

Using the IBM Spectrum Accelerate Family in VMware Environments

IBM XIV, IBM FlashSystem A9000 and IBM FlashSystem A9000R, and IBM Spectrum Accelerate

Oscheka Markus

Bert Dufrasne

Grant Kabobel

Abilio Oliveira



 **Cloud**

Storage



International Technical Support Organization

**Using the IBM Spectrum Accelerate Family in VMware
Environments: IBM XIV, IBM FlashSystem A9000 and
IBM FlashSystem A9000R, and IBM Spectrum
Accelerate**

May 2018

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

Second Edition (May 2018)

This edition applies to Version 12.2.1 of the IBM Spectrum Accelerate family of software with Hyper-Scale Manager Version 5.4 and Spectrum Connect Version 3.4.

This document was created or updated on February 4, 2019.

© Copyright International Business Machines Corporation 2017, 2018. 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

Notices	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	x
Comments welcome	x
Stay connected to IBM Redbooks	x
Chapter 1. IBM Spectrum Accelerate family and VMware	1
1.1 Introduction	2
1.2 IBM Spectrum Accelerate Family integration with VMware	4
1.3 vStorage API for Array Integration support	4
1.4 vSphere deployment flexibility with an IBM storage system	4
1.5 VMware Horizon 7	6
Chapter 2. VMware integration	9
2.1 Integration concepts	10
2.1.1 vSphere storage architectural overview	11
2.1.2 vStorage API for Array Integration	13
2.1.3 VMware thin provisioning	16
2.2 VMware vRealize Suite	22
2.2.1 VMware vCenter Server and vSphere Web Client	22
2.2.2 IBM Storage Enhancements for VMware vSphere Web Client	22
2.2.3 IBM Storage Provider for VMware VASA	22
2.2.4 VMware vSphere Virtual Volumes	23
2.2.5 VMware vRealize Operations Manager	23
2.2.6 VMware vRealize Orchestrator	23
2.3 Assumptions for the use cases in this paper	23
Chapter 3. Attaching VMware ESXi	25
3.1 VMware ESXi 5.5 and 6.0	26
3.1.1 ESXi connectivity	26
3.1.2 Installing host bus adapter drivers	32
3.1.3 Identifying the ESXi host port WWN	33
3.1.4 Creating a host and volume, and mapping new LUNs	34
3.1.5 Attaching an ESXi host	38
3.1.6 Configuring the ESXi host for multipathing	40
3.1.7 Performance tuning tips for ESXi hosts	42
3.2 VMware ESXi 5.0 and 5.1	44
3.2.1 ESXi 5.0 and 5.1 Fibre Channel configuration	44
3.2.2 Performance tuning tips for ESXi 5.0 and 5.1 hosts	45
3.2.3 Creating data stores that are larger than 2 TiB	49
Chapter 4. IBM Spectrum Connect software	51
4.1 IBM Spectrum Connect overview	52
4.1.1 IBM Spectrum Connect management	52
4.1.2 IBM Spectrum Connect advantages	53
4.2 IBM Spectrum Connect first-time configuration	54

4.2.1 Initial Setup Wizard	55
4.2.2 Setting up the VASA credentials.	58
4.2.3 Adding an IBM storage system as a storage array	58
Chapter 5. VMware Virtual Volumes	61
5.1 Introduction to VMware vSphere Virtual Volumes.	62
5.1.1 VMware vSphere Virtual Volumes with IBM XIV.	62
5.1.2 Implementing VMware vSphere Virtual Volumes on an XIV Storage System . . .	63
5.1.3 VVoLs concepts mapping in IBM Spectrum Connect	65
5.2 Defining Virtual Volumes in XIV	65
5.2.1 Prerequisites and configuration	65
5.2.2 XIV configuration.	66
5.2.3 VMware vCenter and IBM Spectrum Connect configuration.	79
Chapter 6. vSphere Web Client	87
6.1 vSphere Web Client illustration	88
6.2 IBM Spectrum Connect configuration for IBM Storage Enhancements for vSphere Web Client	89
6.2.1 Adding the vCenter server in IBM Spectrum Connect	89
6.2.2 Controlling the IBM Storage Enhancements for vSphere Web Client plug-in on vCenter.	90
6.2.3 Defining a storage space	91
6.2.4 Configuring a storage service	92
6.2.5 Adding a storage resource	93
6.3 Using IBM Storage Enhancements for vSphere Web Client.	96
6.3.1 Reviewing the available storage enhancements.	96
6.3.2 Additional storage control and monitoring options	98
Chapter 7. VMware vRealize Operations Manager	101
7.1 Configuration of IBM Spectrum Connect server for vROps.	102
7.2 Install the storage management package onto vROps	103
7.3 IBM XIV dashboards in vROps	106
Chapter 8. IBM Spectrum Connect configuration for vRealize Orchestrator	111
8.1 Configuration of IBM Spectrum Connect Server for vRO	112
8.2 Running workflows in vRO	121
Chapter 9. VMware vCenter Site Recovery Manager	123
9.1 IBM Spectrum Accelerate family and VMware vCenter Site Recovery Manager	124
9.1.1 Remote Mirroring overview	125
9.1.2 VMware vCenter Site Recovery Manager overview	126
9.1.3 Minimum IBM Spectrum Accelerate family and VMware vCenter SRM solution prerequisites	129
9.1.4 VMware vCenter SRM integration with the IBM Spectrum Accelerate Family Storage Replication Adapter	131
9.1.5 VMware vCenter Site Recovery Manager operations.	132
9.2 Installing Spectrum Accelerate Family SRA for VMware vCenter SRM	138
9.2.1 Configuring the storage system for VMware vCenter SRM	139
9.3 VMware vCenter Site Recovery Manager implementation and usage	140
9.3.1 Connecting the sites	141
9.3.2 Setting up inventory mappings	144
9.3.3 Configuring placeholder data stores	149
9.3.4 Adding and configuring array managers	151
9.3.5 Creating protection groups	156

9.3.6 Creating recovery plans	159
9.3.7 Testing recovery plans	162
9.3.8 Cleanup.....	164
9.3.9 Recovery.....	166
9.3.10 The reprotect process.....	168
9.3.11 Failing back to the protected site	170
Related publications	171
IBM Redbooks	171
Help from IBM	171

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

HyperSwap®	IBM Spectrum Protect™	Redbooks (logo)  ®
IBM®	IBM Spectrum Storage™	Storwize®
IBM FlashSystem®	IBM Spectrum Virtualize™	System i®
IBM Spectrum™	Real-time Compression™	System Storage®
IBM Spectrum Accelerate™	Redbooks®	XIV®
IBM Spectrum Control™	Redpaper™	

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 is a brief overview of synergistic aspects between various VMware offerings and the IBM Spectrum™ Accelerate family, including IBM XIV® and IBM FlashSystem® A9000 and IBM FlashSystem A9000R servers.

After reviewing different integration concepts and explaining general implementation aspects for attaching the IBM Spectrum Accelerate™ family to VMware ESXi deployments, the paper focuses on components that are enabled by IBM Spectrum Connect v3.4.

This paper is intended for planning to use or implementing the IBM Spectrum Accelerate family of storage systems in a VMware environment.

Authors

This paper was produced by a team of specialists from around the world.

Oscheka Markus is an IT Specialist for Proof of Concepts and Benchmarks with the Disk Solution Europe team in Germany. He has worked at IBM for 14 years. He has performed many Proof of Concepts with Copy Services on IBM Spectrum Virtualize™ and IBM Spectrum Storage™, and Performance-Benchmarks with IBM Spectrum Virtualize and IBM Spectrum Storage. He has written extensively and acted as project lead for various IBM Redbooks® publications. He has spoken at several System Technical Universities. He holds a degree in Electrical Engineering from the Technical University in Darmstadt.

Bert Dufrasne is an IBM Certified Consulting IT Specialist and Project Leader for IBM System Storage® disk products at the International Technical Support Organization (ITSO), San Jose Center. He has worked at IBM in various IT areas. He has authored many IBM Redbooks publications and has also developed and taught technical workshops. Before joining the ITSO, he worked for IBM Global Services as an Application Architect. He holds a master's degree in Electrical Engineering.

Grant Kabobel is a Storage Services / Architect Specialist with IBM US Lab Services. He has 15 years of IT experience, working as a Storage Admin / Engineer at multiple Fortune 500 companies. He has expertise with Brocade and Cisco SAN switches, Storage Data Overwrite, XIV servers, and flash storage.

Abilio Oliveira is an IBM Certified Expert IT Specialist and works as a Client Technical Specialist in Storage at IBM Asia Pacific. He has 22 years of IT experience. He holds a bachelor degree in Computer Science with a Master in Data Information Security. He has expertise in designing storage solutions with OpenStack, IBM XIV Storage System servers, and storage for cloud technologies. He also specializes in Storage Infrastructure Optimization studies.

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

Ran Harel
Yossi Siles
Dima Isayev
Alon Marx
John Hyams
IBM

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:

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:

ibm.com/redbooks

- Send your comments in an email to:

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>



IBM Spectrum Accelerate family and VMware

This chapter is a high-level introduction to the concepts, features, and business relationships that make the IBM Spectrum Accelerate family and VMware a perfect fit for an optimized storage-server virtualization solution.

This chapter includes the following sections:

- ▶ Introduction
- ▶ IBM Spectrum Accelerate Family integration with VMware
- ▶ vStorage API for Array Integration support
- ▶ vSphere deployment flexibility with an IBM storage system
- ▶ VMware Horizon 7

1.1 Introduction

IT organizations today are looking for a tool that can deliver complete end-to-end IT services to their lines of business. The IT service might be infrastructure, applications, desktops, or something else.

The organization that provides these IT services wants these services to be delivered securely in a multi-tenant environment with the capability of self-service for the users while also ensuring compliance with the business policies.

