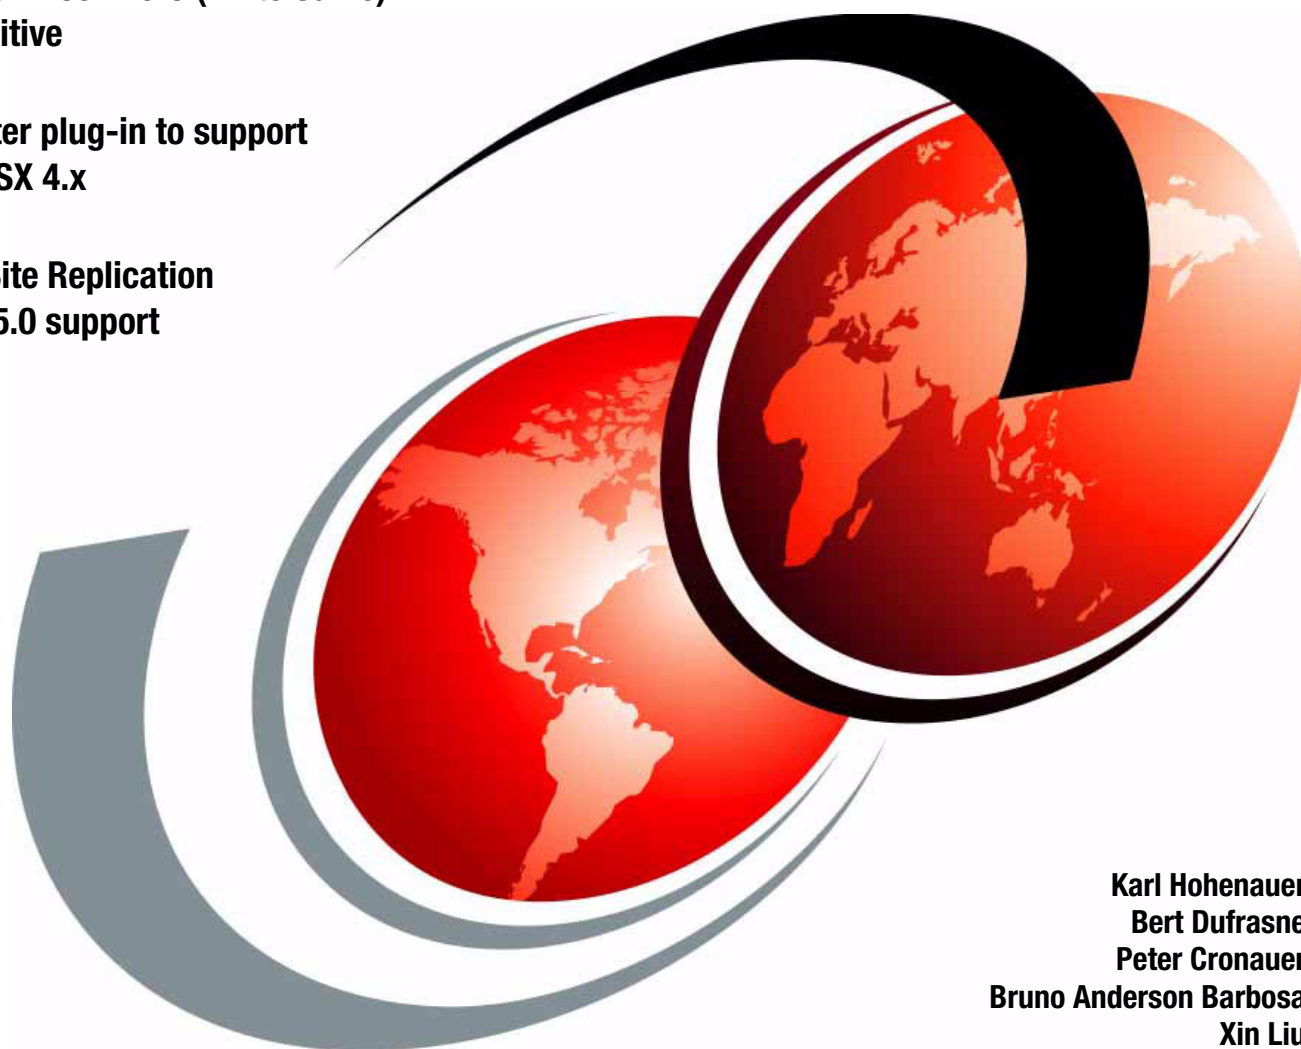


# IBM DS8870 and VMware Synergy

Support for Block Zero (Write Same)  
VAAI primitive

IBM vCenter plug-in to support  
VAAI on ESX 4.x

VMware Site Replication  
Manager 5.0 support



Karl Hohenauer  
Bert Dufrasne  
Peter Cronauer  
Bruno Anderson Barbosa  
Xin Liu





International Technical Support Organization

**IBM DS8870 and VMware Synergy**

August 2013

**Note:** Before using this information and the product it supports, read the information in “Notices” on page v.

**Second Edition (August 2013)**

This edition applies to DS8000 Licensed Machine Code (LMC) 7.7.10.xx.xx (bundle version 87.10.xxx.xx), or later.

This document was created or updated on August 19, 2013.

© Copyright International Business Machines Corporation 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Notices</b> .....	v
Trademarks .....	vi
<b>Preface</b> .....	vii
Authors .....	vii
Now you can become a published author, too! .....	ix
Comments welcome .....	ix
Stay connected to IBM Redbooks .....	ix
<b>Summary of changes</b> .....	xi
August 2013, Second Edition .....	xi
<b>Chapter 1. VMWare introduction</b> .....	1
1.1 Introduction .....	2
1.2 DS8870 overview .....	2
1.3 VMware architecture overview .....	4
1.3.1 VMware vSphere features .....	4
1.3.2 VMware storage concept .....	8
1.3.3 What's new in VMware vSphere for storage .....	9
<b>Chapter 2. VAAI support</b> .....	13
2.1 Introduction .....	14
2.2 VMware vStorage APIs array integration .....	14
2.2.1 vStorage APIs for array integration primitives .....	14
2.2.2 DS8870 support for VAAI .....	16
2.3 Using vStorage APIs for array integration .....	16
2.3.1 Installing the IBM VAAI device driver on an ESXi 4.1 server .....	17
2.3.2 Validating VAAI global settings on the host .....	17
2.3.3 Confirming that VAAI Hardware Acceleration is detected for a volume .....	19
2.3.4 What to do if the Hardware Acceleration status shows as Unknown .....	21
2.4 Testing vStorage APIs for Array Integration .....	21
2.4.1 Atomic Test and Set test .....	22
2.4.2 XCOPY test .....	24
2.4.3 Write Same Test .....	27
<b>Chapter 3. IBM Storage Management Console for VMware vCenter</b> .....	29
3.1 Introduction .....	30
3.2 Installation .....	31
3.3 Adding IBM Storage to the plug-in .....	33
3.4 Manage DS8870 volume properties .....	36
3.5 Creating a datastore .....	37
3.6 Multipathing .....	38
3.7 Host attachment .....	39
<b>Chapter 4. Business continuity with IBM DS8870 Storage System and VMware Site Recovery Manager</b> .....	41
4.1 Introduction to business continuity and disaster recovery solutions .....	42
4.2 VMware Site Recovery Manager .....	42
4.2.1 VMware SRM 5.x Disaster Recovery enhancements .....	43

4.3 Disaster Recovery Solution with DS8870 and SRM 5.x . . . . .	44
4.3.1 DS8870 and SRM Solution Architectural Overview . . . . .	44
4.3.2 IBM DS8870 Support for VMware SRM 5.x . . . . .	45
4.3.3 DS8870 Remote Mirror and Copy functions . . . . .	46
4.3.4 DS8870 Remote Mirror and Copy's Consistency Groups Support . . . . .	47
4.3.5 References for VMware SRM implementation with DS8870 systems . . . . .	48
<b>Related publications</b> . . . . .	49
IBM Redbooks . . . . .	49
Other publications . . . . .	49
Online resources . . . . .	50
Help from IBM . . . . .	50

# Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.*

**The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:** INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.


## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

DS8000®	Power Systems™	Storwize®
Easy Tier®	POWER7®	System Storage®
FlashCopy®	Redbooks®	Tivoli®
Global Technology Services®	Redpaper™	XIV®
IBM®	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.



# Preface

This IBM® Redpaper™ publication describes the support of the IBM System Storage® DS8870 for VMware vSphere in general terms. It also describes in greater detail the VMware vSphere Storage APIs Array Integration (VAAI) primitives, the IBM Storage Management Console for VMware plug-in, and the IBM Site Replication Adapter (SRA) for VMware Site Recovery Manager (SRM).

In addition to a high-level overview of the VMware vSphere suite, this paper emphasizes the interaction of the software components with the IBM System Storage DS8870. It is intended for IT professionals who want an understanding of advantages such as off loading specific tasks to storage arrays, centralized service management, and simplified implementation of disaster recovery processes in virtualized environments.

## Authors

This paper was produced by a team of specialists from around the world working for the International Technical Support Organization (ITSO) at the European Storage Competence Center (ESCC) in Mainz (Germany).

**Karl Hohenauer** is an IT Specialist with the IBM Systems and Technology Group, based in Vienna, Austria. He is a member of the ATS System Storage Europe team working for the European Storage Competence Center (ESCC) in Mainz, Germany. Karl joined IBM in 1995 and worked in different pre-sales and post-sales positions. He has co-authored three IBM Redbooks publications. His main focus is on IBM Storage Virtualization and VMware.

**Bert Dufrasne** is an IBM Certified IT Specialist and Project Leader for System Storage disk products at the ITSO, San Jose Center. He has worked at IBM in various IT areas. He has written many IBM Redbooks® publications and developed and taught technical workshops. Before joining the ITSO, he worked for IBM Global Services as an Application Architect. He holds a Masters degree in Electrical Engineering.

**Peter Cronauer** is a certified SAP consultant who works for IBM in the European Storage Competence Center (ESCC) in Germany. He joined the IBM Advanced Technical Skills (ATS) department and implemented and managed client projects for SAP solutions on IBM Storage. He managed the STG Lab Services storage for Europe and the ATS development support department. Peter is the speaker of the Solution Advisory Board (SAB) and has led Redbooks residencies on various storage topics. He holds a degree in Computer Science and wrote white papers about SAP and IBM storage solutions.

**Bruno Anderson Barbosa** is a Senior Software Support Specialist for Global and Technology Services (GTS) in IBM Brazil. He has seven years of experience working with IBM Power Systems™, storage area network (SAN) and IBM Storage Systems. He holds a degree in IT Systems Analysis, a Business Administration postgraduation degree, and an MBA diploma in Information Technology Management from Fundação Getúlio Vargas (FGV) in Brazil. His areas of expertise include Implementation, Software Support, Problem Determination, and Performance Analysis on IBM Storage Products for Open Systems.

**Xin Liu** is a Senior IT Specialist in IBM China. He has seven years of experience in technical support area for IBM system hardware and software products. He joined IBM Global Technology Services® in 2006 as a field software engineer and supported Power Systems and Storage. In 2010, he became a member of storage ATS team, and since then he worked on storage pre-sales support focusing on high-end storage solutions. Xin holds a Master's degree in Electronic Engineering from Tsinghua University, China.

Thanks to the following people for their contributions to this project:

- ▶ Pedro Henrique dos Santos Principeza

IBM Brazil

- ▶ Craig Gordon
- ▶ Peter Kisich
- ▶ Mark Kremkus
- ▶ Troy Lee
- ▶ Allen Marin
- ▶ Jared Minch
- ▶ Allen Wright

IBM US

- ▶ Brian Sherman

IBM Canada

- ▶ Torsten Rothenwald

IBM Germany

Thanks to the following authors of the previous editions of this paper:

- ▶ Peter Kimmel
- ▶ Akin Sakarcan
- ▶ Roland Wolf

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at this website:

<http://www.ibm.com/redbooks/residencies.html>

## Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this paper or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at this website:

<http://www.ibm.com/redbooks>

- ▶ Send your comments in an email to:

[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>



# Summary of changes

This section describes the technical changes that were made in this edition of the paper and in previous editions. This edition also might include minor corrections and editorial changes that are not identified.

Summary of Changes  
for IBM DS8870 and VMware Synergy  
as created or updated on August 19, 2013.

This edition applies to DS8000 Licensed Machine Code (LMC) 7.7.10.xx.xx (bundle version 87.10.xxx.xx), or later.

## August 2013, Second Edition

This revision reflects the following addition, deletion, or modification of new and changed information:

### **New information**

- ▶ Support for Block Zero (Write Same) VAAI primitive
- ▶ Additional IBM VAAI device driver to support VAAI on ESX/ESXi 4.1
- ▶ Support for the IBM Storage Management Console for VMware vCenter
- ▶ VMware Site Recovery Manager 5.0/5.1 support.

### **Changed information**

- ▶ Added a section to show how to disable / enable VAAI primitives on ESX/ESXi 4.1
- ▶ Updated the VMware storage news section to include VMware vSphere 5.1 latest updates
- ▶ Updated the test section to include the new Write Same VAAI primitive on DS8870.





# VMWare introduction

This chapter provides a high-level introduction to the concepts and features of VMware and includes the following topics:

- ▶ Introduction
- ▶ DS8870 overview
- ▶ VMware architecture overview

## 1.1 Introduction

Virtualization technology is transforming business. Companies are increasingly virtualizing their environments to reduce costs and meet the following goals:

- ▶ Consolidate servers
- ▶ Centralize services and management
- ▶ Implement disaster recovery
- ▶ Set up remote or thin client desktops
- ▶ Create clouds for optimized resource use

Organizations often deploy server virtualization to gain economies of scale by consolidating under-used resources to a new platform. The storage is equally crucial to server virtualization scenarios. Implementing server virtualization without taking storage into account can cause challenges, such as uneven resource sharing and performance and reliability degradation.

Through its extraordinary flexibility, the DS8000 series storage family is designed to effectively and efficiently manage a broad range of storage workloads that exist in today's complex data centers. This flagship IBM disk system can bring simplicity to your storage environment by supporting a mix of random and sequential I/O workloads and interactive and batch applications, regardless of whether they are running on one of today's popular distributed server platforms like VMware vSphere or on the mainframe. With the ability to manage all of these workloads simultaneously, this dynamically self-tuning system can greatly reduce the challenge of managing your increasingly complex storage environment as requirements change.

## 1.2 DS8870 overview

Built on over 50 years of enterprise disk storage expertise, the IBM System Storage DS8000 series is the flagship disk storage platform within the IBM System Storage portfolio. The new IBM System Storage DS8870 represents the latest in this series of high-performance, high-capacity, flexible, and resilient disk storage systems, which are aimed at addressing the needs of the most demanding clients.

The DS8870 represents another generation leap for the DS8000 series: it offers up to three times faster performance and 20% higher energy efficiency than its direct predecessor, the DS8800 system. The DS8000 architecture is server-based. Powerful POWER7 processor-based servers manage the cache to minimize disk I/Os to maximize performance and throughput. The DS8870 product is tremendously scalable, has broad server support, and virtualization capabilities. These features can help simplify the storage environment by consolidating multiple storage systems.



Figure 1-1 shows the IBM System Storage DS8870 product.



Figure 1-1 DS8870 front view

### Did you know?

Full Disk Encryption (FDE) drives are now standard on every DS8870 system. FDE enables organizations to fully protect their data at rest when it is combined with Tivoli Key Lifecycle Manager. FDE drives have the same performance as standard drives.

IBM Easy Tier, a well-proven DS8000 feature, is available at no charge. It dynamically optimizes performance for multi-tiered systems. It can also rebalance data within a single tier to help maintain optimal performance for organizations, which have not yet deployed solid-state drives (SSDs) or nearline drives. The updated Easy Tier fifth generation added the Easy Tier Server, Easy Tier Application, and Easy Tier Heat Map Transfer features to further enhance the performance of the IBM DS8870 system.

A new power supply system that is based on direct current uninterruptible power supply (DC-UPS) is much more energy-efficient than the Power Distribution System of previous models. The DS8870 system is designed to comply with the emerging ENERGY STAR specifications.

For more information about the DS8870 architecture, see *IBM System Storage DS8870 Architecture and Implementation*, SG24-8085 and *DS8870 Product Guide*, TIPS0915.

## 1.3 VMware architecture overview

VMware vSphere is an operating system that is used to virtualize hardware resources and is the base for the VMware cloud-computing infrastructure.

