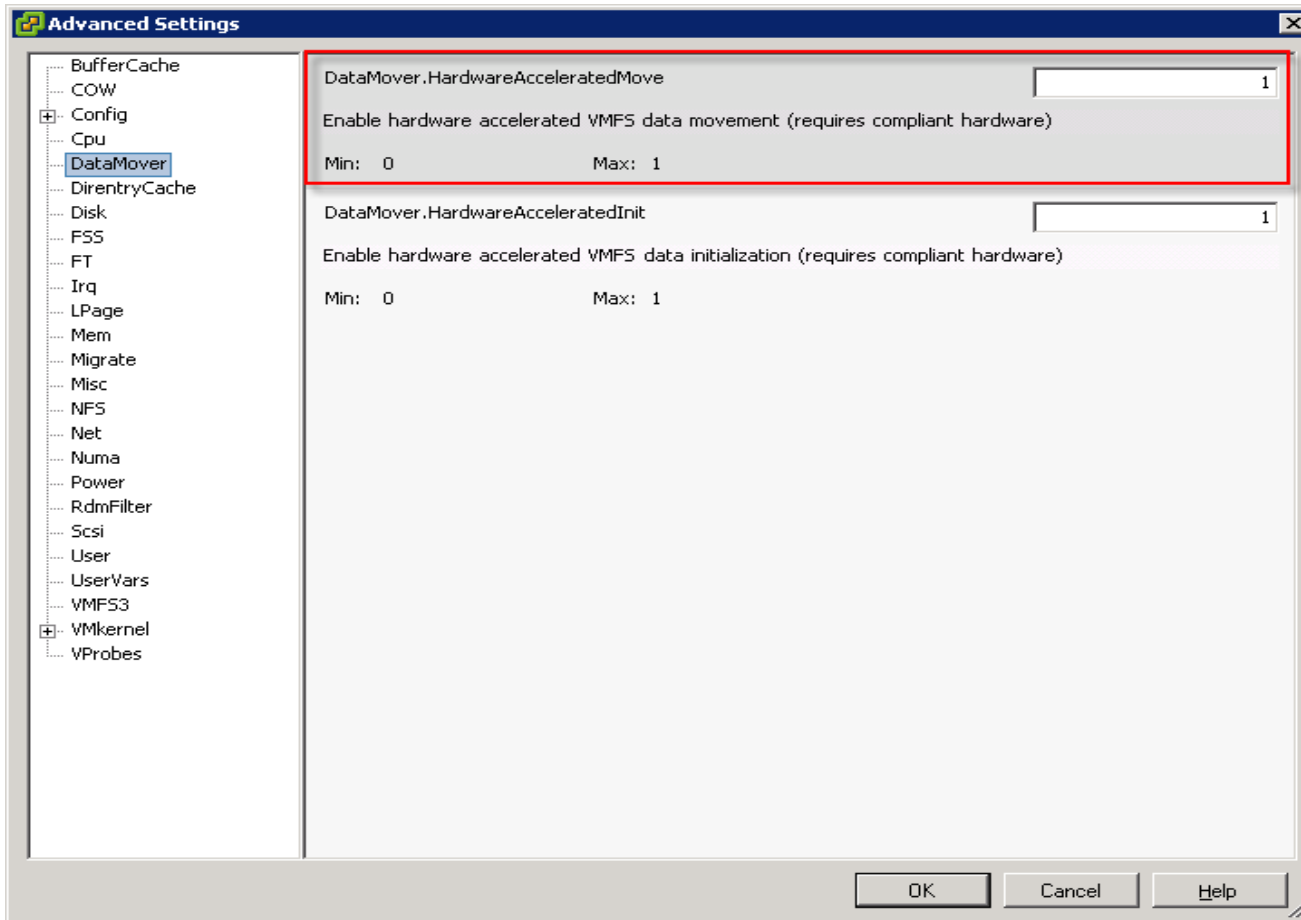
A: Yes!

Hardware-Accelerated Copy

- Without API
 - 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 milliCOPY small block operations
- With API
 - SCSI Extended Copy (Data moved within array)
 - Repeat
 - Order of magnitude reduction in IO operations
 - Order of magnitude reduction in array IOps
- Use Cases
 - Storage VMotion
 - VM Creation from Template



Future ESX Parameter for Hardware-Accelerated Copy



Q: Can this be changed without a reboot?

A: Yes!

Demo of Hardware Accelerated Copy effect

Recap - Customer Benefits of VAAI

API Used	Use Case	Benefit in a nutshell
Hardware Accelerated Locking	<ul style="list-style-type: none">• Big clusters, lots of VMs• Lots of VMware snapshots• VMware View, Lab Manager, Project Redwood	
Hardware Accelerated Zero	<ul style="list-style-type: none">• Reduction in IO for almost all VMware operations	
Hardware Accelerated Copy	<ul style="list-style-type: none">• Faster storage Storage vMotion and Clone/Template• Huge reduction in ESX, array, and network load	

...what Intel VT did for compute in vSphere 4...

...VAAI-enabled EMC arrays do for storage in vSphere.ahem

EMC: The #1 Choice for Virtualized Environments

Why do twice as many customers pick EMC for their VMware environments?*

Best technology
+
Best people
+
Best solutions

EMC² + vmware[®] = Zero Risk

*Source: Goldman Sachs IT Spending Survey, January 7, 2010



Q&A

Thank You!

- Join us every Thursday at 11:00 AM ET for our EMC Solutions for VMware Webcast Series
- Virtual Geek blog: <http://virtualgeek.typepad.com/>
- EMC Solutions for VMware landing page: <http://www.emc.com/themes/nextstartsnow/emc-solutions-vmware/index.htm>
- Everything VMware at EMC Community: <http://www.emc.com/vmwarecommunity>

Would you like to see specific topics in this webcast series?

Let us know ... vmwaresolutions@emc.com

See You at VMworld!

vmworld 2010

August 30-September 2
Moscone Center, San Francisco

October 12-14
Bella Center, Copenhagen

Register Now and Save!

San Francisco



\$250 Savings

Register by June 25

August 30 - September 2

Moscone Center
San Francisco

[Register Now](#)

Copenhagen



€200 Savings

Register by July 23

October 12 - 14

Bella Center
Copenhagen

[Register Now](#)

VMworld 2010, hosted by VMware, gives you unparalleled access to educational opportunities that will show you how to reduce IT complexity, and enable IT as a service through virtualization and cloud computing. Join us to...

Learn

Select from Breakout Sessions and Hands on Labs that cover topics such as leveraging the public cloud, virtualizing enterprise applications and managing desktops as a service.

Engage

Meet one on one with Knowledge Experts or join group discussions to learn and share with peers and industry experts.

Network

Visit industry-leading technology partners showcasing complementary solutions that can help you dramatically reduce your IT complexity.

Social Media & Community

- VMworld on Twitter
- VMworld on Facebook
- VMworld TV on YouTube
- VMworld Discussions
- VMworld Conference Blog

Tell us what you think

What would you like to hear Paul Maritz talk about in General Session?

[Submit](#)

Sponsors & Exhibitors

Get Early Bird registration ... at any time!

**U.S. Event Code:
EMC_EB_US2010**

**EMEA Event Code:
EMC_EB_EMEA2010**

Register: www.vmworld.com/registration.jspa?sponsorCode=EMC