The orchestration of the infrastructure needs multiple entities that are tightly integrated with each other and smartly responding to administrator or user needs, and that is where a software-defined environment (SDE) plays an important role in participating in the overall orchestration. Integration between service delivery, management, orchestration, automation, and hardware systems is becoming a requirement to support the emergence of SDEs.

For SDEs to provide benefits, they must manage all the components of the infrastructure, including storage, and that makes the software-defined storage (SDS) more relevant and important.

The capability of collecting the information from storage systems and providing it to a VMware virtual environment, and in turn implementing the tasks on the storage systems in an automated fashion, demonstrates the software-defined characteristics that are enabled by IBM Spectrum Connect (formerly known as IBM Spectrum Control™ Base Edition).

Virtualization technology is a transforming business. Companies are increasingly virtualizing their environments to meet these goals:

- ▶ Consolidate servers
- ▶ Centralize services
- ▶ Implement disaster recovery (DR)
- ▶ 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 underutilized resources to a new platform. Equally crucial to a server virtualization scenario is the storage itself. Implementing server virtualization without taking storage into account can cause challenges, such as uneven resource sharing and performance and reliability degradation.

IBM Storage solutions provide the speed and performance of ready data access with the agility and efficiency of hybrid cloud and SDS. By connecting data across any architecture, storage from IBM delivers deeper insights faster, which gives you the edge to outshine and outperform your competition and win in the cognitive era.

It also provides the following advantages to help meet your enterprise virtualization goals:

- ▶ End-to-end support for VMware solutions, including vSphere and vCenter
- ▶ Provides hotspot-free server-storage performance
- ▶ Optimal resource use
- ▶ An on-demand storage infrastructure that allows simplified growth

IBM collaborates with VMware on the strategic, functional, and engineering levels. IBM storage systems use this technology partnership to provide robust solutions and release them quickly.

VMware offers a comprehensive suite of products for server virtualization:

- ▶ VMware ESXi server: This production-proven virtualization layer runs on physical servers. It allows processor, memory, storage, and networking resources to be provisioned to multiple virtual machines (VMs).
- ▶ VMware Virtual Machine file system (VMFS): A high-performance cluster file system for VMs.
- ▶ VMware Virtual symmetric multiprocessing (SMP): Allows a single VM to use multiple physical processors simultaneously.
- ▶ VMware Virtual Machine: A representation of a physical system by software. A VM has its own set of virtual hardware on which an operating system and applications are loaded. The operating system sees a consistent, normalized set of hardware regardless of the actual physical hardware components. VMware VMs contain advanced hardware features, such as 64-bit computing and virtual SMP.
- ▶ vSphere Client / Web Client: An interface allowing administrators and users to connect remotely to the VirtualCenter Management Server or individual ESX installations from any Windows PC.
- ▶ VMware vCenter Server: Centrally manages VMware vSphere environments. It gives IT administrators dramatically improved control over the virtual environment compared to other management platforms. Formerly called VMware VirtualCenter.
- ▶ Virtual Infrastructure Web Access: A web interface for VM management and remote consoles access.
- ▶ VMware VMotion: Allows the live migration of running VMs from one physical server to another, one data store to another, or both. This migration has zero downtime, continuous service availability, and complete transaction integrity.
- ▶ VMware vCenter Site Recovery Manager (SRM): A business continuity and DR solution for VMware ESX servers providing VM-aware automation of emergency and planned failover/failback scenarios between data centers incorporating either server or storage-based data store replication.
- ▶ vStorage APIs for Storage Awareness (VASA): An API that facilitates the awareness of specific storage-centric attributes to vCenter. These functional and non-functional characteristics are automatically surfaced by a VASA-compatible storage subsystem and presented to vCenter to enhance intelligent automation of storage resource management with the VMware Profile-Driven Storage resource classification and deployment methodology.
- ▶ VMware Storage Distributed Resource Scheduler (DRS): Facilitates the automated management of initial VMDK placement. It also facilitates continual, dynamic balancing of VMDKs among clustered data stores by identifying the most appropriate resource candidates based on capacity, performance, and functional characteristics that are specific to the requirements of individual VMs or clusters.

Beginning in vSphere 5.0, VMware Storage DRS can take advantage of VASA-based and administrator-based storage resource classifications to realize simplification of heterogeneous storage management based on the concept of Profile-Drive Storage, which organizes diverse storage resources into profiles meeting specific classification criteria.

- ▶ VMware high availability (HA): Provides easy-to-use, cost-effective HA for applications running in VMs. If a server fails, affected VMs are automatically restarted on other production servers that have spare capacity.

- ▶ VMware Consolidated Backup (VCB): Provides an easy-to-use, centralized facility for agent-free backup of VMs that simplifies backup administration and reduces the load on ESX installations. VCB is being replaced by VMware vStorage APIs for Data Protection (VADP).
- ▶ VMware vStorage APIs for Data Protection (VADP): Allows backup software, such as IBM Spectrum Protect™ for Virtual Environments V6.2 or later, optionally with IBM Spectrum Protect Snapshot for VMware V3.1 or later and Spectrum Protect Plus 10.1.0 or later, to perform customized, scheduled centralized backups at the granularity of VMs, and recovery at the data store, VM, or file level. You do not have to run backup tasks inside each VM.
- ▶ VMware Infrastructure software development kit (SDK): Provides a standard interface for VMware and third-party solutions to access VMware Infrastructure.

1.2 IBM Spectrum Accelerate Family integration with VMware

IBM Spectrum Accelerate Family provides end-to-end support for VMware with ongoing support for VMware virtualization solutions as they evolve and are developed. Specifically, an IBM XIV server works in concert with the following VMware products and features:

- ▶ vSphere Hypervisor (ESXi)
- ▶ vCenter Server
- ▶ vStorage APIs for Data Protection (VADP)
- ▶ vSphere vMotion and Storage vMotion
- ▶ vSphere APIs for Storage Awareness (VASA) in concert with VMware DRS and Storage I/O Control (SIOC)
- ▶ vStorage API for Array Integration (VAAI)

1.3 vStorage API for Array Integration support

VAAI helps reduce host usage and increases scalability and the operational performance of storage systems, particularly in densely configured, multi-tenant virtual environments. The traditional ESX operational model with storage systems forces the ESX host to issue many identical commands to complete certain types of operations, including cloning operations. Using VAAI, the same task can be accomplished with far fewer commands, reduced contention, and with the potential to greatly reduce resource consumption at all levels along the I/O path.

1.4 vSphere deployment flexibility with an IBM storage system

The minimum implementation of a VMware virtualization environment by using an IBM storage system requires the deployment of at least one ESXi server to host the VMs and one vCenter server and vSphere client. Also, you need redundancy at both the network and SAN levels.

IBM Spectrum Protect Snapshot for VMware and IBM Spectrum Protect Plus seamlessly integrate with the advanced snapshot technology of IBM storage systems and VADP to implement robust end-to-end centralized backup at the data store or data store cluster level and restore capabilities.

With these solutions, there is no need to deploy OS-specific agents on each VM. This solution can be further enhanced by incorporating IBM Spectrum Protect for Virtual Environments to implement off-host incremental data backup (supporting VMware Changed Block Tracking) and archival processes targeting appropriate nearline or lower-tiered media. For more information about VADP, see the following VMware Knowledge Base article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1021175

To further improve the availability of your virtualized environment, simplify your business continuity and DR solution by using integrated and automated VM-aware failover and failback and implementing ESXi servers, vCenter server, and another IBM storage system products at the recovery site.

Also, install VMware vCenter Site Recovery Manager (SRM), and use the Storage Replication Adapter (SRA) to integrate SRM and VCenter with your IBM storage systems at both sites. The SRM itself can also be implemented as a VM on the ESXi server or run at the vCenter host. Both the primary data center and the DR data centers must incorporate redundant networks and SANs.

Figure 1-1 is a modular architectural overview that summarizes the key vSphere and IBM storage system integration points.

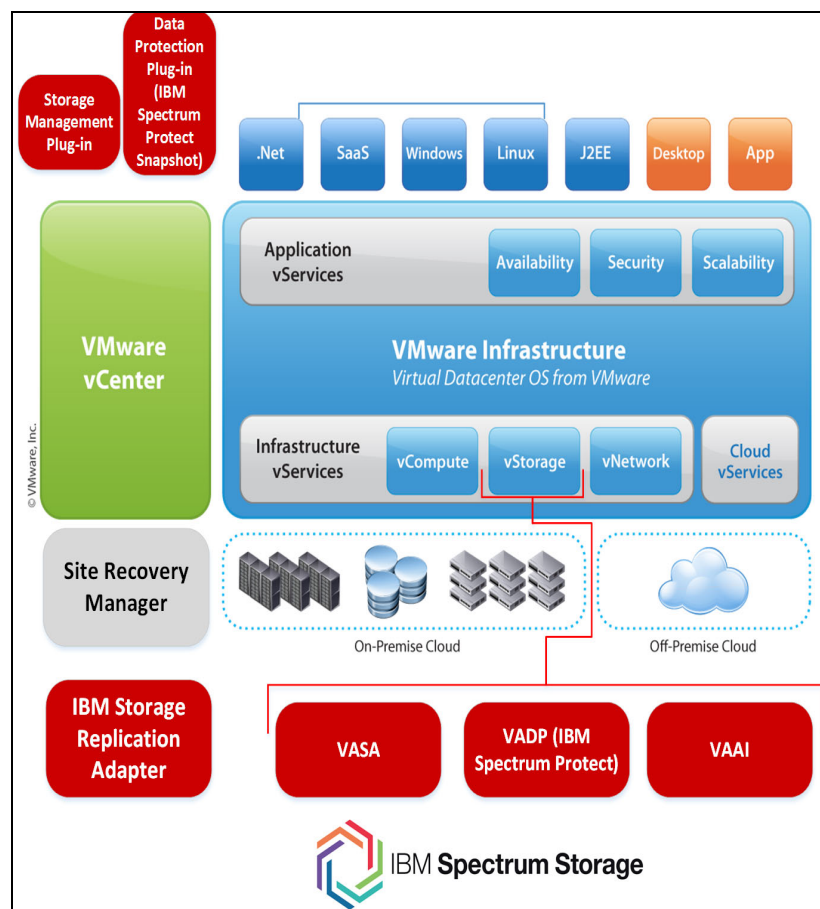


Figure 1-1 IBM System Storage Integration Points with VMware

Deep IBM storage integration with VMware is provided at no additional charge as part of IBM storage licensing. Integration works either ready-for-use, as is the case for VAAI with vSphere 5.0 or higher, or with simple software plug-ins or drivers.