### 1.3.1 VMware vSphere features

There are six core features of the vSphere platform, as shown in Figure 1-2.

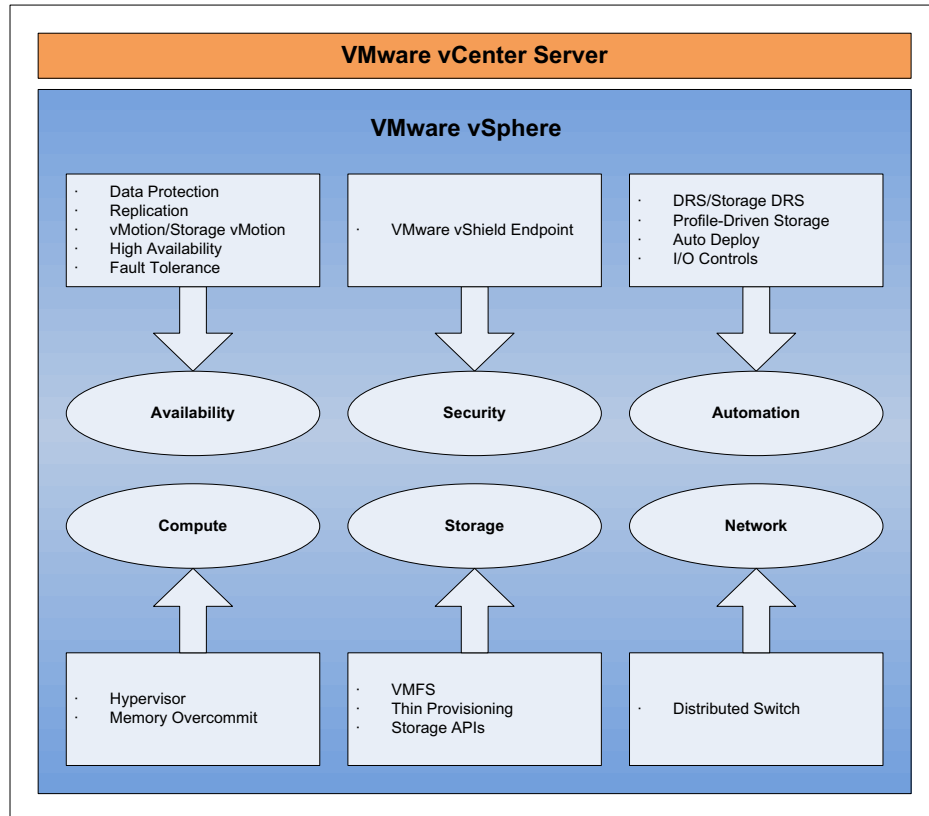


Figure 1-2 vSphere core features

The core features consist of the following components:

- ▶ **Availability:** Increases the availability of all the aspects of the infrastructure, including applications, storage resources, infrastructure, and management. It includes the following components:
  - **Data protection:** Protects the data through fast, agentless backups to disk, with de-duplication to minimize the use of backup disk space.
  - **Replication:** Eliminates third-party replication costs by using built-in vSphere Replication.
  - **vMotion:** By migrating running virtual machines between hosts, vMotion eliminates application downtime for planned server maintenance.
  - **Storage vMotion:** Provides a proactive and nondisruptive storage migration.
  - **High availability:** Provides high availability across the entire virtualized IT environment without the cost or complexity of traditional clustering solutions.

- Fault tolerance: Provides continuous availability for applications with zero data loss in the event of server failures.
- ▶ Security: Provides a robust and secure virtualization platform.
  - vShield endpoint: Eliminates the anti-virus footprint in virtual machines and improves AV scan performance by off loading AV functions to a hardened security virtual machine.
- ▶ Automation
  - Automation provides accurate, consistent, and repeatable solutions that save IT administrators time and effort. It includes the following components:
    - Distributed Resources Scheduler (DRS): Used for aligning compute resources with business priority by automatic load balancing across hosts. DRS optimizes power consumption by turning off hosts during lower load periods.
    - Storage DRS: Facilitates the automated management of initial VMDK placement. It also facilitates continual, dynamic balancing of VMDKs among clustered datastores by identifying the most appropriate resource candidates that are based on capacity, performance, and functional characteristics that are specific to the requirements of individual virtual machines or clusters. Beginning in vSphere 5.0, VMware Storage DRS can make use of VASA-based and administrator-based storage resource classifications to realize simplification of heterogeneous storage management that is based on the concept of Profile-Drive Storage, which organizes diverse storage resources into profiles that meet specific classification criteria, as shown in Figure 1-3.

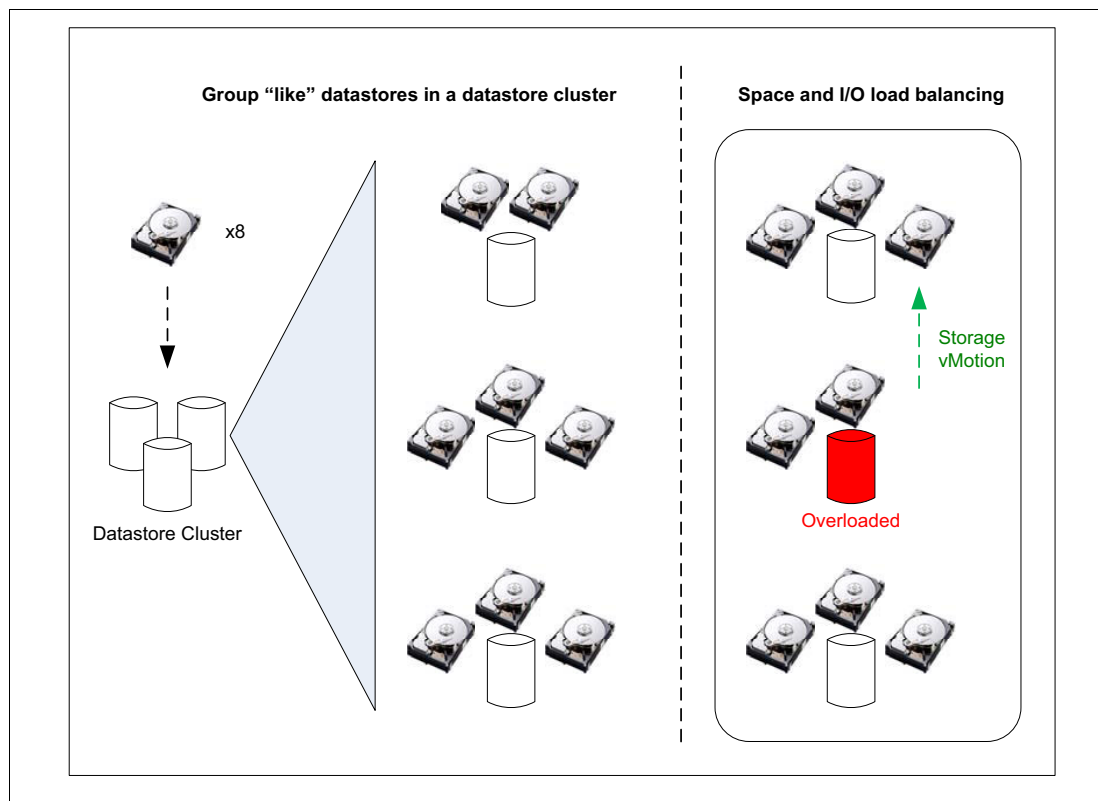


Figure 1-3 Storage DRS

- Profile driven storage: Faster selection of storage resources by grouping storage according to user-defined policies. An example is shown in Figure 1-4 on page 6.

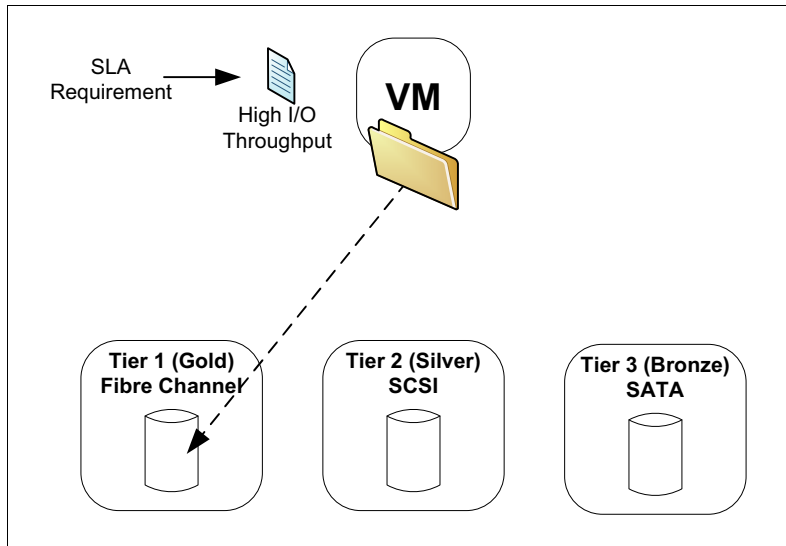


Figure 1-4 Profile Driven Storage

- Auto deploy: Fast deployment and patching of vSphere hosts.
- Network I/O control: Defines priority access to network resources according to established business rules.
- Storage I/O control: Monitors the load on a volume and dynamically allocates resources to virtual machines according to business needs.

► Compute

Virtualizes the server resources and aggregates them into logical pools that can be allocated to multiple workloads. It includes the following components:

- Hypervisor (vSphere ESXi): The virtualization layer that abstracts server hardware resources and can share them with multiple virtual machines.

ESX and ESXi: Before VMware vSphere ESXi version 5.0, clients chose between ESX and ESXi. ESX contains an integrated service console for management. ESXi consists only of the hypervisor and requires the installation of external agents for monitoring and administration. ESX and ESXi provide the same capabilities regarding virtual machines. For more information about the difference between ESX and ESXi, see this website:

<http://www.vmware.com/products/vsphere/esxi-and-esx/compare.html>

- Memory Overcommit: Enables the administrator to provision more memory to virtual machines than are physically installed on the host.

► Storage

Reduces the complexity of back-end storage systems and uses storage in virtual environments. It includes the following components:

- Virtual Machine File System (VMFS): Provides simple virtual machine provisioning and administration with a cluster file system that is optimized for virtual machines.
- Storage thin provisioning: Offers dynamic allocation and intelligent provisioning of physical storage capacity for virtual machine storage usage, as shown in Figure 1-5 on page 7.

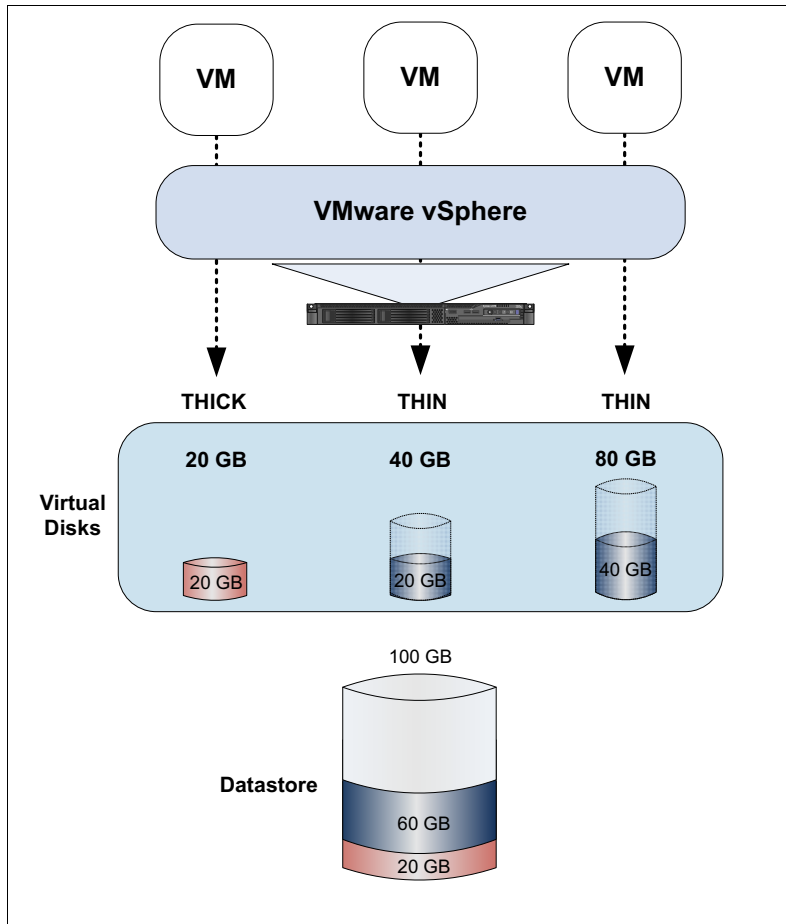


Figure 1-5 Storage thin provisioning

- VMware storage APIs: Storage awareness, data protection, array integration, and multi-pathing capabilities that are offered by third-party storage vendors. It includes the following components:
  - vStorage APIs for Data Protection (VADP): Enables backup applications to perform file or VM-level backup of running virtual machines. It is used by solutions such as Tivoli Storage Manager for VE and Tivoli FlashCopy Manager.
  - vStorage APIs for Multipathing (VAMP): Helps to intelligently control path selection from storage adapters in a host to storage devices, thereby optimizing storage I/O throughput. It is not required with DS8870 because a native VMware multi-pathing policy that uses round-robin is in use.
  - vStorage APIs for Array Integration (VAAI): Enables certain storage tasks to be offloaded from the server hardware to the storage array. VAAI was first implemented for block access protocols and is now native to vSphere 5. More primitives were added to support VAAI also for NFS-based datastores. For more information about VAAI integration with DS8870, see Chapter 2, “VAAI support” on page 13.
  - vStorage APIs for Storage Awareness (VASA): Exposes storage architecture specifications to vSphere to enable policy-based storage management.

- Site Recovery Manager (SRM): A business continuity and disaster recovery solution for VMware ESX servers that provides VM-aware automation of emergency and planned failover or failback scenarios between data centers and incorporates server or storage-based datastore replication. For more information, see Chapter 4, “Business continuity with IBM DS8870 Storage System and VMware Site Recovery Manager” on page 41.

► Network

Provides network services optimization, simplified administration, and management. It includes a distributed switch, which centralizes network provisioning, administration, and monitoring by using datacenter-wide network aggregation.

### 1.3.2 VMware storage concept

Figure 1-6 shows the ESX and ESXi storage concepts. The term *ESXi* is used as a synonym for both. The storage system provides logical unit numbers (LUNs) and grants ESXi servers access to these LUNs. In VMware terminology, LUNs are usually called volumes. These volumes are normally shared between different ESXi servers. A central component of the VMware storage concept is the data store. A data store is an NFS mount or a VMFS file system. VMFS is a clustered file system that can be accessed simultaneously by multiple ESXi systems. Each data store can span multiple physical LUNs or even storage subsystems. Other LUNs can be added to extend a previously created data store without powering down physical servers.

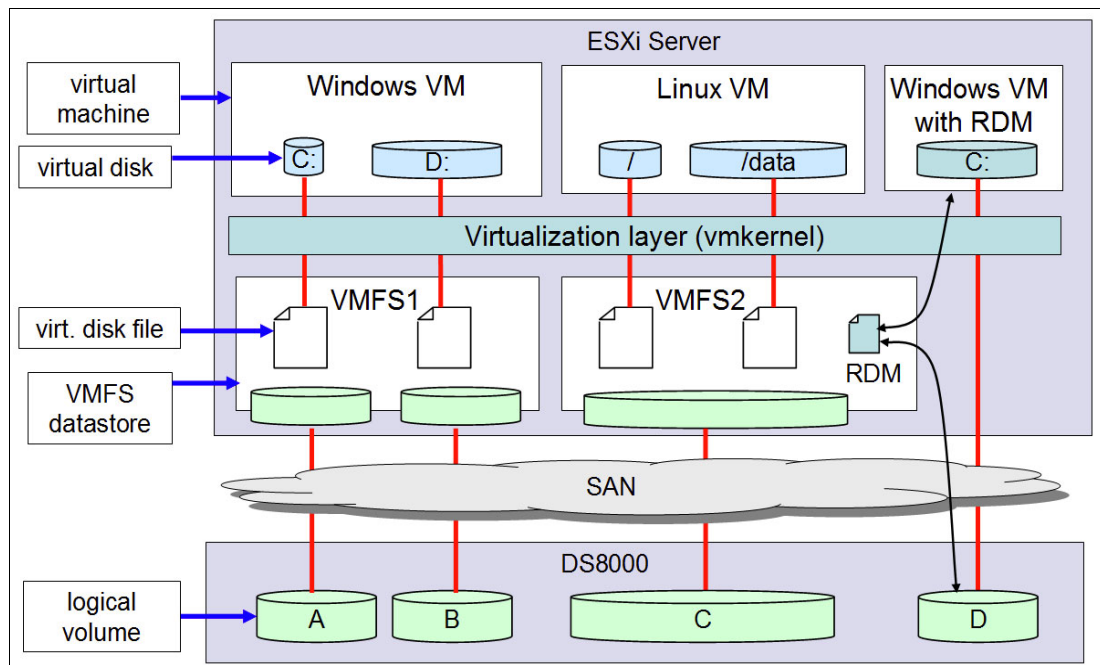


Figure 1-6 Storage concepts of ESXi

VMFS provides a locking mechanism to ensure that the same virtual machine is not powered on by multiple servers at the same time. If a physical server fails, the on-disk lock for each virtual machine can be released so that virtual machines can be restarted on other physical servers.

VMFS also features crash consistency and recovery mechanisms, such as distributed journaling, crash-consistent virtual machine I/O paths, and machine state snapshots.

Instead of the use of the VMFS-based data store, virtual machines can have direct access to raw devices by using Raw Device Mapping (RDM) as a proxy. RDM can be thought of as providing a symbolic link between a VMFS volume and a raw volume (LUN). When a volume is opened for access, VMFS uses the RDM file to get to the correct physical device and performs access checking and locking before accessing the volume. Thereafter, reads and writes go directly to the raw volume rather than going through the mapping file.

A virtual machine typically is stored in a set of files in the VMFS file system or data store. There are files that store the virtual machine's configuration, the state of the virtual machine's basic input/output system (BIOS), and the virtual disk files (.vmdk) which store the content of the virtual machine's hard disk drives, files for metadata, snapshots, logs, and some other files.

Virtual machines can be cloned or moved to another data store by copying these files. Such operations are performed with the help of the vCenter server.

### 1.3.3 What's new in VMware vSphere for storage

VMware is constantly improving the vSphere offering and provides regular updates to the Hypervisor. This section describes storage-related features that were introduced or improved since vSphere 5.

#### News in VMware vSphere 5.1 Update 1

VMware released vSphere 5.1U1 in April 2013.

From a storage perspective, many fixes are included in this release, including the following important changes:

- ▶ The maximum VMFS Heap can now grow up to 640 MB. This is especially useful in installations with big .vmdk files on VMFS volumes that are accessed by single hosts.
- ▶ Renaming of Virtual Machine Files while doing Storage VMotion can be controlled.
- ▶ Fix for incorrectly displayed WWNN and WWPN values for Fibre Channel (FC) host bus adapters (HBAs).

For more information, see the VMware 5.1U1 release notes at this website:

<http://www.vmware.com/support/vsphere5/doc/vsphere-esxi-51u1-release-notes.html#resolvedissuesstorage>

**VMware vSphere functions:** Not all the new VMware vSphere functions for storage that are listed are supported by the DS8870 system. For more information, see the VMware Compatibility Guide, which is available at this website:

<http://www.vmware.com/resources/compatibility/search.php?deviceCategory=san>

#### News in VMware vSphere 5.1

VMware released vSphere 5.1 in September 2012, which includes the following features:

**Important:** VMware vSphere/ESXi version 5.1 is the latest version that is supported by IBM DS8000® System Storage at the time of writing. For the most up-to-date interoperability details, see the IBM System Storage Interoperability Center (SSIC) at this website:

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

- ▶ You can use Storage VMotion now without having a shared storage infrastructure.
- ▶ Support for up to four parallel migrations for a single Storage VMotion.
- ▶ The All Paths Down (APD) feature was enhanced to better detect complex transient APD conditions.
- ▶ Support for Boot from software FCoE initiator. This is similar to the Boot from Software iSCSI initiator that is introduced in ESXi 4.1.
- ▶ Support for jumbo frames for all iSCSI adapters.
- ▶ Support for 16 GB FC HBAs running at 16 GB. Previously, 16 GB FC HBAs were supported only when they were running at 8 GB.
- ▶ It is now possible to automatically set the latency threshold in Storage I/O Control (SIOC).
- ▶ Increased the number of hosts sharing read-only files on a VMFS volume from 8 to 32, which is especially useful in VDI environments when linked clones are used.
- ▶ Enhanced ESXCLI to show statistics for FC, FCoE, iSCSI, and SAS protocols. A new command to monitor local SSDs was introduced. These changes are especially useful when troubleshooting is performed.
- ▶ Introduced Space Efficient Virtual Disk, which allows you to grow and shrink .vmdk files dynamically.

**VMware vSphere 5.1 for storage enhancements:** For more information about the VMware vSphere 5.1 for storage enhancements, see the following VMware white paper, which is available at this website:

<http://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-51-Storage-Technical-Whitepaper.pdf>

## News in VMware vSphere 5.0

VMware released vSphere 5.0 in July 2011. VMware vSphere 5.0 consists of Version 5 of ESXi and Version 5 of the vCenter server.

From a storage perspective, the following significant enhancements in vSphere 5.0 made storage easier to manage:

- ▶ Virtual Machine File System Version 5 (VMFS-5): The following enhancements are provided in VMFS-5:
  - Unified 1 MiB file block size. There is no longer a requirement to use 1, 2, 4, or 8 MB file blocks to create large files (files greater than 256 GB). When a VMFS-5 data store is created, the user is no longer prompted to select a maximum file size.
  - Large single extent volumes. Before VMFS-5, the largest volume that was presented to an ESX server was 2 TiB. With VMFS-5, this limit is increased to 64 TiB (64\*1024\*1024\*1024\*1024 bytes).
  - Increased file count. VMFS-5 introduces support for more than 100,000 files, which are compared to approximately 30,000 in VMFS-3.
  - Atomic Test and Set (ATS). ATS was introduced with VAAI. It is now used throughout VMFS-5 for file locking, which provides improved file locking performance over previous versions of VMFS.
- ▶ Support for larger raw device mappings (RDMs). Pass through RDMs can now be 64 TiB in size (as opposed to 2 TiB in vSphere 4.1).
- ▶ Improvements to storage vMotion by using a new mirror driver mechanism.



- ▶ Removing the requirement to install a vendor-supplied VAAI driver.
- ▶ Storage Distributed Resource Scheduler (SDRS), which allows for intelligent initial placement. SDRS can also load balance by using space usage and I/O latency by using storage vMotion.
- ▶ vSphere Storage APIs for Storage Awareness (VASA), a set of APIs that allows storage devices to report physical characteristics, alerts, and events to vSphere. This set helps with decision making and monitoring. VASA is normally enabled with a vendor-supplied software provider that communicates with the vCenter Server Storage Monitoring Service (SMS).

**VMware vSphere 5.0 for storage enhancements:** For more information about the VMware vSphere 5.0 for storage enhancements, see the following VMware white paper, which is available at this website:

<http://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-50-Storage-Technical-Whitepaper.pdf>





## VAAI support

This chapter provides a brief overview of the VMware vStorage APIs for Array Integration (VAAI) feature and emphasizes its interaction and usage with the IBM System Storage DS8870 system.

This chapter includes the following chapters:

- ▶ Introduction
- ▶ VMware vStorage APIs array integration
- ▶ Using vStorage APIs for array integration
- ▶ Testing vStorage APIs for Array Integration

## 2.1 Introduction

VMware vSphere provides an application programming interface (API) and software development kit (SDK) environment. This environment allows customers and independent software vendors (ISVs) to enhance and extend the functionality of vSphere in specific areas. The VAAI feature offloads specific storage operations from the hosts to be handled by supported disk arrays for improved performance and efficiency.

With VAAI, VMware vSphere can perform key operations faster and use less host processor resources, memory, and storage bandwidth.

The following topics are aimed at giving a clear and comprehensive understanding of the VAAI features and usage with the DS8870 system.

- ▶ VMware vStorage APIs array integration
- ▶ Using vStorage APIs for array integration
- ▶ Testing vStorage APIs for Array Integration

## 2.2 VMware vStorage APIs array integration

ESX and ESXi 4.1 brought a new level of integration capability with storage systems through the introduction of VAAI. Traditional SCSI commands force the ESX host to issue many repetitive commands to complete certain types of operations. VAAI helps reduce host usage. It also increases scalability and operational performance by off loading certain tasks to storage systems if they support the relevant commands.

VAAI support relies on the storage implementing several fundamental operations that are named as *primitives*. These operations are defined in terms of standard SCSI commands, which are defined by the T10 SCSI specification. Each fundamental operation can operate independently of the others, and it is possible to certify the storage as supporting a subset of the available primitives.

### 2.2.1 vStorage APIs for array integration primitives

**Supported primitives:** Not all VMware VAAI primitives are currently supported by the DS8870 system. For more information about the supported primitives, see 2.2.2, “DS8870 support for VAAI” on page 16. For the most up-to-date interoperability details, see the IBM System Storage Interoperability Center (SSIC) and the VMware websites:

- ▶ <http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>
- ▶ <http://www.vmware.com/resources/compatibility>

The VAAI primitives include the following features:

▶ Atomic Test and Set (ATS)

ATS is used to implement hardware-accelerated locking of files. Without support for this primitive, co-operating hosts rely on SCSI reservations, which operate only on the granularity of a full SCSI logical unit number (LUN) (not an individual file). With ATS support, time for creation or access of files from a host can be dramatically improved. With the introduction of ATS, it is now possible to run more VMs on a single volume and to create more VMware-based Snapshots without creating performance problems. To process this command, the storage does the following actions:

- Receives one block (512 bytes) of data from the host.
- Compares that block of data to the data that is already present on the addressed device at a specified logical block access (LBA). If the block data is not byte-for-byte identical, reject the command with sense, which indicates that the data is different.

Otherwise, it receives a second block of data from the host and writes that to the specified logical block address (LBA), which replaces the original data.

These steps are all done automatically so that no other host might alter the data while this operation is in progress. This primitive is also known as *Hardware Accelerated Locking* and it corresponds to the SCSI **Compare And Write** command.

▶ Clone Blocks (XCOPY)

This primitive requests that the storage copy a range of logical blocks from one location to another, either within a single LUN or from one LUN to another. It is used to copy or migrate virtual volumes or machine images from one datastore to another. Without the **XCOPY** command, the host must read all the data from the source blocks and write it to the target blocks, which can be time-consuming for large sets of data. This primitive is also known as *Hardware Accelerated Move* or *full copy* and it corresponds to the SCSI **Extended Copy** command.

▶ Write Same

This primitive is used to initialize new virtual disks. The SCSI **Write Same** command sends a single block of data from the host to the storage, and instructs the storage to replicate that block across a large range of LBAs. Without **Write Same** support, the host zeros the blocks by issuing standard SCSI write commands. With **Write Same** support, the host sends only a single 512-byte block of zeros, which the storage writes to many blocks. The primary advantages are saving bandwidth between the host and the storage, and reducing the processing that is done by the host system. This primitive is also known as *Zero Blocks* or *Hardware Accelerated Init* and it corresponds to the SCSI **Write Same** command.

▶ Thin provisioning

Thin provisioning corresponds to the SCSI **Unmap** or **Write Same** (Unmap) commands. Thin provisioning encompasses the following primitives:

- Out-of-space condition: To pause a running virtual machine when capacity is exhausted.
- Thin provisioning LUN reporting: To enable vSphere to determine the LUN thin provisioning status.
- Quota Exceeded Behavior: Allows vSphere to react before out-of-space conditions occur.
- Unmap: Allows the clearing of unused VMFS space.

## 2.2.2 DS8870 support for VAAI

If you plan on using the VAAI functions with the DS8870 system, the VMware VAAI indicator (Feature Code #0965) must be included in your DS8870 order. The DS8870 with Licensed Machine Code (LMC) 7.7.10.x.xx (bundle version 87.10.x.xx) supports the following VAAI primitives:

- ▶ ATS, also known as *Compare and Write* for hardware-assisted locking.
- ▶ Write Same, also known as *Zero Blocks* or *Hardware Accelerated Init* for hardware assisted initializing of a virtual disk.
- ▶ Clone Blocks, also known as *Extended Copy* or *XCOPY* for hardware-assisted move or cloning.

**Important:** For VAAI XCOPY, IBM FlashCopy® at the track level is used. To take advantage of the VAAI XCOPY primitive, the FlashCopy feature of the DS8870 system must be enabled. This means that a FlashCopy license for Fixed Block capacity is required.

### Restrictions for the use of XCOPY

If you plan to use VAAI XCOPY, the following restrictions apply:

- ▶ A Fixed Block FlashCopy license is required.
- ▶ LUNs used by the ESX and ESXi servers must not be larger than 2 TB.
- ▶ Track Space Efficient volumes and Extent Space Efficient volumes cannot be used.
- ▶ The target of an XCOPY operation cannot be a Remote Mirror and Copy primary volume.

In any case where the XCOPY command is not supported, the DS8870 product indicates to the host that the XCOPY command is not supported, and the host performs the copy tasks. The only impact to applications is that operations that might otherwise use the XCOPY command do not get the benefit of hardware acceleration.

The VAAI primitives that the storage system supports are enabled by default in ESXi5. No intervention is required by the user to obtain the benefit of these features.

## 2.3 Using vStorage APIs for array integration

This chapter describes the VMware VAAI and its support by the DS8870 storage subsystem. For more information about the attachment of storage to VMware ESX and ESXi hosts, see *IBM System Storage DS8000 Host Attachment and Interoperability*, SG24-8887-01.

It is possible to enable and disable specific VAAI features globally on an ESX/ESXi host. Depending on the type of volume you create on the DS8870 system, specific VAAI features are not enabled on the volume even if you turned on the feature globally on the host.

**Important:** In a VMware ESX or ESXi 4.1 environment, another driver must be installed to support VAAI on the DS8870 system. Starting with vSphere 5.0, VAAI is supported natively and no other driver must be installed.

### 2.3.1 Installing the IBM VAAI device driver on an ESXi 4.1 server

The IBM Storage device driver for VMware VAAI is a kernel module that allows the VMware VAAI driver to offload certain storage operations to the DS8870 system. The driver must be installed on every ESX/ESXi 4.1 server and requires that each server is restarted after installation. Updates to the IBM Storage driver also require that each ESX/ESXi server is rebooted. When it is combined with vMotion and vSphere High Availability, this process usually does not require any guest VM outages.

Use the latest version of the driver. For more information about installing the driver on ESX, see the IBM VAAI driver installation guide that is available in the IBM Storage Host Software Solutions Information Center at this website:

[http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strg.hosts.doc%2Fhsg\\_virtualization.html](http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strg.hosts.doc%2Fhsg_virtualization.html)

### 2.3.2 Validating VAAI global settings on the host

The following parameters control whether VAAI is enabled or disabled on an ESX/ESXi host:

- ▶ DataMover.HardwareAcceleratedMove
- ▶ DataMover.HardwareAcceleratedInit
- ▶ VMFS3.HardwareAcceleratedLocking

The default value for all three parameters is 1, which means they are enabled by default. You can check the setting of the parameters in the vSphere Client or via the command-line interface (CLI) on the host. The CLI can be used via a Secure Shell Session to the ESX/ESXi host.

#### 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 all three parameters are set to 1 (enabled), as shown in Figure 2-1 on page 18.