Optional, separately licensed IBM Spectrum Protect products provide additional VMware data protection integration.

1.5 VMware Horizon 7

Users are using new types of devices for work, such as ones that are based on Windows, Linux, iOS, or Android, and because of these new requirements, managing and delivering services to users with traditional PC-centric tools has become increasingly difficult. Data loss and image drift are real security and compliance concerns.

VMware Horizon 7 provides an IT team with a new streamlined approach to deliver, protect, and manage Windows and Linux desktops and applications while containing costs and ensuring that users can work anytime, anywhere, on any device. This task is ensured by delivering virtual or hosted desktops and applications through a single platform to streamline management, easily entitle users, and quickly deliver desktops and applications to users across devices and locations.

Use and extend the experience with vSphere to deliver desktop and application workloads. Horizon 7 extends the power of virtualization with a software-defined infrastructure (SDI).

SDI is a framework for creating and implementing optimized IT infrastructures that can help enterprises attain a competitive advantage by delivering higher value and profitability through speed and efficiency in provisioning IT services. Most enterprise IT architectures already use virtualization to manage growth and improve agility.

IBM Spectrum Accelerate is a key component that supports the SDI framework along with Software-Defined Compute and SDN constructs. Although each of these constructs can be used separately, substantial synergy and value results from an integrated approach, which most organizations should adopt.

SDS is a new storage architecture for a wide variety of data storage requirements based on a set of loosely coupled software and hardware components. It is a model that encompasses traditional workloads (systems of record) and newer types of workload (systems of engagement), and is optimized for interoperability across hardware and software solutions.

This model provides greater flexibility around how customers can receive, consume, and explore different options for data storage, which allows customers to better harness their data for greater business insights. SDS delivers software-based storage services to SDI through these methods:

- ▶ Storage virtualization
- ▶ Automated policy-driven administration for storage management functions
- ▶ Analytics and optimization
- ▶ Backup and copy management
- ▶ Integration and API services
- ▶ Security
- ▶ Massive scale-out architecture
- ▶ Cloud accessibility

Although point SDS solutions are available, it is important to recognize an enterprise-wide SDS implementation will not be realized by installing one product offering, and will not be software only. Generally, software and hardware products and their specific features must be orchestrated to meet specific customer workload requirements in an enterprise SDI. Most IT organizations want an evolutionary transition path into SDS and SDI to gain experience, avoid risk, and preserve existing infrastructure investments.

For more information about the IBM SDI framework, see *IBM Software-Defined Storage Guide*, REDP-5121.



VMware integration

The IBM Spectrum Accelerate family and VMware have jointly designed an architecture that enables full integration between these environments. This architecture and integration offers an attractive and strategic software-defined environment (SDE).

This chapter offers a complete overview of the IBM Spectrum Accelerate family and VMware integration concepts.

It covers the following topics:

- ▶ Integration concepts
- ▶ VMware vRealize Suite
- ▶ Assumptions for the use cases in this paper

2.1 Integration concepts

At a fundamental level, the goal of both the IBM storage system and VMware's storage features is to reduce the complexity of deploying and managing storage resources. With IBM storage, administrators can provide consistent tier-1 storage performance and quick change-request cycles because they perform little planning and maintenance to keep performance levels high and storage optimally provisioned.

The underlying strategies that are devised within the vSphere storage framework to insulate administrators from complex storage management tasks, non-optimal performance, and capacity resource utilization include:

- ▶ Make storage objects much larger and more scalable, reducing the number to be managed by the administrator.
- ▶ Extend specific storage resource-awareness by attaching features and profiling attributes to the storage objects.
- ▶ Help administrators make the correct storage provisioning decision for each virtual machine (VM) or even fully automate the intelligent deployment of VM storage.
- ▶ Remove many time-consuming and repetitive storage-related tasks, including the need for repetitive physical capacity provisioning.

Clearly, vSphere relies upon the storage system to fully support several key integration features to effectively implement these strategies. Compatible storage, such as an IBM XIV, IBM Flash System A9000, IBM Flash System A9000R, or IBM Spectrum Accelerate system is essential.

To provide contrast, consider traditional storage provisioning in vSphere, which typically tasks the vSphere administrator with the following storage-centric responsibilities:

- ▶ Determine the correct data store on which to initially place a VM's virtual disk
- ▶ Continuously monitor data stores for capacity consumption
- ▶ Continuously monitor data stores for performance/latency
- ▶ Repetitively deploy physical LUN as capacity consumption grows
- ▶ Ensure that a VM remains backed by a suitable storage resource throughout its lifecycle

Additional concerns for vSphere administrators can include:

- ▶ Possible mistrust of thin provisioning due to an out-of-space condition
- ▶ Possible mistrust of physical capacity usage reporting of thin-provisioned LUNs

The remainder of this chapter addresses each of these hurdles, demonstrating the concepts and operational practices that are necessary to derive maximal value from the unique vSphere-specific capabilities of the IBM Spectrum Accelerate family. As an introduction to essential storage principles in the vSphere environment, a brief overview precedes the discussion of integration principles and preferred practices.

For more information about the interoperability options, see IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STWMS9/landing/vmware_compatibility_matrix.html

2.1.1 vSphere storage architectural overview

First, consider the vSphere storage architecture, including the physical and logical storage elements that are shown in Figure 2-1.

Although this figure is not intended to thoroughly explore vSphere storage concepts and terminology, the essential components and their relationships provide the foundational framework that is necessary to understand the upcoming integration principles.

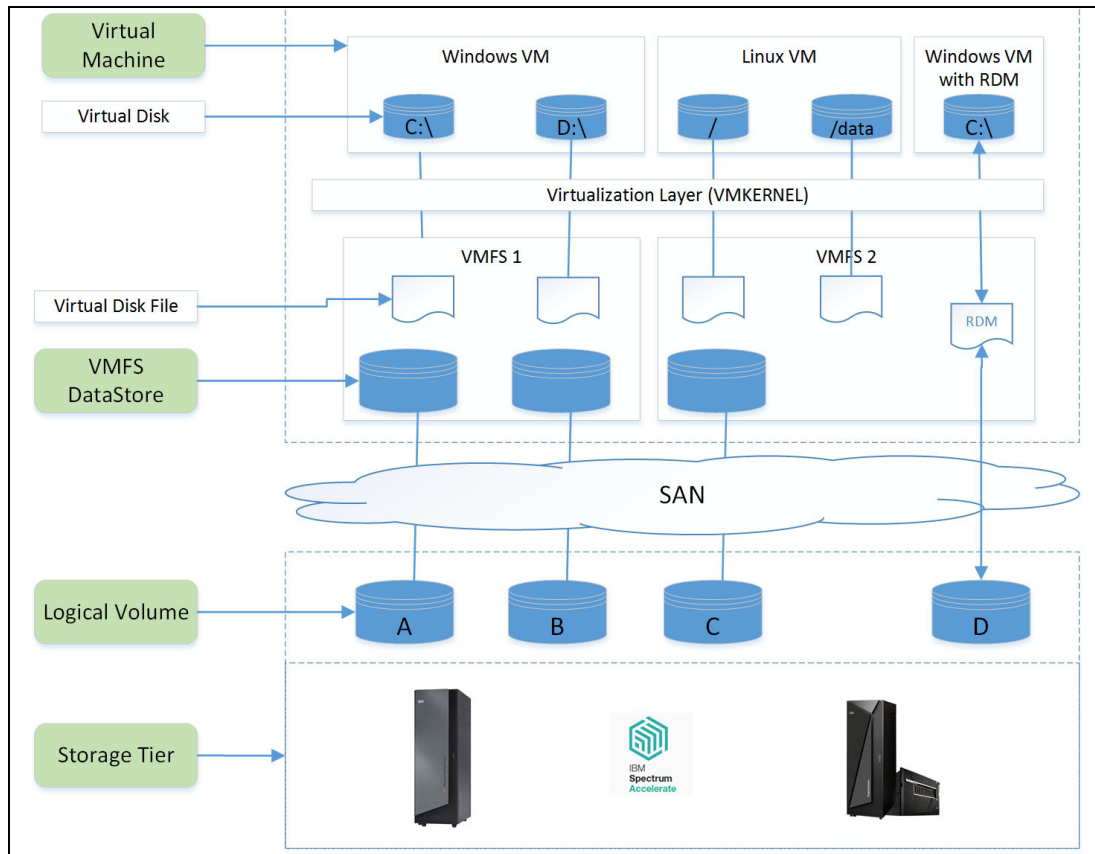


Figure 2-1 ESXi Basic Storage Elements in the vSphere Infrastructure

One of the main features that are available, starting with vSphere 6.0, is the VMware Virtual Volumes (VVol). VVols are encapsulations of VM files, virtual disks, and their derivatives. VVols are stored natively inside a storage system that is connected through Ethernet or SAN.

Three important objects in particular that are related to VVoLs are the storage provider, the protocol endpoint (PE), and the storage container. The relationship between them is illustrated in Figure 2-2.

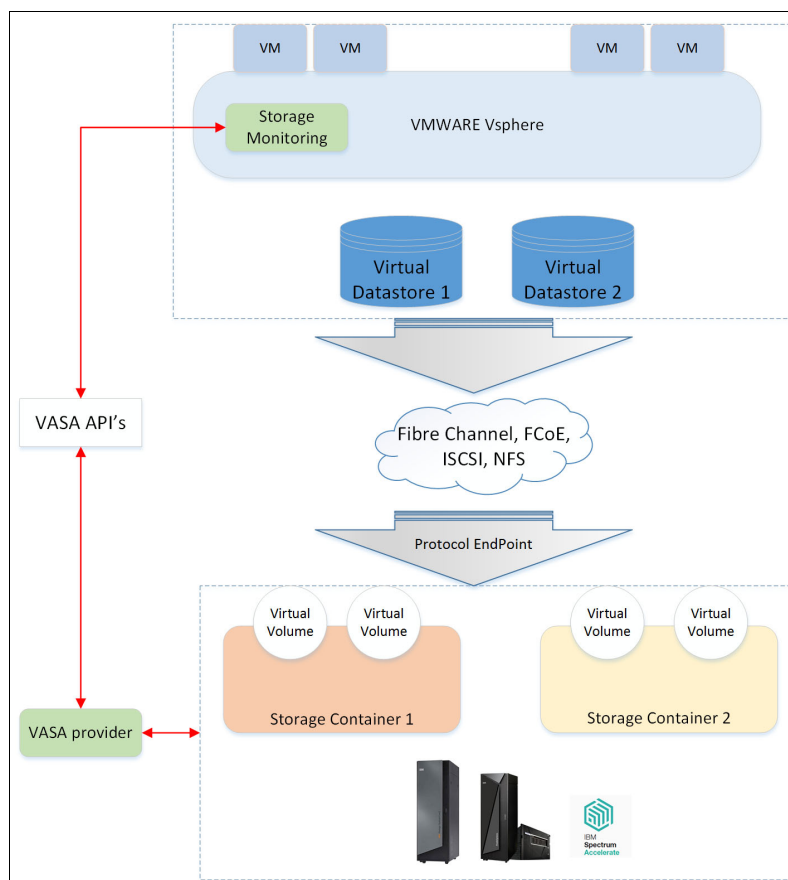


Figure 2-2 ESXi 6.0 Basic Storage Elements in the vSphere Infrastructure

VMware File System

The VMware File System (VMFS) is the central abstraction layer that acts as a medium between the storage and the hypervisor layers. The current generation of VMFS evolved to include the following distinguishing attributes, among others:

- ▶ *Clustered file system*: Purpose-built, high-performance clustered file system for storing VM files on shared storage (Fibre Channel (FC) and Internet Small Computer Systems Interface (iSCSI)). The primary goal of VMFS's design consists of functioning as an abstraction layer between the VMs and the storage to efficiently pool and manage storage as a unified, multi-tenant resource.
- ▶ *Shared data file system*: Enables multiple vSphere hosts to read and write from the same data store concurrently.
- ▶ *Online insertion or deletion of nodes*: Adds or removes vSphere hosts from VMFS volume with no impact to adjacent hosts or VMs.
- ▶ *On-disk file locking*: Ensure that the same VM is not accessed by multiple vSphere hosts concurrently.

The following sections examine concepts and preferred practices that are crucial to building an adaptable, efficient, and high-performance vSphere infrastructure with the IBM XIV Storage System's inherently cloud-optimized design and deep vSphere integration capabilities at its foundation.

2.1.2 vStorage API for Array Integration

vStorage API for Array Integration (VAAI) is an application program interface (API) framework that allows certain I/O operations to be off-loaded from the ESXi to the physical array. There are SCSI primitives that are available for block copy and block zeroing, which are used by VM Snapshots, cloning operations, Storage vMotion, and by virtual disks that are built with the Eager Zeroed Thick (EZT) option. There is another primitive called *Atomic Test and Set* (ATS), which is a superior alternative to SCSI Reservations when it comes to metadata locking on VMFS.

All of the block VAAI commands are based on standard T10 commands. For more information about T10 standards, see the following website:

<http://www.t10.org/>

VAAI helps reduce host resource utilization impact while running common vSphere operations. It also increases scalability and operational performance by offloading certain storage-centric tasks to storage systems that support the relevant commands. In contrast, traditional SCSI commands force the ESXi host to issue many repetitive commands to complete certain types of operations. These operations include cloning a VM and creating a new VM with the thick provision eager zeroed option. For example, the zeroed option writes zeros across the new virtual disk. Using VAAI, the same task can be accomplished with far less effort on the part of the ESX/ESXi server.

An IBM storage system with the correct firmware release supports the T10-compliant SCSI commands (also called primitives) to achieve this new level of integration. These commands are described in the following sections.

Hardware Accelerated Move

Hardware Accelerated Move, also known as FULL COPY or XCOPY, offloads copy operations from VMware ESXi host to the storage system. This process allows for rapid movement of data when performing copy, move, and VMware snapshot operations within the storage system. It reduces the processor and host bus adapter (HBA) workload of the ESXi server. Similarly, it reduces the volume of traffic moving through the SAN when performing VM deployment. It does so by synchronizing individual VM level or file system operations, including clone, migration, and snapshot activities, with the physical storage level operations at the granularity of individual blocks on the devices. The potential scope in the context of the storage is both *within and across* LUNs. This command has the following benefits:

- ▶ Expedites copy operations, including:
 - Cloning of VMs
 - Migrating VMs from one data store to another (Storage vMotion) on the same storage system
 - Provisioning from template
- ▶ Minimizes host processing/resource allocation
 - Copies data from one LUN to another without reading/writing through the ESXi server and network
- ▶ Reduces SAN traffic

It is important to note that the Hardware Accelerated Move SCSI primitive is used by vSphere only when the source and target LUNs are on the same storage system. For the remaining cases, vSphere implements a standard host-centric data movement process. In this case, the implication is that the SAN, the source and target hosts, and in most cases the network are all again in-band. Figure 2-3 provides a conceptual illustration contrasting a copy operation both with and without hardware acceleration.

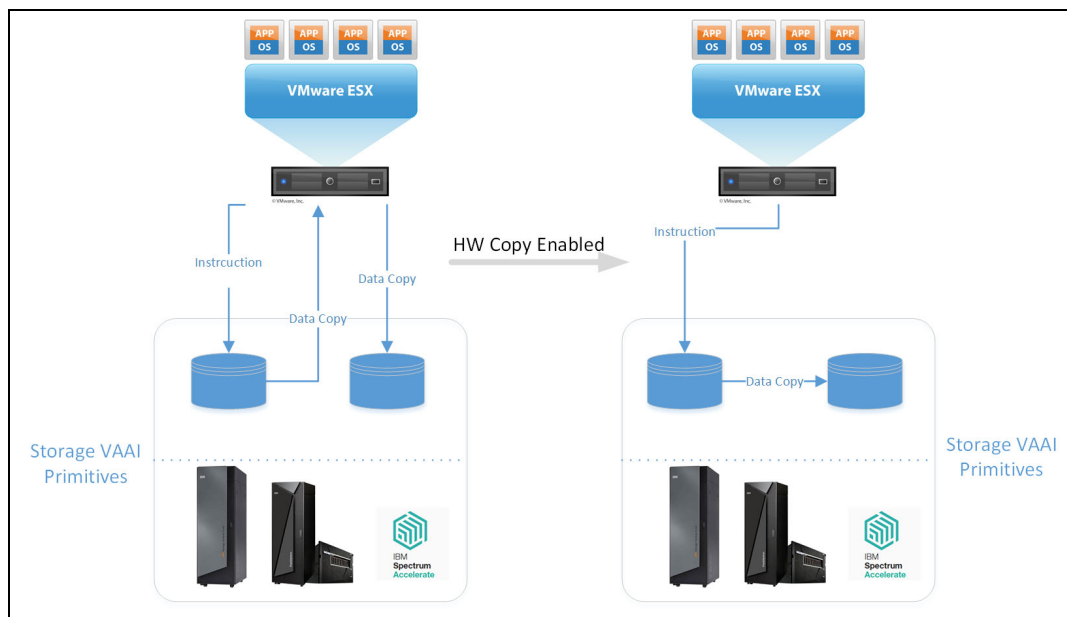


Figure 2-3 Hardware Accelerated Move primitive

Hardware Accelerated Initialization

Hardware Accelerated Initialization, or Block Zeroing, exploits the **WRITE_SAME** command to issue a chain of identical write transactions to the storage system, thus almost entirely eliminating server processor and memory utilization by eliminating the need for the host to run repetitive identical write transactions. It also reduces the volume of host HBA and SAN traffic when performing repetitive block-level write operations within VM disks to the storage system. Similarly, it allows the storage system to minimize its own internal bandwidth consumption.

For example, when provisioning a VMDK file with the **eagerzeroedthick** specification, the Zero Block's primitive issues a single **WRITE_SAME** command that replicates zeros across the capacity range that is represented by the difference between the VMDK's provisioned capacity and the capacity that is consumed by actual data. The alternative requires the ESXi host to issue individual writes to fill the VMDK file with zeros.

The Spectrum Accelerate family system further augments this benefit by flagging the capacity as "zeroed" in metadata without the requirement to physically write zeros to the cache and the disk. The scope of the Zero Blocks primitive is the VMDK creation within a VMFS data store, and the scope of the primitive is generally within a single LUN on the storage subsystem, but can possibly span LUNs backing multi-extent data stores.

In summary, Hardware Accelerated Initialization offers the following benefits:

- ▶ Offloads initial formatting of Eager Zero Thick (EZT) VMDKs to the storage system.
- ▶ Assigns zeros to large areas of storage without writing zeros from the ESXi server.
- ▶ Speeds the creation of new VMs – EZT VMDKs that are available immediately.
- ▶ Reduces elapsed time, server workload, and network workload.

Figure 2-4 provides a conceptual illustration contrasting the deployment of an EZT VMDK both with and without Hardware Accelerated Initialization.