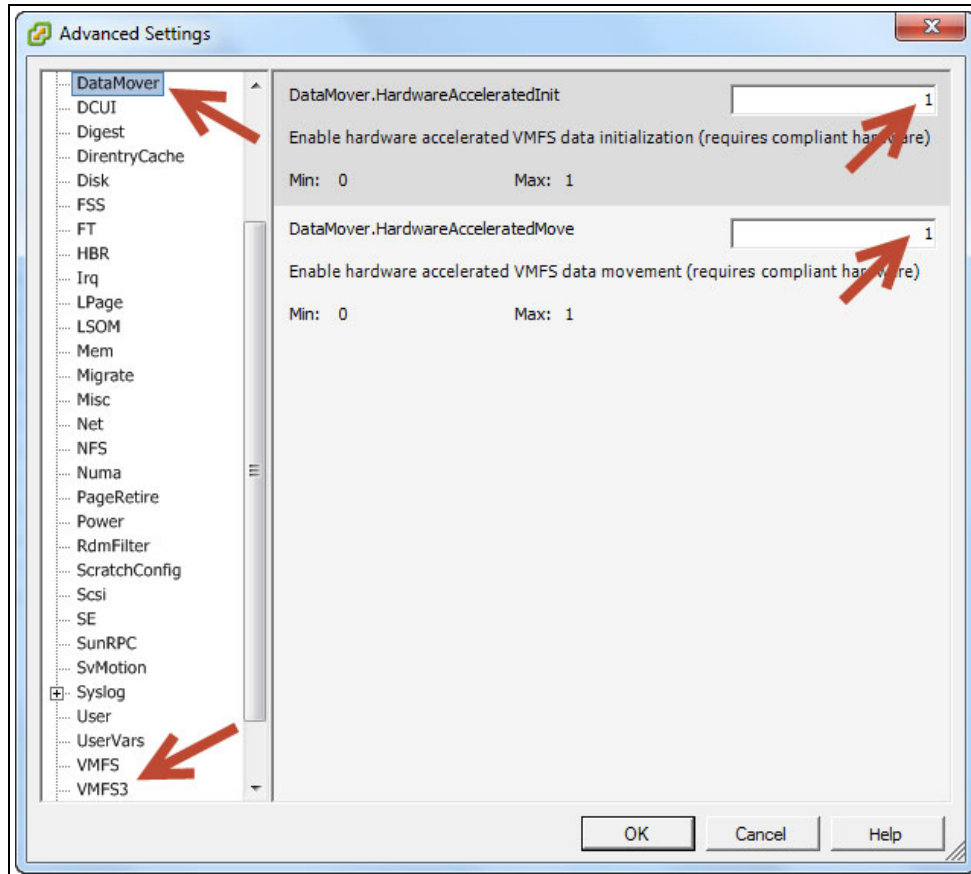


Figure 2-1 Check VAAI in vSphere

## Using the vSphere CLI with ESX and ESXi 4.1

If the service console is used to control VAAI, the commands that are shown in Example 2-1 can be used on ESX/ESXi 4.1 and on ESXi 5.0/5.1. The first three commands show the status of VAAI. If the value that is returned for each function is 0, that function is disabled. If the value returned is 1, the function is enabled.

*Example 2-1 Check if VAAI function is enabled or disabled*

---

```
esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove
esxcfg-advcfg -g /DataMover/HardwareAcceleratedInit
esxcfg-advcfg -g /VMFS3/HardwareAcceleratedLocking
```

---

The commands that are shown in Example 2-2 show how to disable each VAAI function (changing each value to 0).

*Example 2-2 Disable VAAI functions*

---

```
esxcfg-advcfg -s 0 /DataMover/HardwareAcceleratedMove
esxcfg-advcfg -s 0 /DataMover/HardwareAcceleratedInit
esxcfg-advcfg -s 0 /VMFS3/HardwareAcceleratedLocking
```

---



The commands that are shown in Example 2-3 change each VAAI function (change each value to 1).

*Example 2-3 Enable VAAI functions*

---

```
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedInit
esxcfg-advcfg -s 1 /VMFS3/HardwareAcceleratedLocking
```

---

## Using the vSphere CLI with ESXi 5.0 and 5.1

For ESXi 5.0 and 5.1, there also is a new syntax that is available that can be used to control VAAI. Again, the first three commands that are shown in Example 2-4 show the status of VAAI. If the value returned is 0, that function is disabled. If the value returned is 1, the function is enabled.

*Example 2-4 Check if VAAI function is enabled or disabled on ESXi 5.0/5.1*

---

```
esxcli system settings advanced list -o /DataMover/HardwareAcceleratedMove
esxcli system settings advanced list -o /DataMover/HardwareAcceleratedInit
esxcli system settings advanced list -o /VMFS3/HardwareAcceleratedLocking
```

---

The commands that are shown in Example 2-5 disable each VAAI function (changing each value to 0).

*Example 2-5 Disable VAAI functions on ESXi 5.0/5.1*

---

```
esxcli system settings advanced set --int-value 0 --option /DataMover/HardwareAcceleratedMove
esxcli system settings advanced set --int-value 0 --option /DataMover/HardwareAcceleratedInit
esxcli system settings advanced set --int-value 0 --option /VMFS3/HardwareAcceleratedLocking
```

---

The commands that are shown in Example 2-6 enable each VAAI function (changing each value to 1).

*Example 2-6 Enable VAAI functions on ESXi 5.0/5.1*

---

```
esxcli system settings advanced set --int-value 1 --option /DataMover/HardwareAcceleratedMove
esxcli system settings advanced set --int-value 1 --option /DataMover/HardwareAcceleratedInit
esxcli system settings advanced set --int-value 1 --option /VMFS3/HardwareAcceleratedLocking
```

---

### 2.3.3 Confirming that VAAI Hardware Acceleration is detected for a volume

You can confirm if VAAI is enabled on a volume by using the vSphere client or the CLI.

#### Using the vSphere Client

From the vSphere Client, verify whether a datastore volume is VAAI capable by viewing the hardware acceleration status from the Configuration tab, as shown in Figure 2-2 on page 20. Possible states are Unknown, Supported, and Not Supported.

View: **Datstores** **Devices**

**Datstores** Refresh Delete Add Storage... Rescan A

Identification	Status	Device	Hardware Acceleration	Drive Type	Capacity
datastore1	Normal	LSILOGIC Serial Attached SCSI Disk (naa.600508e000000000925c820...	Not supported	Non-SSD	273.50
DS8870_VMFS06	Normal	IBM Fibre Channel Disk (naa.6005076303ffd5aa000000000004915);1	Supported	Non-SSD	49.75 G
DS8870_VMFS07	Normal	IBM Fibre Channel Disk (naa.6005076303ffd5aa0000000000004916);1	Supported	Non-SSD	49.75 G

Figure 2-2 VAAI supported

## Using the vSphere CLI with ESX/ESXi 4.1

In ESX/ESXi4.1, you can use the command that is shown in Example 2-7 to check if VAAI is supported on individual volumes.

Example 2-7 ESX/ESXi 4.1 CLI commands to check the support level of VAAI primitives

```
~ # esxcfg-scsidevs -l | egrep ".Display Name:|VAAI Status:"
Display Name: VM_2Tiers_st
VAAI Status: supported
Display Name: VM_3Tiers_st
VAAI Status: supported
```

## Using the vSphere CLI with ESXi 5.0/5.1

In ESXi 5.x, use the **esxcli storage core device** commands, as shown in Example 2-8. You first must find the device number for the DS8870 LUN in vSphere or run the command without the **-d** option to show all devices.

Example 2-8 ESXi 5.0/5.1 CLI commands to check the support level of VAAI primitives

```
~ # esxcli storage core device list -d naa.6005076303ffd5aa0000000000001004
naa.6005076303ffd5aa0000000000001004
Display Name: DS8870_VAAI1
Has Settable Display Name: true
Size: 409600
Device Type: Direct-Access
Multipath Plugin: NMP
Devfs Path: /vmfs/devices/disks/naa.6005076303ffd5aa0000000000001004
Vendor: IBM
Model: 2107900
Revision: .277
SCSI Level: 5
Is Pseudo: false
Status: on
Is RDM Capable: true
Is Local: false
Is Removable: false
Is SSD: false
Is Offline: false
Is Perennially Reserved: false
Queue Full Sample Size: 0
Queue Full Threshold: 0
Thin Provisioning Status: unknown
Attached Filters:
VAAI Status: supported
```

```
Other UIDs: vm1.02000600006005076303ffd5aa000000000001004323130373930
Is Local SAS Device: false
Is Boot USB Device: false
```

```
~ # esxcli storage core device vaa1 status get -d naa.6005076303ffd5aa0000000000001004
VAAI Plugin Name:
  ATS Status: supported
  Clone Status: supported
  Zero Status: supported
  Delete Status: unsupported
```

---

### 2.3.4 What to do if the Hardware Acceleration status shows as Unknown

ESXi 5.0/5.1 uses an ATS command when it detects a new LUN to determine whether hardware acceleration is possible. For ESX/ESXi 4.1, the initial hardware acceleration status of a datastore or device normally shows as Unknown. The status changes to Supported after ESX/ESXi performs a VAAI offload function. If the attempt by ESX/ESXi to use an offload command fails, the state changes from Unknown to Not Supported. If it succeeds, it changes from Unknown to Supported.

**Hint:** One way to prompt this change is to clone a virtual disk that is on that datastore. You also can copy a virtual disk to a new file in the relevant datastore in the vSphere Client.

## 2.4 Testing vStorage APIs for Array Integration

Three VAAI primitives, ATS, XCOPY, and Write Same (see 2.2.1, “vStorage APIs for array integration primitives” on page 14), are inspected by making some functional tests.

For the test, we used an IBM BladeCenter H with HS21 Blades. ESXi 4.1 and ESXi 5.1 were installed on the HS21 Blades. LUNs from a DS8870 (2107-961) with 256 GB cache were mapped to the ESXi host and formatted with VMFS5. IBM B80 SAN Switches were used to connect the IBM BladeCenter H and the IBM DS8870 product.

Figure 2-3 on page 19 shows a VAAI test setup.




Figure 2-3 VAAI test setup

### 2.4.1 Atomic Test and Set test

In this test, we test the steps to verify that the ATS feature is working by showing its effect before and after the VM is started. Complete the following steps:

1. Create a volume on the DS8870 system and then create a data store that uses that volume. (By using this process, you can run your tests on a data store that has no competing traffic.)
2. From the vSphere Client home page, go to **Hosts and Clusters**.
3. Right-click the selected ESX or ESXi server and select **New Virtual Machine**.
4. When prompted to select a configuration, leave it on Typical.
5. Give the new virtual machine a name.
6. For a data store, select the data store that you created.
7. When you are prompted for a Guest Operating System, leave it on default.
8. When prompted to create a disk, leave the default size (40 GB). If you are using vSphere Client 5.0, select the **Thick Provision Eager Zero** option. This option formats the Virtual Machine Disk (VMDK) with zeros.

**Virtual Machine Disk file format:** VMDK is a file format that is used for virtual appliances that are developed for VMware products. The format is a container for virtual hard disk drives to be used in virtual machines.

9. After the virtual machine is created, connect to the CLI on the ESXi over ssh and issue the **esxtop** command.

**Secure Shell:** Secure Shell (SSH) is a UNIX based command interface and protocol for securely getting access to a remote computer. Linux and Mac OS X operating systems include the SSH client Terminal by default. Windows-based operating systems do not include an SSH client by default. PuTTY, a free and open source terminal emulator application that can act as a client for the SSH, can be used for Windows environments. For information and a download link, see this website:

<http://www.putty.org/>

The VMware `esxtop` utility allows monitoring and collection of data for all system resources (like processor, memory, disk, and network). For information about using `esxtop` and identifying storage performance issues with `esxtop`, see this website:

[http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=1008205](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1008205)

10. Enter `u` for Disk Devices and `f` to add and remove fields. Select the **Device** and **VAAIStats** fields, as shown in Figure 2-4.

```
Current Field order: AbcdefghijklmnOp

* A:  DEVICE = Device Name
  B:  ID = Path/World/Partition Id
  C:  NUM = Num of Objects
  D:  SHARES = Shares
  E:  BLKSZ = Block Size (bytes)
  F:  QSTATS = Queue Stats
  G:  IOSTATS = I/O Stats
  H:  RESVSTATS = Reserve Stats
  I:  LATSTATS/cmd = Overall Latency Stats (ms)
  J:  LATSTATS/rd = Read Latency Stats (ms)
  K:  LATSTATS/wr = Write Latency Stats (ms)
  L:  ERRSTATS/s = Error Stats
  M:  PAESTATS/s = PAE Stats
  N:  SPLTSTATS/s = SPLIT Stats
* O:  VAAIStats= VAAI Stats
  P:  VAAILATSTATS/cmd = VAAI Latency Stats (ms)

Toggle fields with a-p, any other key to return:
```

Figure 2-4 `esxtop` settings

With these settings, you can monitor the VAAI statistics on the ESX host.

When an ATS call is made to the volume, the ATS counter increases. If ATS is working, it is expected that the counter increases during creation and locking of files on the VMFS volumes. Figure 2-5 on page 24 shows that the counter increased after a VM was started in the lab test.

DEVICE	CLONE RD	CLONE WR	CLONE F	MBC RD/s	MBC WR/s	ATS	ATSF	ZERO	ZERO F	MBZERO/s	DELETE	DELETE F	MBDEL/s
mpx.vmhba0:C0:T0:L0	0	0	0	0.00	0.00	0	0	0	0	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004924	0	0	0	0.00	0.00	1295	0	0	78	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004925	0	0	0	0.00	0.00	731	0	0	32	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004926	0	0	0	0.00	0.00	2216	0	0	41	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004927	0	0	0	0.00	0.00	54	0	0	0	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492a	0	0	0	0.00	0.00	3318	0	0	69	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492b	0	0	0	0.00	0.00	46	0	0	0	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492c	0	0	0	0.00	0.00	4163	0	0	95	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492d	0	0	0	0.00	0.00	26	0	0	0	0.00	0	0	0.00
naa.600508e0000000005c638200aef6300f	0	0	0	0.00	0.00	0	0	0	0	0.00	0	0	0.00
naa.600508e000000000925c8200f0f5300f	0	0	0	0.00	0.00	0	0	0	0	0.00	0	0	0.00

Figure 2-5 Before VM was started

Figure 2-6 shows that the counter increased after a VM was started in the lab test.

DEVICE	CLONE RD	CLONE WR	CLONE F	MBC RD/s	MBC WR/s	ATS	ATSF	ZERO	ZERO F	MBZERO/s	DELETE	DELETE F	MBDEL/s
mpx.vmhba0:C0:T0:L0	0	0	0	0.00	0.00	0	0	0	0	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004924	0	0	0	0.00	0.00	1295	0	0	78	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004925	0	0	0	0.00	0.00	731	0	0	32	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004926	0	0	0	0.00	0.00	2295	0	0	42	0.00	0	0	0.00
naa.6005076303ffd5aa0000000000004927	0	0	0	0.00	0.00	54	0	0	0	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492a	0	0	0	0.00	0.00	3318	0	0	69	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492b	0	0	0	0.00	0.00	46	0	0	0	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492c	0	0	0	0.00	0.00	4163	0	0	95	0.00	0	0	0.00
naa.6005076303ffd5aa000000000000492d	0	0	0	0.00	0.00	26	0	0	0	0.00	0	0	0.00
naa.600508e0000000005c638200aef6300f	0	0	0	0.00	0.00	0	0	0	0	0.00	0	0	0.00
naa.600508e000000000925c8200f0f5300f	0	0	0	0.00	0.00	0	0	0	0	0.00	0	0	0.00

Figure 2-6 After VM was started

## 2.4.2 XCOPY test

In this test, the VMware Storage vMotion feature is used to show the performance effect on the host before and after the HardwareAcceleratedMove feature of VAAI is enabled.

Complete the following steps to run the test:

1. Create two volumes on the DS8870 system and then create two data stores by using these volumes. (With this process, you can run your tests on data stores that have no competing traffic.)
2. From the vSphere Client home page, go to Hosts and Clusters.
3. Right-click your selected ESX or ESXi server and select **New Virtual Machine**.
4. When prompted to select a configuration, leave it on Typical.
5. Give the new virtual machine a name.
6. For a data store, select one of the data stores that you created.
7. When prompted for a Guest Operating System, leave it on default.
8. When prompted to create a disk, leave the default size (40 GB). If you are using vSphere Client 5.0, select the **Thick Provision Eager Zero** option. This option formats the VMDK with zeros.
9. Disable the HardwareAcceleratedMove option. From the vSphere Client home page, go to Hosts and Clusters and select your ESX or ESXi server. Click **Configuration** and on Advanced Settings in the Software section, change the DataMover.HardwareAcceleratedMove parameter to 0, as shown in Figure 2-7 on page 25.




Figure 2-7 DataMover value set to 0

10. Right-click the virtual machine and select **Migrate**. In the next window, choose **Change datastore** and click **Next**. Now select the target data store and start the migration.
11. In the vSphere performance monitor, select **Advanced**. In Chart Options, select the read and write rate graphs for the two test volumes. You see the migration traffic is handled by the ESX host. In the test setup, an approximately 400 MBps throughput was observed from the ESX host during migration. This throughput was limited by the 4 Gbps FC HBAs in the ESX host. The migration completed in approximately 5 minutes, as shown in Figure 2-8.