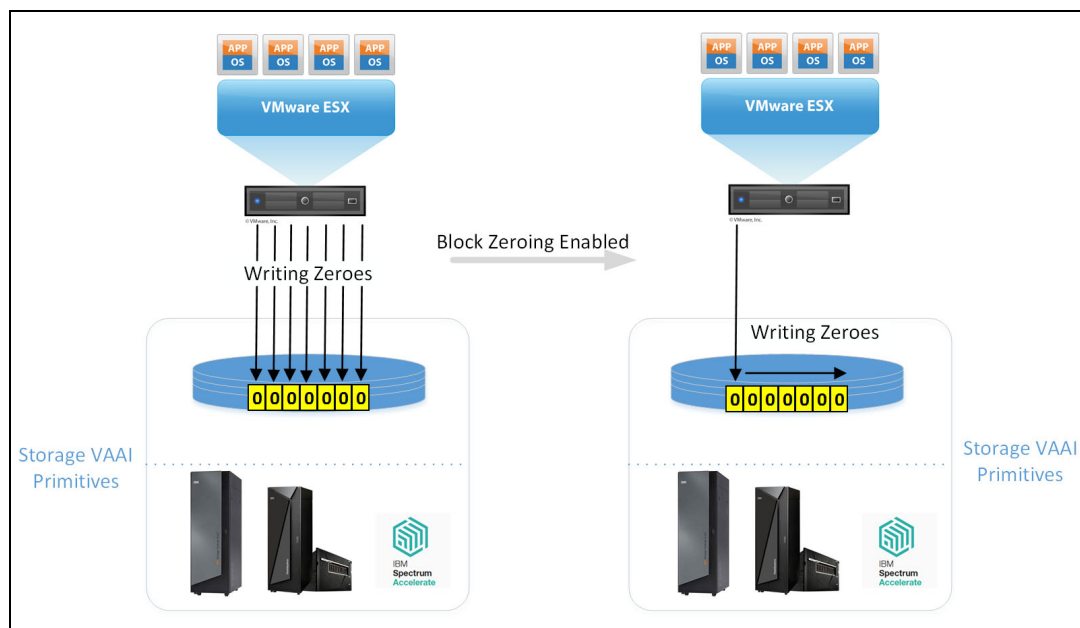


Figure 2-4 Hardware Accelerated Initialization primitive

Hardware Assisted Locking

Hardware Assisted Locking, also known as ATS, intelligently relegates resource access serialization down to the granularity of the block level during VMware metadata updates. It is the key feature of the VMware HA cluster solution.

The concept is to handle conflicting requests from multiple hosts instead of using a mature SCSI2 reserve that serializes access to adjacent ESXi hosts with a minimum scope of an entire LUN. An important note is that the VMFS version 3 or higher uses ATS in a multi-tenant ESXi cluster. This cluster shares capacity within a VMFS data store by serializing access only to the VMFS metadata that is associated with the VMDK or by using a file update that is needed through an on-disk locking mechanism. As a result, the function of ATS is identical whether implemented to grant exclusive access to a VMDK, another file, or even a Raw Device Mapping (RDM). The ATS primitive has the following advantages, which are obvious in enterprise environments where LUNs are used by multiple applications or processes at one time:

- ▶ Significantly reduces SCSI reservation contentions by locking a range of blocks within a LUN rather than issuing a SCSI reservation on the entire LUN
- ▶ Enables parallel storage processing
- ▶ Reduces latency for multiple ESXi servers that access the same LUN during common vSphere operations involving VMFS metadata updates, including the following:
 - VM/VMDK/template creation or deletion
 - VM Snapshot creation/deletion
 - VM migration and Storage vMotion migration (including when invoked by Distributed Resource Scheduler)
 - VM Power on/off
- ▶ Increases cluster scalability by greatly extending the number of ESX/ESXi hosts and VMs that can viably co-reside on a VMFS data store

Note: The currently implemented VMFS versions and the history of VMFS version deployment within a vSphere environment have important implications in the context of the scope where VMFS activities use the ATS primitive. For more information about this topic, see the following website:

<https://kb.vmware.com/s/article/1021976>

Figure 2-5 provides a conceptual illustration contrasting the scope of serialization of access both with and without Hardware Assisted Locking.

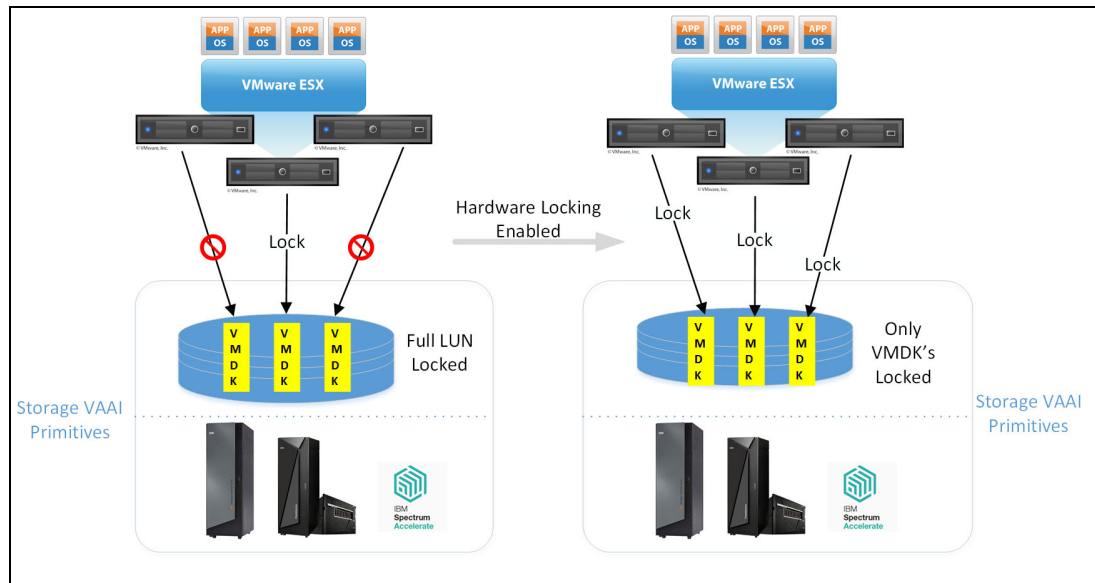


Figure 2-5 Hardware Assisted Locking primitive

2.1.3 VMware thin provisioning

In this section, the topics of thin provisioning at both the IBM Spectrum Accelerate family and at the VMware VMFS level are presented conceptually, followed by an examination of preferred practices that are necessary to most effectively combine the benefits of thin provisioning in both storage and VMFS.

Storage thin provisioning conceptual overview

Thin provisioning is a method for optimizing storage utilization by allocating to a volume only the space that is required to hold its data, deferring space allocation to the time it is needed. Fundamentally, thinly provisioned LUNs are unique in that they report a capacity to a host that is not matched by the physical capacity backing the LUN, with the result that the thin provisioning status of the LUN is transparent to the host.

To demonstrate the motivation for thinly provisioned storage, consider a simple example: As a result of the rapid growth of data, a newly deployed application exceeds a capacity utilization threshold that triggers an alert to the storage administrator to plan for expanding the available capacity. Because there is no history of data growth for the new application, the storage administrator must make a decision that involves weighing the potentially complex and time-consuming process of iteratively deploying storage capacity on an “as needed” basis, versus grossly over-provisioning the capacity to minimize effort and risk at the penalty of higher cost and the utilization inefficiencies that are incurred by deploying large “siloes” storage pools.

At a high level, thin provisioning addresses this issue by simplifying the management of storage capacity, reducing cost, and optimizing the utilization of available capacity in the context of a more broadly shared resource pool. Because the capacity that is presented to the server can be larger than the actual capacity (physical capacity) that is consumed by the server in the shared storage pool, the storage administrator can assign larger thin-provisioned volumes to the server and add the physical capacity only whenever it is necessary. As you can see, there are significant benefits to flexibility and efficiency of deployment when provisioning of capacity at the host level is decoupled from the provisioning of physical capacity.

Using IBM Spectrum Accelerate family thin provisioning

The decision to use thin provisioning is made at the storage-pool level, either regular or thinly provisioned. All volumes in a given storage pool inherit the storage pool type (*regular* or *thin*).

For thin pool, the administrator specifies:

- ▶ Pool soft size
- ▶ Pool hard size
- ▶ Additional parameters specifying behavior regarding snapshots

Changing a resource from *regular* to *thin* constitutes a simple change in designation:

- ▶ The volume type might be changed from regular to thin by moving from a regular to a thin storage pool. This is a dynamic, immediate change.
- ▶ The storage pool type might be changed from regular to thin. This is a dynamic, immediate change.

The following additional changes to the thinly provisioned resources are possible and occur dynamically and immediately:

- ▶ Change storage pool soft size.
- ▶ Change storage pool hard size.
- ▶ Move volume into/out of storage pool.

There is *zero* performance impact to these actions because IBM Spectrum Accelerate family volumes are always written thinly.

This topic is explored in depth in the following IBM Redbooks publications:

- ▶ *IBM XIV Storage System Architecture and Implementation*, SG24-7659
- ▶ *IBM FlashSystem A9000 and IBM FlashSystem A9000R Architecture and Implementation*, SG24-8345
- ▶ *IBM Spectrum Accelerate Deployment, Usage, and Maintenance*, SG24-8267

VMFS thin provisioning conceptual overview

Thin provisioning within the context of the file system follows the same basic principles as thinly provisioned volumes within storage pools, except that the provisioned elements and the associated “container” are now the VMFS files and the data store, respectively. Because the data store itself might be backed by a thinly provisioned LUN, one more layer of abstraction was added, as has one more opportunity to over-commit real capacity, in this case to the VMs themselves.

The following three format options exist for creating a virtual disk within the VMFS file system:

- ▶ Eager Zeroed Thick (EZT): Required for the best performance and for VMs classified as Fault Tolerant:
 - Space is reserved in a data store, which means unused space in the VMDK might *not* be used for other VMDKs in the same data store.
 - A VMDK is not available until formatted with zeros, either as a metadata representation in the case of the XIV Storage System or by physically writing to disk in case of storage systems that do not flag this type of activity.
 - With the VAAI WRITE_SAME (Zero Blocks) primitive, the process of zeroing the VMDK is offloaded to the storage subsystem.
- ▶ Lazy Zeroed Thick (LZT):
 - Unused space in VMDK might *not* be used for other VMDKs in the same data store.
 - The VMDK is immediately available upon creation. The VMkernel attempts to dynamically initiate the allocation of physical capacity within the storage pool by pre-emptively writing zeros to the LUN for each VM-generated write targeting new blocks. This is the default provisioning type.
- ▶ Thin:

Unused space in VMDK can be used for other VMDKs in the same data store, which adds another threshold that must be carefully monitored to prevent service interruption as a result of the VMs sharing the data store and collectively consuming all of the LUN capacity backing the data store:

 - This is possible because the specified VMDK size represents the *provisioned* size, which is what is presented to the VM itself. However, only the *used* size, or *hard* size in XIV terms, is what is subtracted from the data store’s capacity.
 - The capacity utilization percentage at the data store-level is based on the blocksize and the data that is previously written for each VMDK that is co-resident in the data store.

Like the LZT provisioning option, the VMDK is immediately available upon creation. The VMkernel attempts to dynamically initiate the allocation of physical capacity within the storage pool by pre-emptively writing zeros to the LUN for each VM-generated write that targets new blocks.

Using VMFS thin provisioning

When considering whether to thinly provision VMDKs within a VMFS data store, weigh the following advantages and disadvantages specific to the vSphere environment being implemented:

► Advantages:

- Unless the administrator effectively synchronizes capacity reclamation between VMFS data stores and the associated LUNs on the IBM storage system, which is a manual process for ESXi versions before 6.5, the potential to use thin-provisioning efficiency at the VMFS level might exceed the thin-provisioning efficiency that is possible over time within the IBM storage system because the VMFS is aware of data that was moved or deleted while the same capacity remains consumed within the storage system until capacity reclamation can occur. However, if real capacity consumption is not properly managed, the potential benefits achievable by over-representing physically backed capacity to the VMs are greatly reduced.
- Over-provisioned conditions at the VMFS level can be less frequent and generate fewer alerts because fluctuations in VMDK sizes within a data store and the associated data store capacity utilization are dynamically reflected in vCenter due to the awareness of the data consumption within the file system.

► Disadvantages:

- For vSphere releases before vSphere 5, when a thin-provisioned disk grows, the ESX host must make a SCSI reservation to serialize access to an entire LUN backing the data store. Therefore, the viability of dense VM multi-tenancy is reduced because implementing thinly provisioned VMDKs to increase multi-tenancy incurs the penalty of reducing potential performance by increasing congestion and latency.
- Compared to storage pool-based thin provisioning within IBM storage, thin provisioning at the VMDK-level has the following drawbacks:
 - There are more objects to monitor and manage because the VMDKs are thinly provisioned; therefore, they must be monitored with co-resident VMDKs in the data store. Furthermore, this must be done for all data stores. In contrast, thin-provisioning resource management can be better consolidated at the level of the XIV Storage System, thus providing global awareness of soft versus hard capacity consumption and facilitating ease of management activities, including physical capacity deployment where it really matters, that is, in the storage subsystem itself.
 - Consider the scenario of balancing physical, or hard, capacity among a group of data stores that are backed by LUNs within a storage pool whose hard capacity cannot be expanded, for example, by decreasing the size of a LUN that is associated with a given data store in favor of increasing the size of a LUN deemed to have priority. Redistributing capacity among data stores is possible, but cannot be accomplished as a single operation in vSphere at the time of writing.

In contrast, by managing the capacity trade-offs among data stores at the storage level, it is trivial to expand the soft size of both the LUN and the storage pool. The net effect is that the LUNs backing the data store that needs more hard capacity can now effectively borrow that capacity from the pool of unused hard capacity that is associated collectively with all of the LUNs in the storage pool without the need to contract the soft size of any LUNs.

Obviously, if 100% of the physical capacity in the storage pool is already consumed, this requires a coordinated expansion of capacity of the data store, LUN, and finally the physical capacity in the storage pool. If hard capacity is available in the system, the latter task can be easily accomplished within seconds using the HSM interface of the Spectrum Accelerate family system. Otherwise, it still adds more capacity. Again, capacity monitoring at all levels is a mandatory requirement to anticipate this condition.

- The scope of potential capacity utilization efficiency is relatively small at the individual data store level. Using thinly provisioned LUNs in the storage system dramatically increases the potential scope of savings by expanding the sphere of capacity provisioning to include all of the data stores that are co-resident in a storage pool because the potential savings resulting from thin provisioning is effectively proportional to the scale of the capacity pool containing thinly provisioned resources.

Thin provisioning prerequisites

Successful thin provisioning requires a “thin-friendly” environment at all levels of software in the stack:

- ▶ **File system**

VMware environments require consideration of the file systems in use by the guest operating systems and the VMFS version.

- ▶ **Database**

- ▶ **Application**

Thin-friendly file systems, databases, and applications have the following attributes:

- ▶ **Physical locality of data placement:** If data is placed randomly across the LUN, the storage system interprets the interspersed free space as being consumed.
- ▶ **Wherever possible, reuse previously freed space.** Writes are issued to previously used and deleted space before being issued to “never-used” space.
- ▶ **Provision for the file system to communicate deleted space to the storage subsystem for reclamation.**

If these properties are not pervasive across these elements, implementation of thin provisioning might have little benefit and might even incur additional penalties that are compared to regular provisioning.

In addition, be aware that the following user options and activities might affect the success of thin provisioning:

- ▶ **LUN format options.**
- ▶ **Defrag processes:** Swapping algorithms can defeat thin provisioning by touching unused space.
- ▶ **“Zero file” utilities can enable space reclamation for storage systems with zero detect or scrubbing capabilities.**

Thin provisioning general guidelines

Consider the following guidelines:

- ▶ Ensure that the following classes of applications are not included as candidates for thin provisioning:
 - Applications that are not thin-friendly.
 - Applications that are risk-averse.
 - In terms of general storage preferred practices, highest transaction applications must be excluded from consideration for thin provisioning. However, the sophisticated data distribution characteristics of the IBM Spectrum Accelerate family are designed with high transaction applications in mind, so thin provisioning can be effectively used for an expansive set of applications.
- ▶ Automate monitoring, reporting, and notifications, and set thresholds according to how quickly your business can respond.
- ▶ Plan procedures in advance for adding space, and decide whether to automate them.
- ▶ Use VAAI and the latest version of VMFS:
 - The VAAI ATS primitive limits the impact of SCSI2 reservations when thin provisioning is used.
 - Improves performance.

Reclaiming VMFS deleted blocks on thin-provisioned LUNs

vSphere 5.5 introduced a new command in the `esxcli` namespace that allows deleted blocks to be reclaimed on thin-provisioned LUNs that support the VAAI UNMAP primitive. The command can be run without a maintenance window, and the reclaim mechanism has been enhanced in the following ways:

- ▶ The reclaim size can be specified in blocks instead of a percentage value to make it more intuitive to calculate.
- ▶ Dead space is reclaimed in increments instead of all at once to avoid possible performance issues.

With the introduction of 62 TB VMDKs, **UNMAP** can now handle much larger dead space areas. However, **UNMAP** operations are still manual.

Starting with VMFS 6, space reclamation happens automatically.

It is possible to reclaim unused storage blocks on a VMFS data store for a thin-provisioned device by using the command line.

Tip: It is possible to reclaim unused storage blocks on a VMFS data store for a thin-provisioned device by running the following command:

```
[root@localhost:~] esxcli storage vmfs unmap -l 'datastore name'
```

For more information about how to use the **unmap** command, see the VMware knowledge base, found at:

<https://kb.vmware.com/s/article/2057513>

2.2 VMware vRealize Suite

IBM Spectrum Connect software enables a simplified deployment and more efficient integration of IBM storage systems and the VMware vCloud suite.

2.2.1 VMware vCenter Server and vSphere Web Client

The VMware vCenter Server component is the central management and monitoring of VMware vSphere environments. It gives IT administrators dramatically improved control over the virtual environments through a single console.

Starting with vSphere vCenter 5.5, all new features are available only through the vSphere Web Client. These features include vCenter Single Sign-on, vSphere Web Client, vCenter Inventory Service, VMware vCenter Site Recovery Manager (SRM), and vCenter Server database.

2.2.2 IBM Storage Enhancements for VMware vSphere Web Client

The IBM storage enhancements for VMware vSphere Web Client integrate into the VMware vSphere Web Client platform and enable VMware administrators to independently and centrally manage their storage resources on IBM storage systems.

Depending on the IBM storage system in use, VMware administrators can self-provision volumes (LUNs) in selected storage pools that were predefined by the storage administrators. The volumes are mapped to ESXi hosts, clusters, or data centers as logical drives that can be used for storing VMware datastores (VM data containers).

The IBM Storage Enhancements for vSphere Web Client represents a better alternative to IBM Storage Management Console for VMware vCenter. Compared to the IBM Storage Management Console for VMware vCenter, which is a plug-in that runs on each vCenter server, the IBM Storage Enhancements for vSphere Web Client are installed only on the vSphere Web Client Server, allowing multiple vCenter servers to use IBM storage resources. In addition, storage pool attachment and detachment operations are performed on the IBM Spectrum Connect side, rather than on the vSphere Client side.

IBM Storage Enhancements for VMware vSphere Web Client is provided as a plug-in on the VMware vSphere Web Client Server, and then communicates with IBM Spectrum Connect.

2.2.3 IBM Storage Provider for VMware VASA

VMware vSphere APIs for Storage Awareness (VASA) can provide information about an IBM storage-centric topology, capabilities, attributes, and storage events to vCenter Server, in real time. VMware VASA improves the ability to monitor and automate storage-related operations on VMware platforms. These functional and non-functional characteristics are automatically surfaced by a VASA-compatible storage system, and presented to vCenter to enhance intelligent automation of the storage resource management with the VMware Profile-Driven Storage resource classification and deployment methodology.