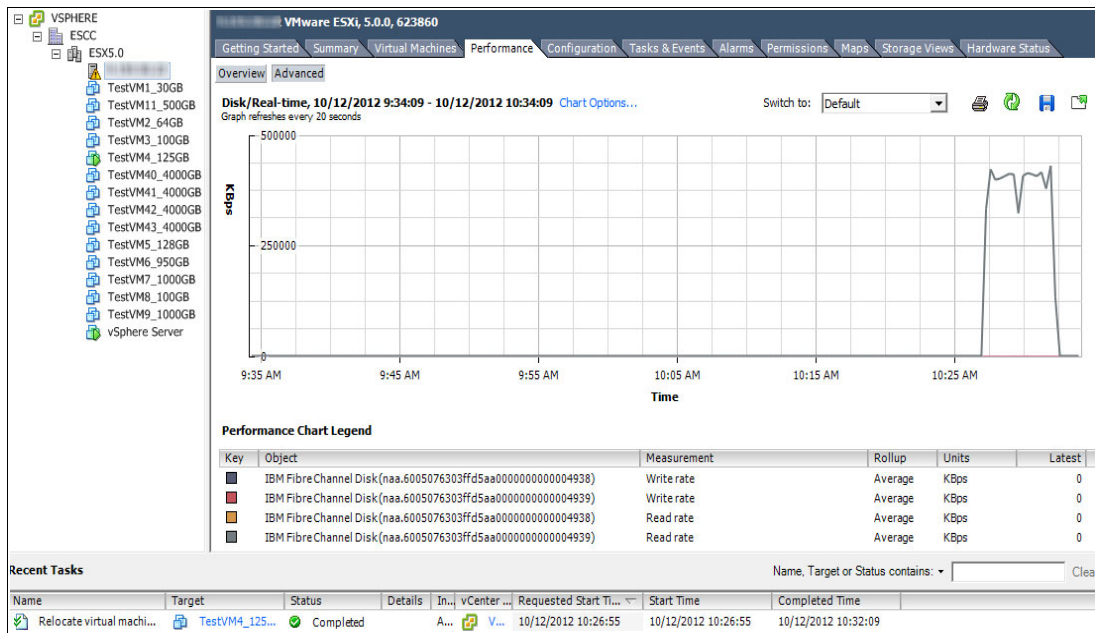


Figure 2-8 Before the DataMover of VAAI is enabled

12. Enable the HardwareAcceleratedMove option. From the vSphere Client home page, go to Hosts and Clusters and select your ESX or ESXi server. Click **Configuration** and on Advanced Settings in the Software section, change the DataMover.HardwareAcceleratedMove parameter to 1.

13. Repeat the test. No traffic on the ESX hosts is observed, as shown in Figure 2-9. This observation means that the migration traffic was offloaded to storage. In the test setup, a peak value of 2 MBps throughput was observed from the ESX host during migration. The migration completed in approximately 1 minute.

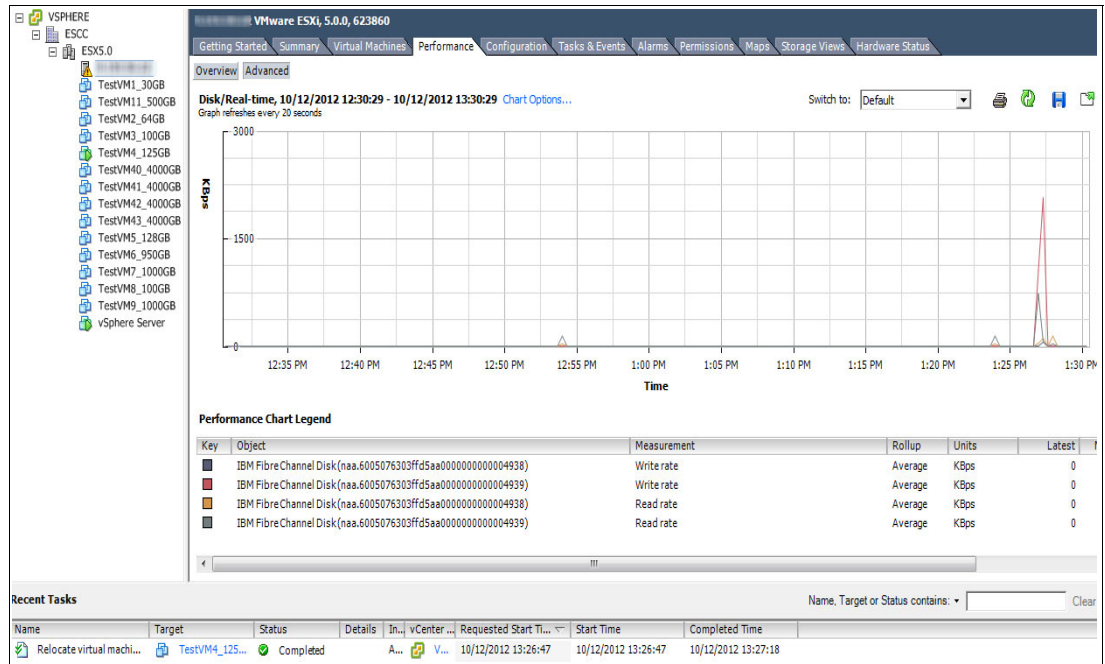


Figure 2-9 After the DataMover of VAAI is enabled

You can also use the DS8000 CLI to see that FlashCopy is used for XCOPY. For more information about how to install and use the DS8000 CLI, see *IBM System Storage DS8870 Architecture and Implementation*, SG24-8085.

The `lsflash` command with the `-dataset` option can show if FlashCopy relations exist, as shown in Example 2-9.

Example 2-9 `lsflash` command to show dataset FlashCopy relations

```
dsccli> lsflash -dataset 8800-8801
Date/Time: October 15, 2012 11:10:32 AM EDT
ID          SrcLSS SequenceNum Timeout ActiveCopy Recording Persistent Revertible
=====
8800:8800 - - - - - - - -
8801:8801 - - - - - - - -
```

**Tip:** The `esxtop` command, as described in the ATS test, can also be used to check that VAAI is used to clone or migrate the data.



## 2.4.3 Write Same Test

In this test, we create a new Virtual Machine in vSphere. While the virtual disk is created for the new VM, we choose the Thick Provisioned Eager Zeroed disk. When a Thick Provisioned Eager Zeroed disk is used, the space that is required for the virtual disk is allocated immediately and all the blocks are zeroed out on the physical device. The ESX Host issues a Write Same Command to the DS8870 system and the zeroing task is offloaded from the Host to the Storage Subsystem.

Complete the following steps to perform the test:

1. Create a volume on the DS8870 product and then create a datastore that uses that volume.
2. From the vSphere Client home page, go to Hosts and Clusters.
3. Right-click your selected ESX or ESXi server and select **New Virtual Machine**.
4. When prompted to select a configuration, leave it on Typical.
5. Give the new virtual machine a name.
6. For a data store, select the new data store that you created.
7. When prompted for a Guest Operating System, leave it on default.
8. When prompted to create a disk, select the **Thick Provision Eager Zeroed** option as shown in Figure 2-10. This option formats the VMDK with zeros.

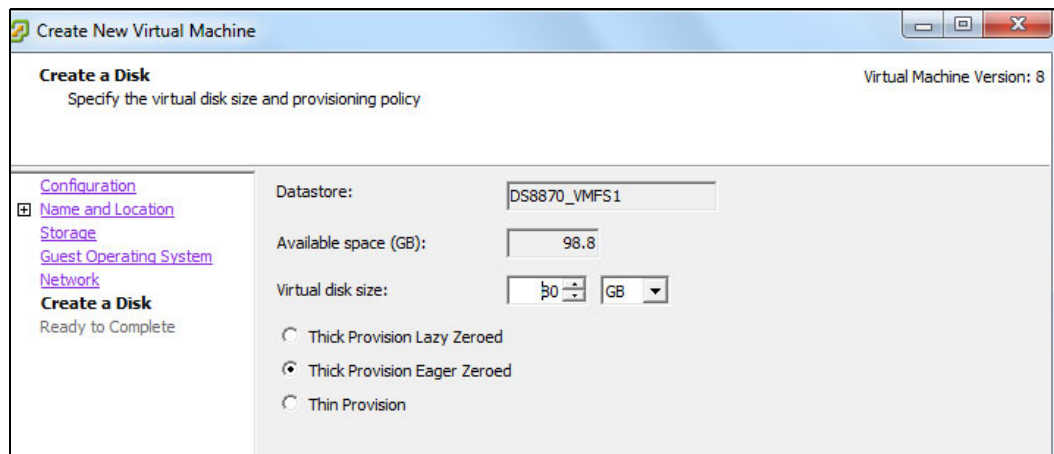


Figure 2-10 Create Thick Provision Eager Zeroed disk

9. Before you finish the virtual machine creation, connect to the CLI on the ESXi server with **ssh** and start the **esxtop** command.
10. Enter **u** for Disk Devices and **f** to add and remove fields. Select the **Device** and **VAAISTATS** fields. Figure 2-4 on page 23 shows the details for the **esxtop** settings.
11. Complete the VM creation in the vSphere Client and monitor the **esxtop** output. The counter **ZERO** is increased every time a block is zeroed and the VAAI **Write Same** command is used. Monitor the **ZERO** counter, as shown in Figure 2-11 on page 28.

DEVICE	CLONE RD	CLONE WR	CLONE F	MBC	RD/s	MBC WR/s	ATS	ATSF	ZERO	ZERO F	MBZERO/s
naa.6005076303ffd5aa0000000000001000	0	0	0	0.00	0.00	1774	0	40960	0	0.00	
naa.6005076303ffd5aa0000000000001001	0	0	0	0.00	0.00	7069	0	0	24	0.00	
naa.6005076303ffd5aa0000000000001002	0	0	0	0.00	0.00	10	0	0	0	0.00	
naa.6005076303ffd5aa0000000000001003	0	0	0	0.00	0.00	0	0	0	0	0.00	
naa.6005076303ffd5aa0000000000001004	38400	0	0	0.00	0.00	3934	0	2	0	0.00	
naa.6005076303ffd5aa0000000000001005	0	38400	0	0.00	0.00	6175	0	102401	0	0.00	
naa.6005076303ffd5aa0000000000001100	0	0	0	0.00	0.00	4	0	0	0	0.00	
naa.6005076303ffd5aa0000000000001101	0	0	0	0.00	0.00	17	0	0	0	0.00	
naa.600508e00000000014f23aa3368bf107	0	0	0	0.00	0.00	0	0	0	0	0.00	

Figure 2-11 esxtop Output for Write Same



# IBM Storage Management Console for VMware vCenter

This chapter introduces the IBM Storage Management Console for VMware vCenter. It is sometimes also called VMware vCenter Storage Management Plug-in, vCenter Storage Plug-in, and vCenter Plug-in.

This chapter includes the following topics:

- ▶ Introduction
- ▶ Installation
- ▶ Adding IBM Storage to the plug-in
- ▶ Manage DS8870 volume properties
- ▶ Creating a datastore
- ▶ Multipathing
- ▶ Host attachment

### 3.1 Introduction

The IBM Storage Management Console for VMware vCenter software is basically a software plug-in that integrates into the VMware vCenter server platform and enables VMware administrators to independently and centrally manage their storage resources on IBM storage systems. These resources include DS8000, Flex System V7000, Scale Out Network Attached Storage, Storwize V3500/V3700, Storwize V7000, Storwize V7000 Unified and SAN Volume Controller.

The plug-in runs as a Microsoft Windows Server service on the vCenter server. Any VMware vSphere Client that connects to the vCenter server detects the service on the server. The service then automatically enables the IBM storage management features on the vSphere Client.

Figure 3-1 shows how IBM storage systems are accessed and controlled through the VMware environment. Also shown is the primary relationships and interaction between the VMware components and the IBM storage systems that provide the storage pools and file systems in which the volumes and file shares are created.

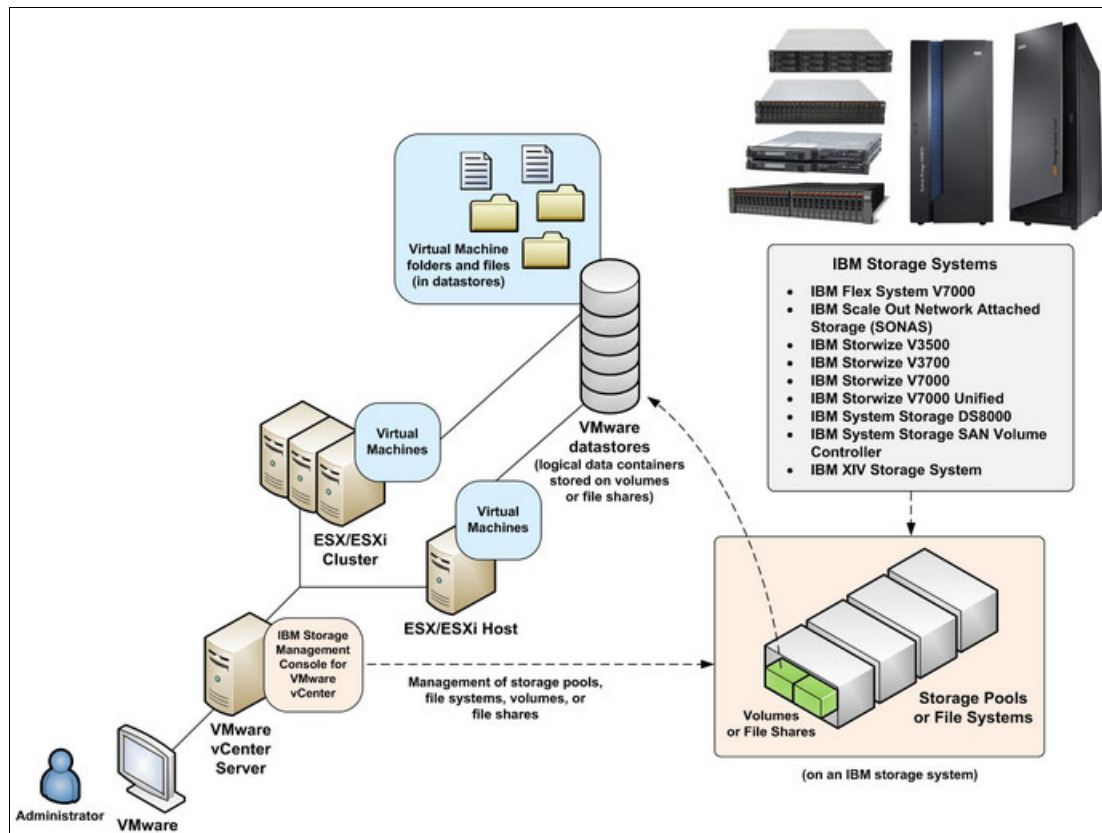


Figure 3-1 IBM Storage Management Console for VMware vCenter concept

**Important:** The IBM Storage Management Console for VMware vCenter can use only predefined storage pools or file systems. New storage pools or file systems cannot be created from the IBM Storage Management Console.

## 3.2 Installation

Install the IBM Storage Management Console for VMware vCenter onto the Windows server that is running VMWare vCenter version 4.0, 4.1, 5.0, or 5.1. There are separate installation packages for x86 and x64. Check the user guide for all prerequisites before installing the plug-in. The user guides for different versions of the plug-in can be found in the IBM Storage Host Software Solutions Information Center at this website:

<http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strg.hosts.doc%2Fvcplugin-homepage.html>

During package installation, you are prompted to complete the following steps:

1. Confirm which language you want to use.
2. Accept license agreements.
3. Select an installation directory location (a default location is offered).
4. When the installation process is complete, a command-line configuration wizard starts, as shown in Example 3-1. Normally, you can safely accept each prompt in the wizard. When prompted for a user name and password, you need a user ID that can log on to the VMware vSphere Client. If you do not change or accept the SSL certificate, you receive Certificate Warning windows when the Client is started. For more information about how to replace the SSL certificate, see the plug-in user guide.

*Example 3-1 IBM Storage Management Console for VMware vCenter Configuration wizard*

---

```
-----  
Welcome to the IBM Storage Management Console for VMware vCenter setup wizard, version 3.2.0.  
Use this wizard to configure the IBM Storage Management Console for VMware vCenter.  
Press [ENTER] to proceed.
```

```
-----  
The Wizard will now install the Management Console service and register the extension in the  
vCenter server.
```

```
Do you want to continue? [default: yes ]:
```

```
-----  
The IBM Storage Management Console for VMware vCenter requires a valid username for connecting to  
the vCenter server.
```

```
This user should have permission to register the Plug-in in the Plug-ins Manager.
```

```
Please enter a username : Administrator
```

```
-----  
Please enter the password for user Administrator :
```

```
-----  
The IBM Storage Management Console for VMware vCenter web component requires a valid network  
port number.
```

```
Please enter a port number for the web component [default: 8880 ]:
```

```
-----  
Please wait while configuring the service and registering the extension
```

```
-----  
The IBM Storage Management Console for VMware vCenter is now configured.
```

```
This product is using an SSL certificate which is not signed.
```

```
Please consult the User Guide for SSL certificate replacement instructions.
```

```
Press [ENTER] to exit.
```

---

5. After you configure the IBM Storage Management Console for VMware vCenter, restart the client if you were already logged on. A new IBM Storage icon and a new IBM Storage tab with all their associated functions are added to the VMware vSphere Client.