IBM Spectrum Connect and IBM Storage Provider for VMware VASA provide a standard interface for any connected VMware vCenter server that uses the VASA.

Version 2.0 of the VASA protocol introduces a new set of APIs specifically for VVOLs that are used to manage storage containers and VVOLs.

2.2.4 VMware vSphere Virtual Volumes

With vSphere 6.0, VMware officially released support for the vSphere VVoLs architecture. VVoLs allow more efficient operations and control of external storage resources, such as the XIV Storage System, running XIV Software V11.6 or later. At the time of writing, VVoLs are not supported by IBM Flash System A9000 V12.2.1, IBM Flash system A9000R V12.2.1, or IBM Spectrum Accelerate V11.5.4.

With VVoLs, the XIV Storage System becomes aware of individual VMDK files, and data operations such as snapshot and replication can be performed directly by the storage at the VMDK level rather than the entire VMFS data store.

For updates about the interoperability between VMWARE and IBM Storage, see IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/en/STWMS9/landing/vmware_compatibility_matrix.html

2.2.5 VMware vRealize Operations Manager

VMware vRealize Operations Manager (vROps) is VMware's integrated operations suite that converges performance, capacity, and configuration management. The integration of IBM Spectrum Connect with vROps allows monitoring and analysis of IBM Spectrum Accelerate family systems' health, performance, and capacity. It can dynamically cope with policy-governed workflows to maintain Service Level Agreements (SLAs).

After successfully configuring IBM Spectrum Connect with vROps, it periodically starts sending the storage system information to vROps. You can view the detailed IBM storage dashboards with the graphical relationships between the storage elements (storage systems, ports, storage pools, and volumes) and virtual elements (data stores, VMs, and hosts) in a drill-down interactive style.

2.2.6 VMware vRealize Orchestrator

Use the VMware vRealize Orchestrator (vRO) server to create workflows for VMware environments that further automate administrative actions and prevent inconsistent configurations. This approach allows for more self-service functions. vRO integration through IBM Spectrum Connect provides the ability to create, extend, map, and delete volumes on the IBM storage systems without any VMware or storage administrator actions.

2.3 Assumptions for the use cases in this paper

The use cases that are described in this paper use the IBM Spectrum Accelerate family and are based on the following versions:

- ▶ IBM Spectrum Connect V3.4
- ▶ VMware 5.5, 6.0, and 6.5
- ▶ IBM XIV Storage System V11.6.2.a
- ▶ IBM Flash System A9000 V12.2.1
- ▶ IBM Flash System A9000R V12.2.1

The use cases that are presented here have these assumptions:

- ▶ The storage area network (SAN) switches that are used in the use cases are already running and the zones are configured.
- ▶ The installation and configuration of any IBM storage system that is used in this paper is not covered. They are already running in the test environment.
- ▶ The VMware environment deployment is not covered. This paper is based on the assumption that the VMware is already installed.

This paper is not intended as any type of formal certification. For more information about hardware capability and supported configurations, see the VMware hardware compatibility list and IBM System Storage Interoperation Center (SSIC) websites:

- ▶ VMware Compatibility Guide:

<http://www.vmware.com/resources/compatibility/search.php>

- ▶ IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STWMS9/landing/IBM_Spectrum_Control_Base_Edition_welcome_page.html

- ▶ SSIC:

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

If a configuration that you want is not available for selection on the SSIC website, submit a Solution for Compliance in a Regulate Environment (SCORE) to IBM, requesting approval. Another name for this procedure is submitting a Request for Price Quotation (RPQ).

Note: To submit a SCORE or RPQ, contact your IBM System Sales Representative.



Attaching VMware ESXi

This chapter describes the considerations and implementation steps that are involved when attaching an IBM Spectrum Accelerate family system to VMware ESXi hosts 5.x and 6.x.

For each version, this chapter describes preferred practices for multipathing and performance tuning.

Note: In this chapter, when referring to information that relates to IBM XIV Storage System, IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM Spectrum Accelerate, the generic *storage system* is used unless specified otherwise.

This chapter covers the following topics:

- ▶ VMware ESXi 5.5 and 6.0
- ▶ VMware ESXi 5.0 and 5.1

3.1 VMware ESXi 5.5 and 6.0

This section describes attaching ESXi 5.5 and 6.0 hosts to the storage system through Fibre Channel (FC) and Internet Small Computer Systems Interface (iSCSI).

3.1.1 ESXi connectivity

You can create ESXi connectivity to the storage systems. This section provides general diagrams that describe an ESXi host that is connected to a storage system by FC and iSCSI. For more information about host connectivity to IBM Spectrum Accelerate family systems, see *IBM FlashSystem A9000*, *IBM FlashSystem A9000R*, and *IBM XIV Storage System Host Attachment and Interoperability*, SG24-8368 and *IBM Spectrum Accelerate Deployment, Usage, and Maintenance*, SG24-8267.

FC connectivity

Zoning is mandatory when connecting an FC host to a storage system. Zoning is configured on the SAN switch, and it isolates and restricts FC traffic to only those host bus adapters (HBAs) within a specific zone. The zones can be created either as hard or soft zones. Each type of zone has its merits, and you must determine the type of zone that is suitable for your environment.

Important: IBM XIV Storage System, IBM FlashSystem A9000, and IBM FlashSystem A9000R do not support direct fibre attachment. At the time of writing, IBM Spectrum Accelerate does not support FC attachment.

The logical FC connectivity from the ESXi host to IBM XIV Storage System is shown in Figure 3-1. The ESXi host should be zoned to multiple separate interfaces for redundancy.




Figure 3-1 ESXi IBM XIV Storage System FC host connectivity

For the IBM FlashSystem A9000 and IBM FlashSystem A9000R systems, there are several configurations for Fibre connectivity to the ESXi host. Each varies in terms of their reliability, degree of flexibility, and performance.

Tip: The preferred overall multipath configuration consists of six paths per LUN.

Because the IBM FlashSystem A9000 system has only three grid controllers, the host should be zoned to all the grid controllers for redundancy, as shown in Figure 3-2.




Figure 3-2 IBM FlashSystem A9000 Fibre Channel redundant configuration

Production environments must always have a redundant (high availability) configuration. Avoid single points of failure. Assign as many HBAs to hosts as needed to support the operating system, application, and overall performance requirements.

The remaining of this section details three typical FC configurations that are supported and offer redundancy. All of these configurations have no single point of failure:

- ▶ If a grid controller fails, each host remains connected to all other grid controllers.
- ▶ If an FC switch fails, each host remains connected to multiple grid controllers.
- ▶ If an HBA fails, each host remains connected to multiple grid controllers.
- ▶ If a host cable fails, each host remains connected to multiple grid controllers.

Figure 3-3 shows an example of a configuration for an IBM FlashSystem A9000R system with six grid controllers. Each ESXi host has dual HBAs with either a single port or dual ports, and is connected to the SAN switch.

- ▶ Each host is equipped with dual HBAs. Each HBA (or HBA port) is connected to one of two FC switches in separate fabrics.
- ▶ Each of the FC switches has a connection to a separate FC port of each of the six grid controllers (in FlashSystem A9000R).
- ▶ Each fabric has 2 zones (noted by the different colors of the connections), with three paths per fabric, per zone for each host, giving a total of six paths per volume. If a fabric fails, all grid controllers that are connected are still used.

If the system had 8 grid controllers such as in a fully configured FlashSystem A9000R Model 425, each of the FC switches would have 2 zones, each zone having a connection to four separate grid controllers.

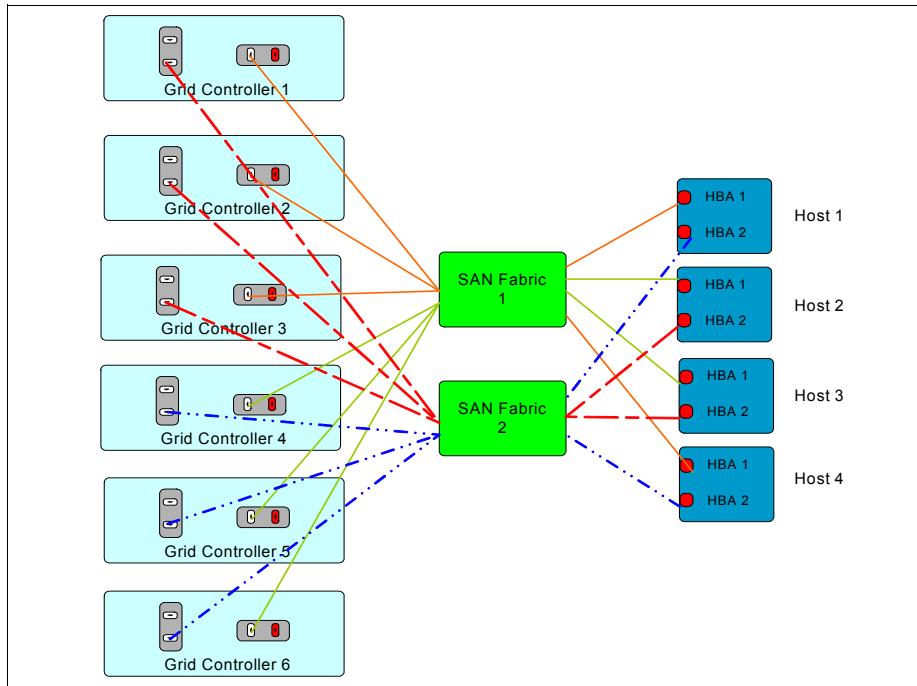


Figure 3-3 IBM FlashSystem A9000R Fibre Channel redundant configuration

Important: The configuration shown in Figure 3-3 is a good overall multipathing configuration consisting of six paths per LUN, for a system with 6 grid controllers.