You can access the IBM Storage icon from the Home view, as shown in Figure 3-2.

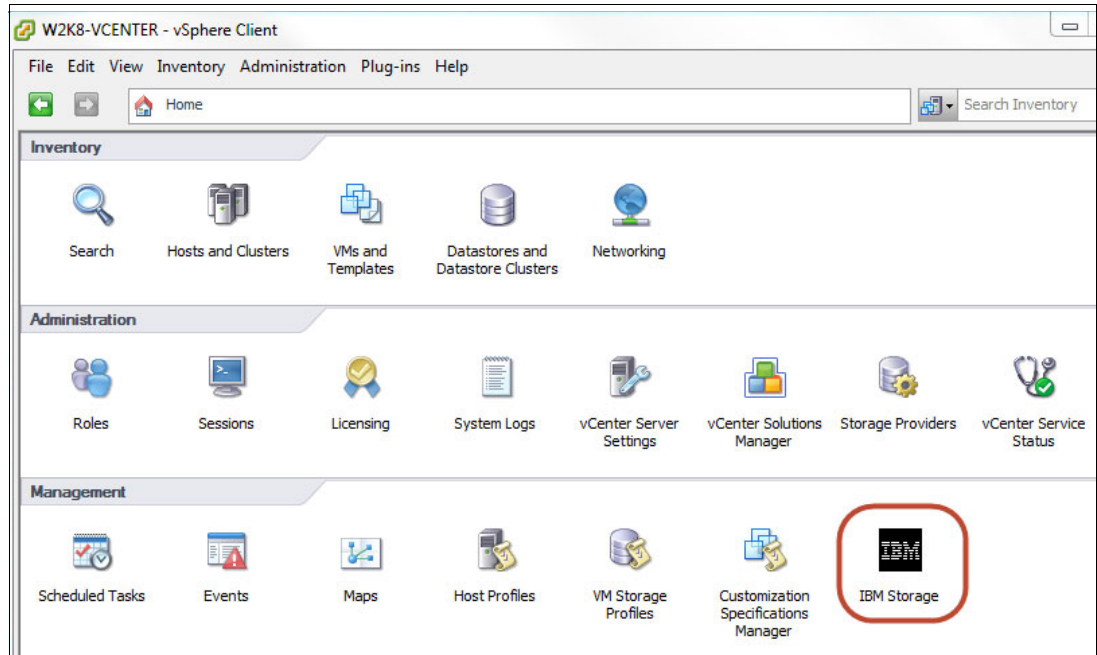


Figure 3-2 IBM Storage plug-in from Home menu

### 3.3 Adding IBM Storage to the plug-in

Complete the following steps to add IBM storage to the plug-in:

1. From the Home view of the vSphere Client, double-click the IBM Storage icon.
2. The Storage Systems view opens, which shows all defined IBM Storage Systems. Select the option to add a Storage System, as shown in Figure 3-3.

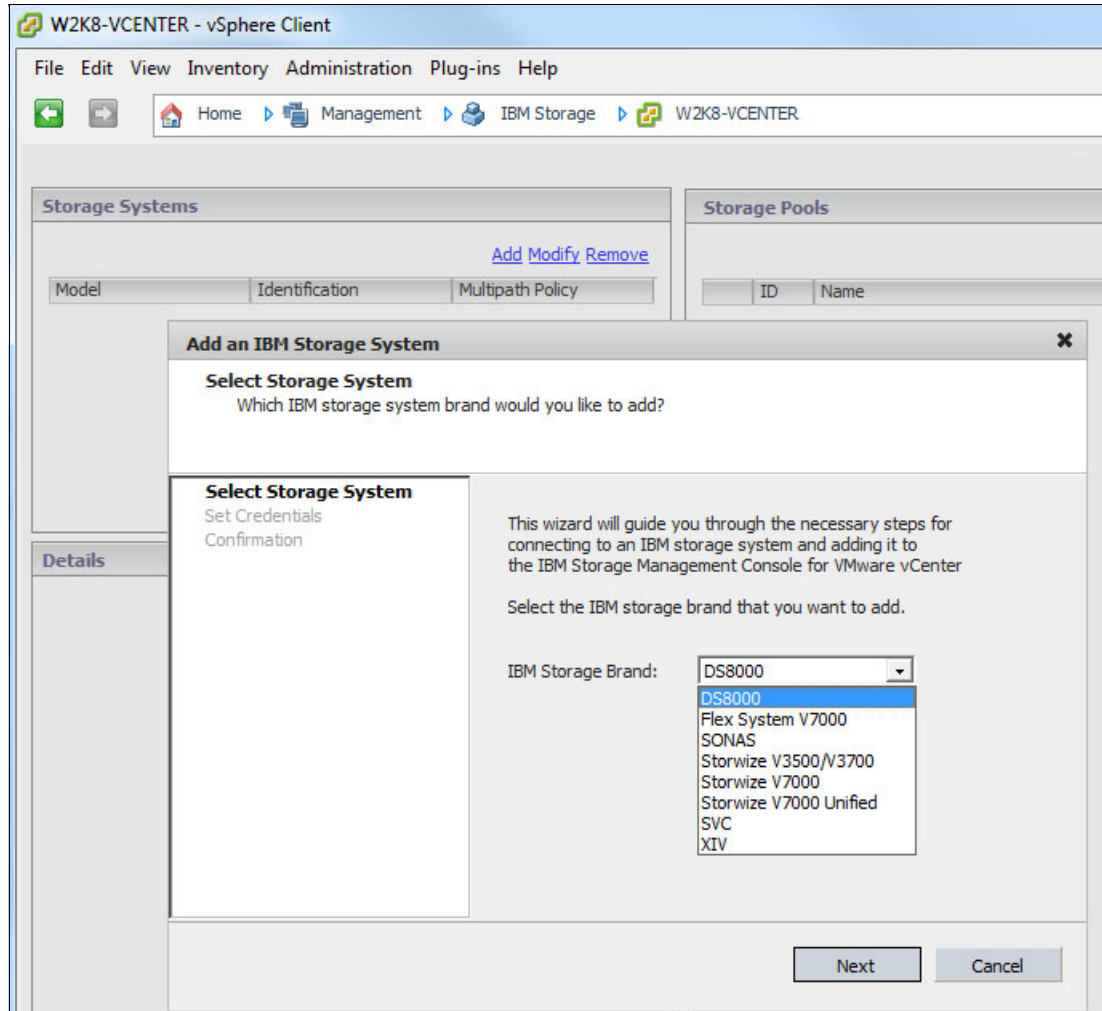


Figure 3-3 Selecting the Add option

3. In the next window (see Figure 3-4), enter an IP address, user name, and password to connect to the IBM DS8870 product.

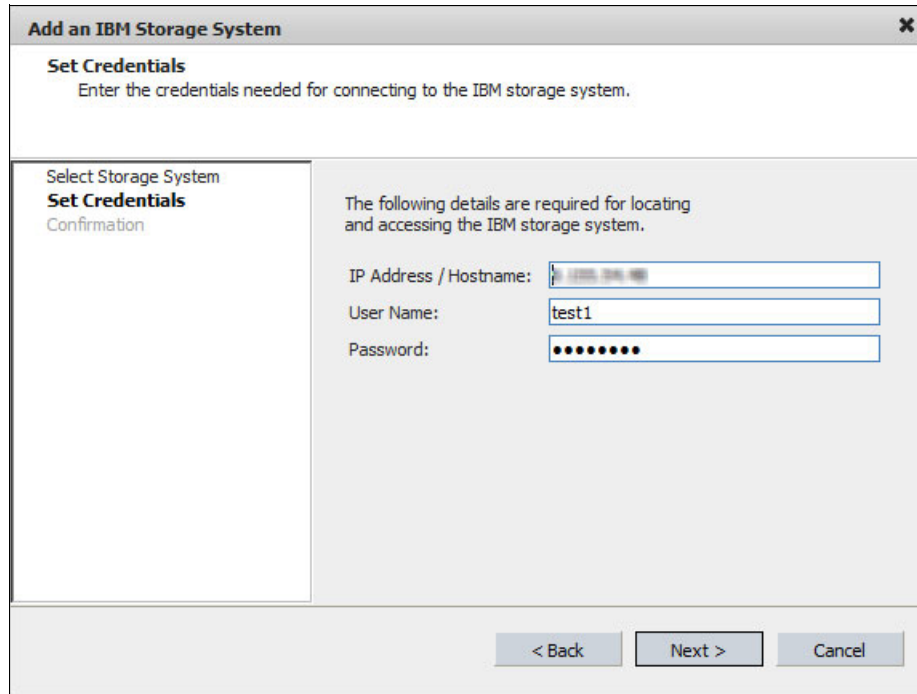


Figure 3-4 Adding a DS8870 system to the plug-in

4. After the vSphere Client connects to the DS8870 system successfully, you see a list of storage pools on the DS8870 system. Select the pools that the VMware administrator allocates storage from, as shown in Figure 3-5. Other pools can be manually added later, if required.

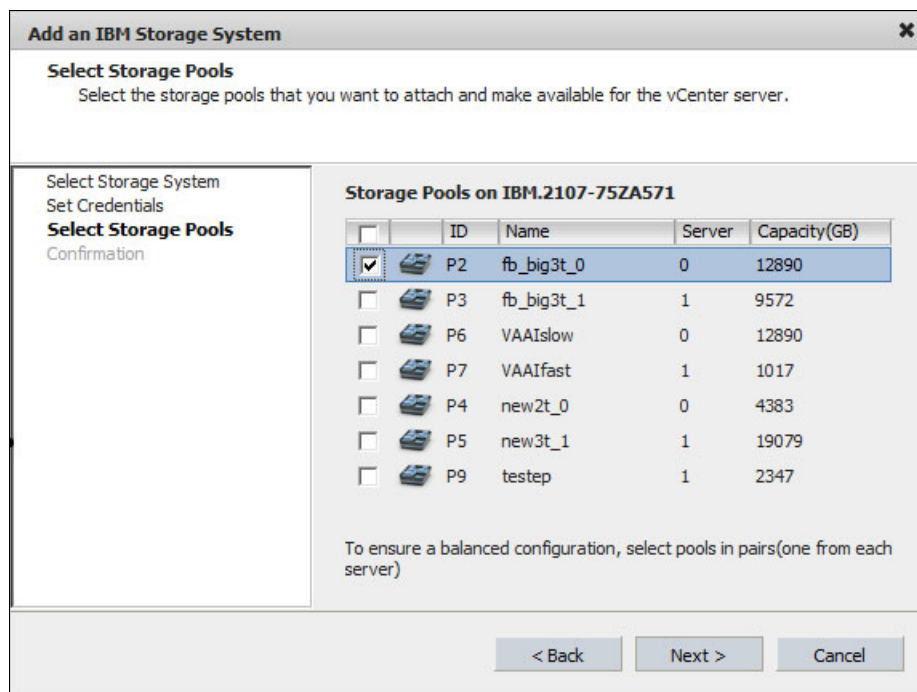


Figure 3-5 Selecting a storage pool from the plug-in



5. The DS8870 product is displayed in the list of storage systems, as shown in Figure 3-6.

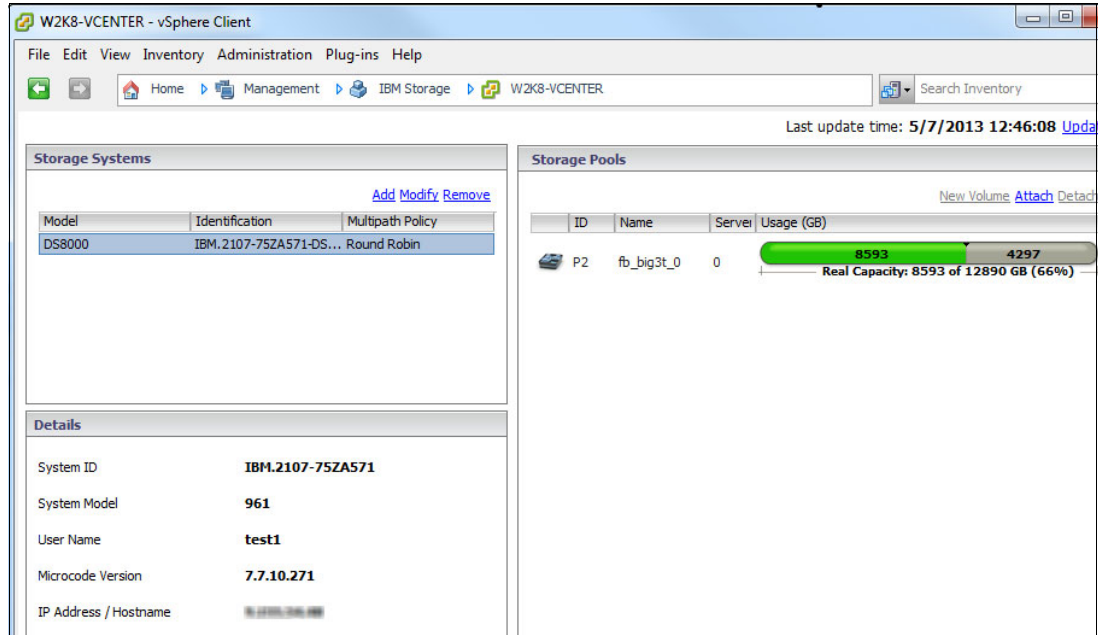


Figure 3-6 vSphere IBM Storage plug-in

6. Select a DS8870 system from the Storage Systems section and then select a pool from the Storage Pools section. Use the **New Volume** option to create a volume and complete the fields in the Create New Volume window, as shown in Figure 3-7. You can create standard or thin-provisioned volumes on the DS8870 system.

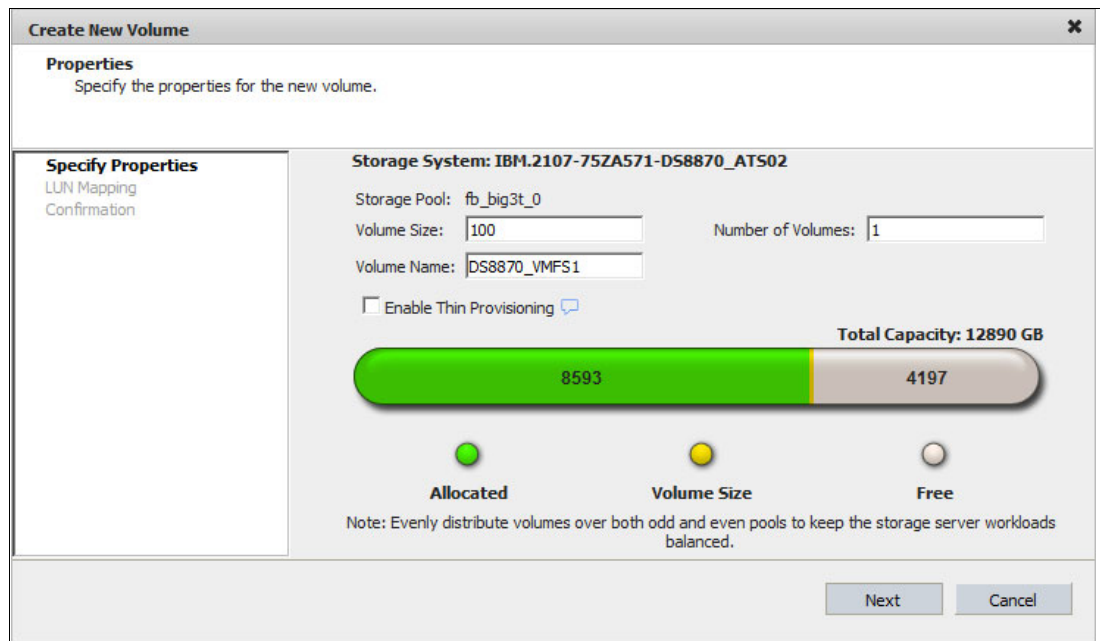


Figure 3-7 Create New Volume with plug-in

7. You are prompted to map the volume to individual VMware servers or the whole cluster, as shown in Figure 3-8.

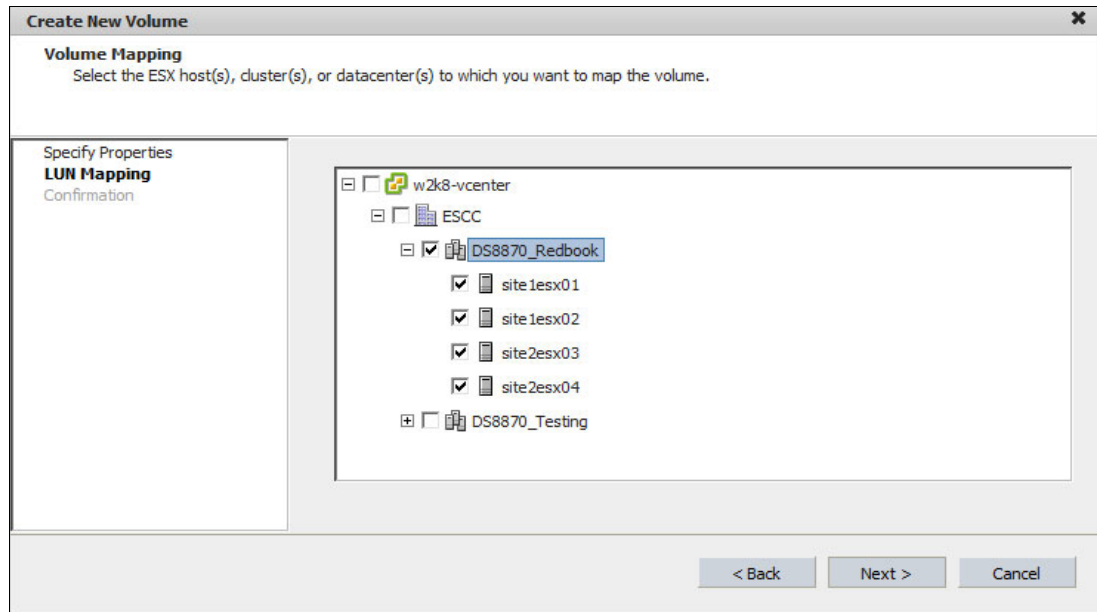


Figure 3-8 Mapping the volumes to hosts and clusters

Normally you select the entire cluster. Ensure that the zoning is correct. The plug-in is used to create the volumes on the DS8870 system, create hosts on the DS8870 system if necessary, and map the volumes to the hosts. You can use the plug-in or the DS8870 system to change the mapping at any time.

## 3.4 Manage DS8870 volume properties

Use the IBM Storage tab to identify the properties of the volume. The IBM Storage tab allows you to perform many useful storage tasks. From this tab, as shown in Figure 3-9, you can perform the following tasks:

- ▶ Extend a volume. This task allows you to grow an existing volume and then later extend the datastore using that volume.
- ▶ Rename a volume. Use this task to ensure that the DS8870 volume name and the datastore name are the same.
- ▶ Query which datastore is which DS8870 volume.
- ▶ Query the status of all of the volumes, snapshots, and mirrors.
- ▶ Migrate volumes between pools.
- ▶ Unmap and delete a volume if it is not used by a datastore. This process allows a VMware administrator to return a volume the DS8870 system.

Figure 3-9 shows the IBM storage tab, which displays more information about the volume. This information includes the storage pool for the volume and the volume ID. It also shows FlashCopy or mirroring relationships for the selected volume and indicates whether it is a standard or thin-provisioned volume.

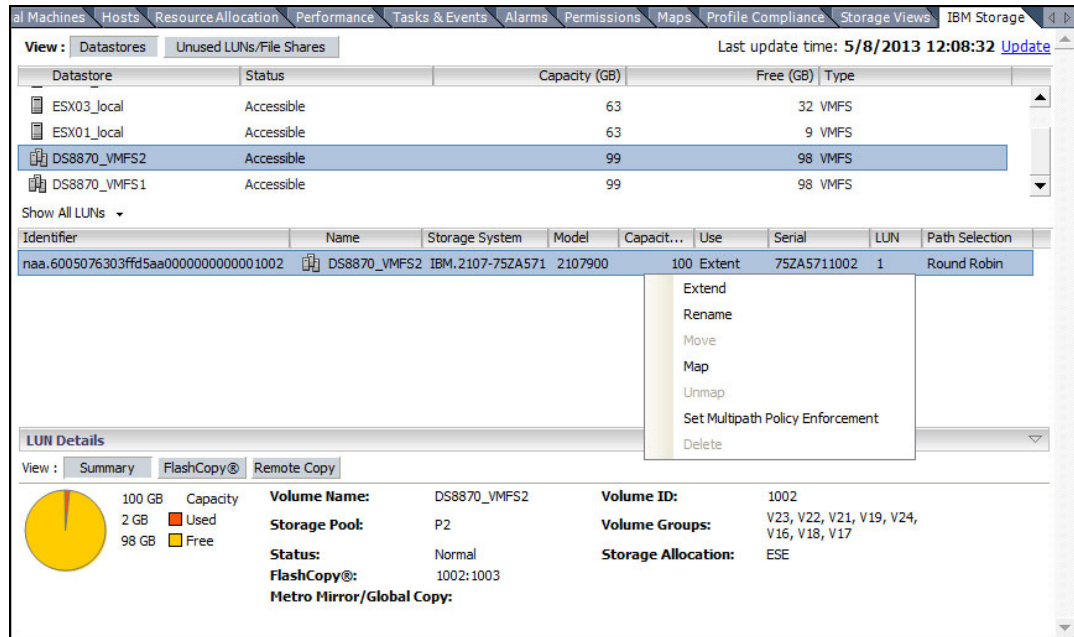


Figure 3-9 The IBM Storage tab

### 3.5 Creating a datastore

Before a datastore is added, check if the logical unit number (LUN) is available through the FC HBAs of the hosts. You can do that by going to your ESX Hosts and check under the Configuration tab in the Storage Adapters menu if your LUNs are accessible through the HBAs, as shown in Figure 3-10.

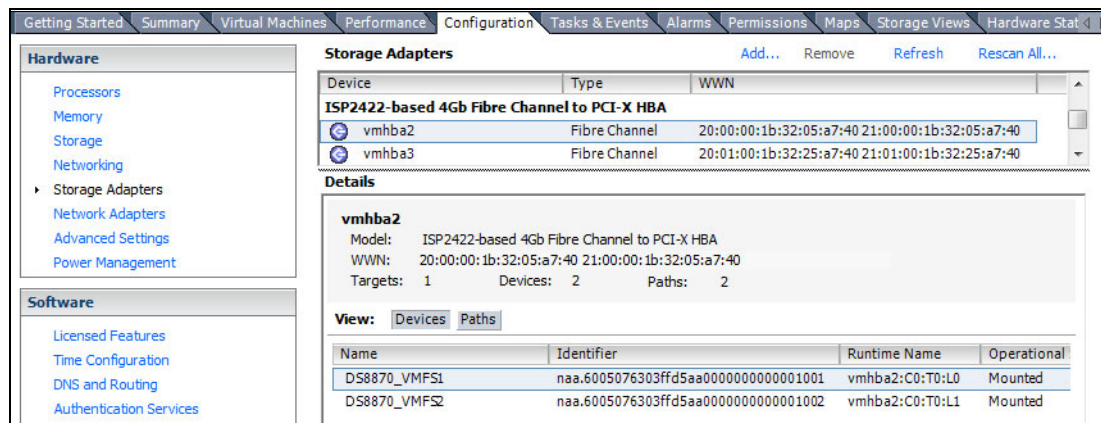


Figure 3-10 Checking attached LUNs on an HBA

After you verify access to your LUN, select **Storage** under the configuration tab and click **Add Storage**. Complete the following steps:

1. Select Disk/LUN as your Storage Type.
2. Select the Device on which you want to create a datastore. Make sure that the Device Name and the Datastore Name are the same. This helps later when you are trying to find relationships between VMware datastores and DS8870 Volume Names.
3. You are prompted to create a VMFS-5 or VMFS-3 file system. If you do not use VMFS-5, you cannot create a datastore larger than 2 TiB. If you use VMFS-5, you are not prompted to define a maximum block size as this is no longer possible with VMFS-5.

## 3.6 Multipathing

In VMware you can select three different path policies within the Path Selection plug-in (PSP), which is a software component in the Plugable Storage Architecture (PSA). The PSA is a special VMkernel layer that manages storage multipathing.

The following options are available to set the multipathing policy in VMware:

► **Fixed**

All I/O is always routed over one of the available paths. The user can set one path as the preferred path that is always used. If this path fails, a failover to one of the remaining paths is done. If the preferred path returns, a failback is done.

► **Most Recently Used (MRU)**

Similar to fixed, all I/O is always routed over one of the available paths. In case of a failure, a failover to another path is done. But if the initial path is recovered, no failback occurs.

► **Round Robin (RR)**

This PSP is the only one that evenly uses all available paths. I/O is routed over all active paths alternating. If one path fails, it is skipped. After its recovery, it is used for I/O again automatically.

There is also an option within the vSphere plug-in to set the default multipathing policy for new devices. You can change this setting by right-clicking the DS8870 storage subsystem and selecting **Set Multipath Policy Enforcement**, as shown in Figure 3-11.

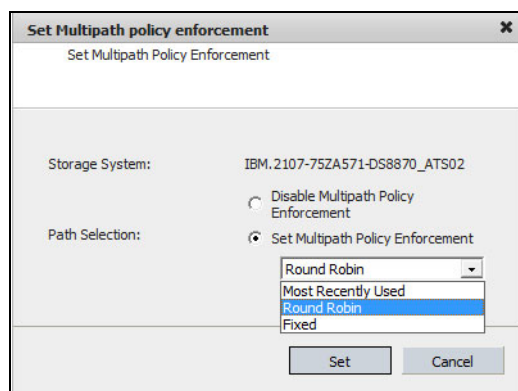


Figure 3-11 Setting Multipath Policy Enforcement

**Round robin:** The recommended setting for multipath policy is round robin.

The following configuration maximums are documented for vSphere 5.0 and vSphere 5.1:

- ▶ The maximum number of LUNs per server is 256.
- ▶ The maximum number of paths per server is 1024.
- ▶ The maximum number of paths per LUN is 32.

These facts have some important design considerations. For example, if each DS8870 volume can be accessed through eight fabric paths, the maximum number of volumes is 128. Dropping the paths to four increases the maximum LUN count to 256.

For more information, see this website:

<http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf>

## 3.7 Host attachment

In this publication, we do not address how to attach VMware hosts to the DS8870. For more information, see *IBM System Storage DS8000: Host Attachment and Interoperability*, SG24-8887.





# Business continuity with IBM DS8870 Storage System and VMware Site Recovery Manager

This chapter describes the combined solution of the IBM DS8870 Storage System and VMware Site Recovery Manager for business continuity and disaster recovery.

For the convenience of those who want to implement the solution, the last section of the chapter contains useful information and references about how to deploy a DS8870 and SRM solution.

This chapter includes the following topics:

- ▶ Introduction to business continuity and disaster recovery solutions
- ▶ VMware Site Recovery Manager
- ▶ Disaster Recovery Solution with DS8870 and SRM 5.x

## 4.1 Introduction to business continuity and disaster recovery solutions

As IT systems have become increasingly critical to the smooth operation of companies, the importance of ensuring their continued operation (including rapid recovery when subsystems fail) has increased.

Planning for all contingencies that are necessary to meet essential business-driven objectives during events ranging from site-wide disasters to equipment outages constitutes one of the greatest challenges faced by IT professionals. This challenge is amplified when applied to complex, densely-provisioned virtual infrastructures.

Fundamentally, these complexities and challenges stem from the fact that they necessitate a seamless, end-to-end scope of integration that coordinates the underlying mechanics of multiple, dependent components that span all layers of the infrastructures.

Clearly, the scope of potential complexity that is involved in realizing the rapid, reliable transition of data center services between geographically dispersed sites means that even relatively trivial errors in planning or execution can be devastating to the business. This is especially true in the absence of a comprehensive disaster recovery framework featuring powerful orchestration with turn key execution.

VMware Site Recovery Manager (SRM) with IBM DS8870 advanced replication technology provides a comprehensive, service-oriented framework that jointly spans disaster recovery and avoidance methodologies while realizing predictable, reliable, and efficient delivery of business continuity services and objectives

For planning the disaster recovery solution, the organization's business continuity plan should indicate the key metrics of Recovery Point Objective (RPO) and Recovery Time Objective (RTO) for various business processes. The metrics that are specified for the business processes must then be mapped to the underlying IT systems and infrastructure that support those processes. After the RTO and RPO metrics are mapped to IT infrastructure, the Disaster Recovery (DR) planner can determine the most suitable recovery strategy for each system.

RPO refers to the point in time to which data must be recovered as defined by the organization. It reflects the maximal amount of data (within the specified RPO time frame) that is acceptable to lose upon failure or unavailability of the main peer. While RTO refers to the duration of time and a service level within which a business process must be restored following a disaster or disruption to avoid unacceptable consequences that are associated with a break in business continuity. RTO defines how much downtime is acceptable for a particular application or infrastructure.

## 4.2 VMware Site Recovery Manager

VMware vCenter Site Recovery Manager (SRM) represents a purpose-built business continuity solution that empowers administrators to implement robust, customized, end-to-end disaster recovery, disaster avoidance, and data center migration strategies.

SRM includes features for seamless automation, simplified management, and functionality that addresses testing and execution of the planned or unplanned relocation of data center services from one site to another.



VMware vCenter SRM uses the IBM DS8870 Copy Services capabilities to create a copy of the data at a secondary location. Software for replicating data on the DS8870 Storage System is included with every system. For more information about DS8870 Copy Services and licensing requirements, see 4.3.3, “DS8870 Remote Mirror and Copy functions” on page 46.

Protected and recovery sites are noteworthy concepts that are applied to SRM solutions with regard to where data is located. In a typical SRM installation, the protected site provides business-critical data center services. The recovery site provides an alternative facility to which these services can be migrated. The protected site can be any site where virtual infrastructure supports a critical business need. The recovery site can be located thousands of miles away or in the same site.

Still, to this extent, a container for VMs and templates that use the same replicated data store group is called *Protection group*. It consists of pointers to the replicated vSphere data stores that contain collections of VMs that get failed over from the protected site to the recovery site during actual disaster recovery or testing operations. Conceptually, SRM protection groups specify the relationship of protected virtual machines in the same way that DS8870 consistency groups specify recovery relationships among logical volumes.

Within SRM configuration, a recovery plan specifies how the virtual machines in a specified set of protection groups are recovered.

For more information about SRM, see the following VMware Support website:

[http://www.vmware.com/support/pubs/srm\\_pubs.html](http://www.vmware.com/support/pubs/srm_pubs.html)

## 4.2.1 VMware SRM 5.x Disaster Recovery enhancements

VMware Site Recovery Manager 5.x enhances the ability to build, manage, and execute reliable disaster recovery plans spanning bi-direction relocation of data center services, and thus provides unprecedented levels of protection. The sophistication of SRM disaster recovery strategies has expanded through the addition of the following capabilities:

- ▶ **Planned Migration**

With version 5.0, VMware introduced a new workflow that is designed to deliver migration from a protected site to a recovery site through execution of a planned migration workflow. Planned migration ensures an orderly and pretested transition from a protected site to a recovery site while minimizing the risk of data loss. Planned migration runs in a highly controlled fashion. It halts the workflow if an error occurs, which differs from a traditional disaster-event recovery plan.

- ▶ **Automated re-protection**

Automated re-protect enables the environment at the recovery site to establish replication and protection of the environment back to the original protected site through a single click. The re-protection capability allows SRM to take advantage of the advanced DS8870 remote mirroring function to wrap the roles of the primary and DR sites following a failover, thereby continuing data center protection while meeting RPOs and RTOs regardless of which site is currently operating in the role of the primary.

- ▶ Automated failback

Automated failback returns the entire environment to the originally protected primary site. This can only happen after re-protection ensures that data replication and synchronization were established to the original primary site. Failback runs the same workflow that was used to migrate the environment to the protected site, which ensures that the critical systems that are encapsulated by the recovery plan are returned to their original environment.

- ▶ Enhanced dependency definition

This feature organizes the failover workflow to enforce the order in which virtual machines are restarted at the DR site. Starting with SRM 5.x, the number of priority groups that are available to vSphere administrators was increased, which permits SRM to recognize virtual machine dependencies within a priority group.

## 4.3 Disaster Recovery Solution with DS8870 and SRM 5.x

With the combined solution of the IBM DS8870 System Storage and VMware vCenter Site Recovery Manager, the activities that are required for meeting business continuity objectives before, during, and after a critical event that results in site-wide interruption of service are vastly simplified by making use of the following features:

- ▶ Drag-and-drop storage replication configuration
- ▶ Automatic discovery of replicated volumes
- ▶ Policy-driven power up procedures at the Disaster Recovery (DR) site
- ▶ Flexible failover and failback
- ▶ High performance incremental flashcopies for testing DR without disruption to existing replication

### 4.3.1 DS8870 and SRM Solution Architectural Overview

DS8870 is integrated to VMware SRM by means of a Storage Replication Adapter (SRA), which is the software that is required for Site Recovery Manager to issue replication commands to the DS8870 array.

As shown in Figure 4-1 on page 45, the IBM DS8870 SRA functions as an interface between the DS8870 Storage Systems and SRM.

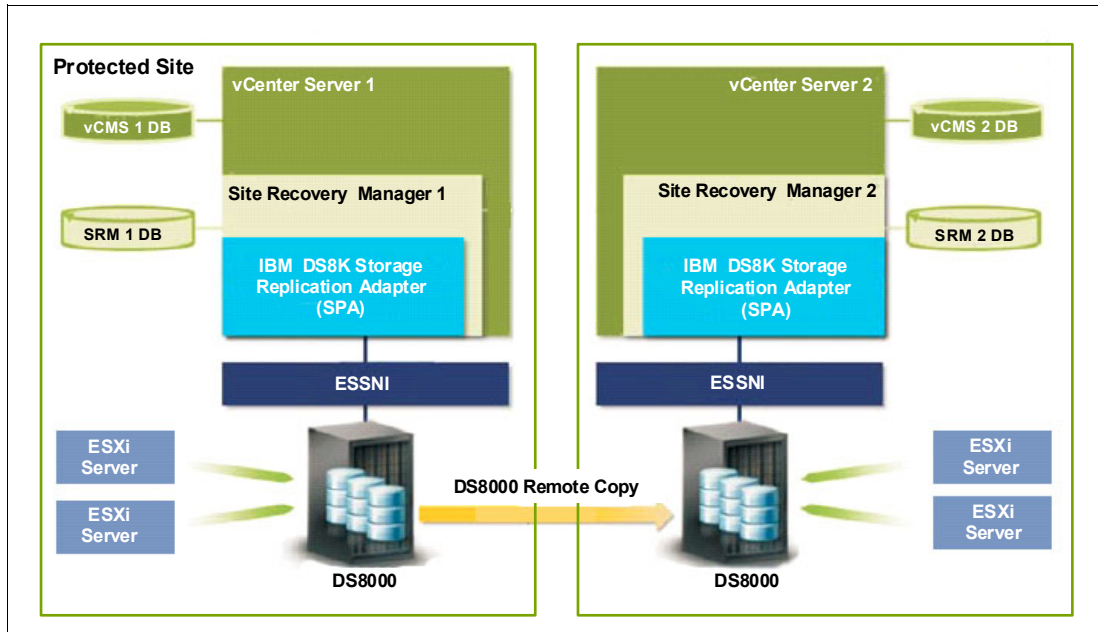


Figure 4-1 DS8870 and SRM Solution Architectural Overview

For example, SRA enables SRM commands that encompass the execution of the following vSphere-centric tasks and workflows (including querying replicated data stores and promoting replicated data stores) to proceed in concert with DS8870-specific functions including remote mirroring and snapshot management:

- ▶ Discovering LUNS and their associations in the context of the remote mirroring relationship (or relationships) that spans DS8870 systems at the primary and disaster recovery sites.
- ▶ Executing test failover, which is used to test the implementation of the planned disaster recovery workflows by starting DS8870 snapshots to create copies of the data stores without impacting operations at the primary or disaster recovery environments.
- ▶ Automating the failover of a DS8870 system at the primary SRM site to a DS8870 system at a recovery (secondary) SRM site.

Immediately upon a failover, the ESX/ESXi servers at the secondary SRM site start using the replicated data stores on the mirrored volumes of the secondary DS8870 system.

- ▶ Starting re-protect for the entirety of the VMware environment or a subset.
- ▶ Executing a failback following a previously completed failover and re-protect.

The IBM DS8870 SRA extends SRM capabilities and allows it to seamlessly use DS8870 replication and mirroring as part of the SRM comprehensive Disaster Recovery Planning (DRP) solution.

### 4.3.2 IBM DS8870 Support for VMware SRM 5.x

IBM DS8000-series Storage Systems that are running microcode releases 2.x, 3.x, 4.x, 5.x or 6.x started supporting VMware SRM versions 4.0 and 4.1 with SRA 1.30.xxxx. For Site Recovery Manager 5.x, SRA version 2.1.x was released and it supports DS8870 running microcode version 7.x and DS8800 starting on microcode R6.3.

**Important:** At the time of writing, IBM DS8870 System Storage is fully supported by VMware Site Recovery Manager latest version, 5.1, with SRA 2.1.x. Also, it supports the previous SRM versions 5.0 Update 2, 5.0 Update 1 and 5.0, by using the same SRA version, 2.1.x.

With DS8870 and SRM 5.x, you can make full use of features like Planned Migration, Automated Re-protection, and Automated Failback, which were described in 4.3.1, “DS8870 and SRM Solution Architectural Overview” on page 44.

For the most up-to-date compatibility information about the SRM and DS8870 integration, see this website:

<http://www.vmware.com/support/srm/srm-storage-partners.html>

### 4.3.3 DS8870 Remote Mirror and Copy functions

VMware Site Recovery Manager relies on DS8870 Remote Mirror and Copy functions to have data copied from the primary to the recovery site. For this purpose, SRM supports DS8870 Metro Mirror, Global Mirror, and Metro/Global Mirror on this inter-site replication. Also, SRM uses DS8870 point-in-time copy function FlashCopy.

The DS8870 Remote Mirror and Copy (RMC) functions provide a set of flexible data mirroring techniques that allow replication between volumes on two or more storage systems. Remote Mirror and Copy functions are optional features of the DS8870. To use them, you must purchase the DS8870 Series Function Authorization with the appropriate feature codes.

**Tip:** VMware Site Recovery Manager with DS8870 supports Tivoli Storage Productivity Center for Replication managing the Remote Mirror and Copy functions.

#### **Metro Mirror**

Metro Mirror provides real-time mirroring of logical volumes between two DS8000s that can be located up to 300 km (186.41 miles) from each other. It is a synchronous copy solution where write operations are completed on both copies (local and remote site) before they are considered to be complete. It is typically used for applications that cannot suffer any data loss in the event of a failure.

As data is synchronously transferred, the distance between the local and the remote disk storage systems determines the effect on application response time.

#### **Global Mirror**

Global Mirror provides a long-distance Remote Copy feature across two sites by using asynchronous technology. With Global Mirror, the data that the host writes to the Storage System at the local site is asynchronously shadowed to the Storage System at the remote site. A consistent copy of the data is automatically maintained on the Storage System at the remote site.

Global Mirror operations provide the benefit of supporting operations over unlimited distances between the local and remote sites, which are restricted only by the capabilities of the network and the channel extension technology. It can also provide a consistent and restartable copy of the data at the remote site, which is created with minimal impact to applications at the local site.

The ability to maintain an efficient synchronization of the local and remote sites with support for failover and failback modes helps to reduce the time that is required to switch back to the local site after a planned or unplanned outage.

### **Metro/Global Mirror**

Metro/Global Mirror combines a Metro Mirror and a Global Mirror to provide the possibility to implement a three-site disaster recovery solution. The production system is using the storage at the local site, which is replicated synchronously by using Metro Mirror to an intermediate site. The secondary volumes of the Metro Mirror relationships are used further on as the primary volumes to the cascaded Global Mirror relationships, which replicate the data to the remote disaster recovery site.

This configuration provides a resilient and flexible solution for recovery in various disaster situations. The customer also benefits from a synchronous replication of the data to a close location that acts as the intermediate site. It also enables the possibility to copy the data across almost unlimited distances, where data consistency can be provided in any time in each location.

### **FlashCopy**

By using the FlashCopy function, you can make point-in-time, full-volume copies of data with the copies immediately available for read or write access.

FlashCopy creates a copy of a source volume on the target volume. This copy is called a *point-in-time copy*. When you start a FlashCopy operation, a FlashCopy relationship is created between a source volume and target volume. A FlashCopy relationship is a mapping of the FlashCopy source volume and a FlashCopy target volume. This mapping allows a point-in-time copy of that source volume to be copied to the associated target volume. The FlashCopy relationship exists between this volume pair from the time that you start a FlashCopy operation until the storage unit copies all data from the source volume to the target volume or you delete the FlashCopy relationship, if it is a persistent FlashCopy.

**Important:** FlashCopy is supported on SRM as a Copy or NoCopy relationships. FlashCopy SE is also supported on SRM environments.

For more information about IBM DS8000 Remote Mirror and Copy functions, see *IBM System Storage DS8000 Copy Services for Open Systems*, SG24-6788.

## **4.3.4 DS8870 Remote Mirror and Copy's Consistency Groups Support**

A consistency group consists of volumes that share the same atomic point in time (RPO). When a recovery is issued at the remote site, all volumes recover at the same point. In Site Recovery Manager environments, a consistency group must only be used if actual dependencies across DS8870 volumes exist (such as data stores with multiple extents).

The Consistency Group is applied in the storage layer. All DS8870-replicated volumes that participate in SRM and belong to the same Remote Mirror Consistency Group are displayed under a single Local Consistency Group in SRM. Take extra care while replicated volumes are added to Consistency Groups. All volumes that are used by a single virtual machine or an application on the VM should be added to the same Consistency Group.

### 4.3.5 References for VMware SRM implementation with DS8870 systems

If you want to implement a Disaster Recovery solution with VMware Site Recovery Manager and IBM DS8870 System Storage, you must adhere to the following basic guidelines:

1. Ensure you meet all of the following prerequisites for installing the IBM DS8000 SRA:
  - The local DS8000 system at the protected site has mirroring connectivity with the target DS8000 system at the recovery site.
  - Hosts on the Recovery (secondary) site have an FC connection to the DS8000 on that site and are defined as hosts on this DS8000.
  - The name of each DS8000 system is unique in the protected and recovery sites.
  - The storage extent pools that contain the replicated volumes at the protected and recovery sites have sufficient free capacity for creating the snapshots of all replicated volumes concurrently.
  - Remote mirroring is defined for all DS8000-based volumes that you intend to protect as part of your Disaster Recovery Plan (DRP).
  - Your protected (primary) volumes are mapped to the protected ESX hosts.
2. Connect the SRM instances at both sites by specifying the IP addresses and authentication credentials of the vCenter Servers.
3. Set up the following Inventory Mappings between sites:
  - Resource Mappings specify the resources that are recovered for protected virtual machines.
  - Folder Mappings correlate the folders between the sites.
  - Network Mappings link the vSphere data center networks of the two sites.
4. Assign Placeholder Data stores at each site:. These data stores serve as repositories for small virtual machine files that function as placeholders, which can be registered and activated when SRM restores operations at the recovery site.
5. Specify alarms and permissions, as required.
6. Add and configure Array Managers.
7. Create Protection Groups.
8. Create a Recovery Plan.
9. Test and Cleanup Recovery Plans.
10. Start Recovery.
11. Start Reprotect (SRM 5.X only).

For more information, see the following resources:

- ▶ *IBM System Storage DS8000 for VMware vCenter Site Recovery Manager 5.x User Guide (Version 2.1.0)*, SC27-4232-00, which is available at this website:  
[http://pic.dhe.ibm.com/infocenter/strhosts/ic/topic/com.ibm.help.strhosts.doc/PDFs/DS8000\\_for\\_VMware\\_SRM\\_2.1.0\\_UG.pdf](http://pic.dhe.ibm.com/infocenter/strhosts/ic/topic/com.ibm.help.strhosts.doc/PDFs/DS8000_for_VMware_SRM_2.1.0_UG.pdf)
- ▶ *Site Recovery Manager Installation and Configuration: vCenter Site Recovery Manager 5.1*, which is available at this website:  
<http://pubs.vmware.com/srm-51/topic/com.vmware.ICbase/PDF/srm-install-config-5-1.pdf>

# Related publications

The publications that are listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

## IBM Redbooks

The following IBM Redbooks publications provide more information about the topic in this document. Note that some publications referenced in this list might be available in softcopy only:

- ▶ *IBM System Storage DS8870 Architecture and Implementation*, SG24-8085
- ▶ *DS8870 Product Guide*, TIPS-0915
- ▶ *IBM System Storage DS8000 Host Attachment and Interoperability*, SG24-8887
- ▶ *IBM System Storage DS8000 Copy Services for Open Systems*, SG24-6788

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft, and other materials at this website:

<http://www.ibm.com/redbooks>

## Other publications

The following publications are also relevant as further information sources:

- ▶ VMWare: *What's New in VMware vSphere@ 5.0 – Storage*:  
<http://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-50-Storage-Technical-Whitepaper.pdf>
- ▶ VMWare: *What's New in VMware vSphere@ 5.1 – Storage*:  
<http://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-51-Storage-Technical-Whitepaper.pdf>
- ▶ IBM: Implement a Disaster Recovery solution with VMware Site Recovery Manager and IBM DS8870 System Storage:  
*IBM System Storage DS8000 for VMware vCenter Site Recovery Manager 5.x User Guide (Version 2.1.0)*:  
[http://pic.dhe.ibm.com/infocenter/strhosts/ic/topic/com.ibm.help.strghosts.doc/PDFs/DS8000\\_for\\_VMware\\_SRM\\_2.1.0\\_UG.pdf](http://pic.dhe.ibm.com/infocenter/strhosts/ic/topic/com.ibm.help.strghosts.doc/PDFs/DS8000_for_VMware_SRM_2.1.0_UG.pdf)
- ▶ VMWare: *Site Recovery Manager Installation and Configuration: vCenter Site Recovery Manager 5.1*:  
<http://pubs.vmware.com/srm-51/topic/com.vmware.ICbase/PDF/srm-install-config-5-1.pdf>

## Online resources

The following websites are also relevant as further information sources:

- ▶ BM System Storage Interoperability Center (SSIC):  
<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>
- ▶ VMware Compatibility Guide:  
<http://www.vmware.com/resources/compatibility>
- ▶ VMware ESXi 5.1 Update 1 Release Notes:  
<http://www.vmware.com/support/vsphere5/doc/vsphere-esxi-51u1-release-notes.html#resolvedissuesstorage>
- ▶ IBM VAAI driver installation guide available on IBM Information Center:  
[http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strhosts.doc%2Fhsg\\_virtualization.html](http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strhosts.doc%2Fhsg_virtualization.html)
- ▶ IBM Storage Management Console for VMware vCenter user guides for different versions of the plug-in can be found in IBM Information Center:  
<http://pic.dhe.ibm.com/infocenter/strhosts/ic/index.jsp?topic=%2Fcom.ibm.help.strhosts.doc%2Fvcplugin-homepage.html>
- ▶ Refer to VMware vCenter Site Recovery Manager Storage Partner Compatibility Matrix for the most up-to-date compatibility information on the SRM and DS8870 integration:  
<http://www.vmware.com/support/srm/srm-storage-partners.html>

## Help from IBM

IBM Support and downloads:

<http://www.ibm.com/support>

IBM Global Services:

<http://www.ibm.com/services>







# IBM DS8870 and VMware Synergy



**Support for Block Zero (Write Same) VAAI primitive**

**IBM vCenter plug-in to support VAAI on ESX 4.x**

**VMware Site Replication Manager 5.0 support**

This IBM Redpaper publication describes the support of the IBM System Storage DS8870 for VMware vSphere in general and more details on VMware vSphere Storage APIs Array Integration (VAAI) primitives, the IBM Storage Management Console for VMware plug-in and the IBM Site Replication Adapter (SRA) for VMware Site Recovery Manager (SRM).

Beside a high level overview of the VMware vSphere suite the paper emphasizes on the interaction of the software components with the IBM System Storage DS8870. It is intended for IT professionals who want an understanding of advantages like off-loading specific tasks to storage arrays, centralized service management and simplified implementation of disaster recovery processes in virtualized environments.

## **INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION**

### **BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:**  
[ibm.com/redbooks](http://ibm.com/redbooks)