An even simpler redundant configuration, which still assumes six grid controllers, is illustrated in Figure 3-4.

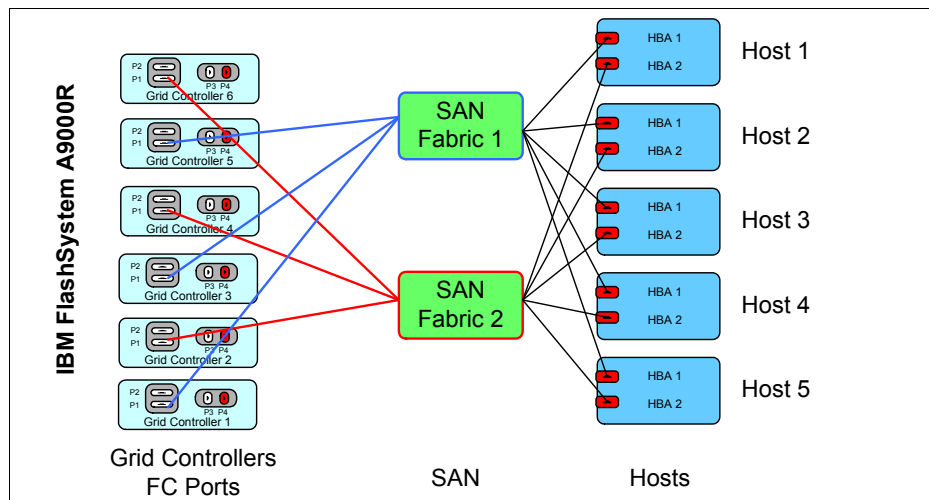


Figure 3-4 Fibre Channel simple redundant configuration

The configuration in Figure 3-4 has the following characteristics:

- ▶ Each host is equipped with dual HBAs. Each HBA (or HBA port) is connected to one of two FC switches in separate fabrics.
- ▶ Each of the FC switches has a connection to three separate grid controllers.
- ▶ Each volume has six paths.

If the system had 8 grid controllers such as in a fully configured FlashSystem A9000R Model 425, each of the FC switches would have a connection to four separate grid controllers.

In a fully configured FlashSystem A9000R Model 415, 12 grid controllers are present. This type of configuration can benefit from attaching each physical connection to each grid controller's port 1 and port 3, as illustrated in Figure 3-5. It is important to equally distribute the workload over all grid controllers.

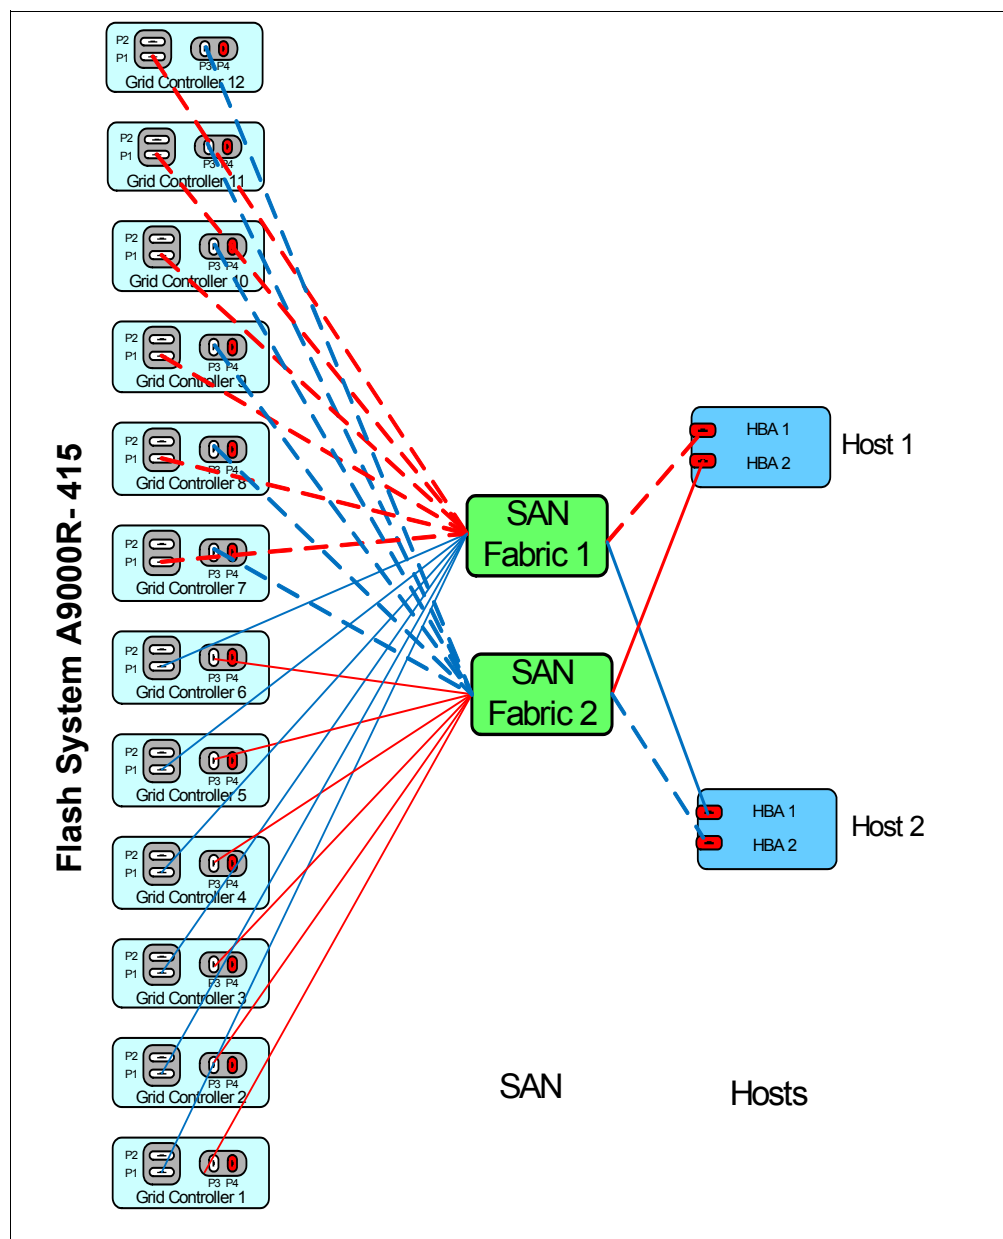


Figure 3-5 IBM FlashSystem A9000R model 415 Full Fibre Channel redundant configuration

The diagram shows the following configuration:

- ▶ Host 1 has HBA1 connected to SAN Fabric 1 and HBA2 connected to SAN Fabric 2
- ▶ Host 2 has HBA1 connected to SAN Fabric 1 and HBA2 connected to SAN Fabric 2

SAN Fabric1 has the following configuration:

- ▶ Six paths to Port1 (FC adapter 1) in grid controllers 7 - 12 (Zone 1)

- ▶ Six paths to Port1 (FC adapter 1) in grid controllers 1 - 6 (Zone 3)

SAN Fabric2 has the following configuration:

- ▶ Six paths to Port3 (FC adapter 2) in grid controllers 1 - 6 (Zone 2)
- ▶ Six paths to Port3 (FC adapter 2) in grid controllers 7 - 12 (Zone4)

Host 1 is zoned to each Grid Controller with 12 paths in total, through these items:

- ▶ HBA 1 zoned to port 1 (FC adapter 1) in grid controllers 7 - 12 through SAN fabric 1
- ▶ HBA 2 zoned to port 3 (FC adapter 2) in grid controllers 1 - 6 through SAN fabric 2

Host 2 is zoned to each Grid Controller with 12 paths in total, through these items:

- ▶ HBA 1 zoned to port 1 (FC adapter 1) in grid controllers 1 - 6 through SAN fabric 1
- ▶ HBA 2 zoned to port 3 (FC adapter 2) in grid controllers 7 - 12 through SAN fabric 2

iSCSI connectivity

iSCSI is an Internet Protocol (IP) based storage system networking standard that provides block-level access to storage systems over a TCP/IP network. IBM XIV Storage System, IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM Spectrum Accelerate systems can communicate to the host through iSCSI. When you implement iSCSI connectivity for the storage systems in a vSphere environment, adhere to the following practices:

- ▶ There is one VMkernel port group per physical network interface card (NIC):
 - The VMkernel port is bound to the physical NIC port in vSwitch, which creates a “path”.
 - Creates a 1-to-1 “path” for VMware NMP.
 - Uses the same Path Selection Plug-in (PSP) as for FC connectivity.
- ▶ Enable jumbo frames for throughput-intensive workloads (must be done at all layers).
- ▶ Use a round-robin PSP to enable load balancing across all modules. Each initiator should see a target port on each module.
- ▶ The queue depth can also be changed on the iSCSI software initiator. If more bandwidth is needed, the LUN queue depth can be modified.

Depending on your storage system configuration, each system has some ports for the use of iSCSI:

- ▶ IBM XIV Storage System

Table 3-1 shows the specification.

Table 3-1 IBM XIV Storage System iSCSI specification

Maximum number of iSCSI over Gigabit Ethernet ports	6 - 22 (1-Gbps Ethernet ports) 12 (10-Gbps Ethernet ports)
iSCSI rates	1 - 10 Gbps

Figure 3-6 on page 31 shows a connectivity example.




Figure 3-6 IBM XIV Storage System i@ SCSI connection

- IBM FlashSystem A9000 and IBM FlashSystem A9000R systems
 - Based on host system interfaces (per grid controller). Table 3-2 shows the specification.

Table 3-2 IBM FlashSystem A9000 and IBM FlashSystem A9000R iSCSI specification

Storage system with FC capabilities.	Four 16 FC + two 10 iSCSI ports
Storage system with iSCSI (Ethernet) capabilities only.	Four 10 Gb iSCSI ports

Figure 3-7 shows a connectivity example.

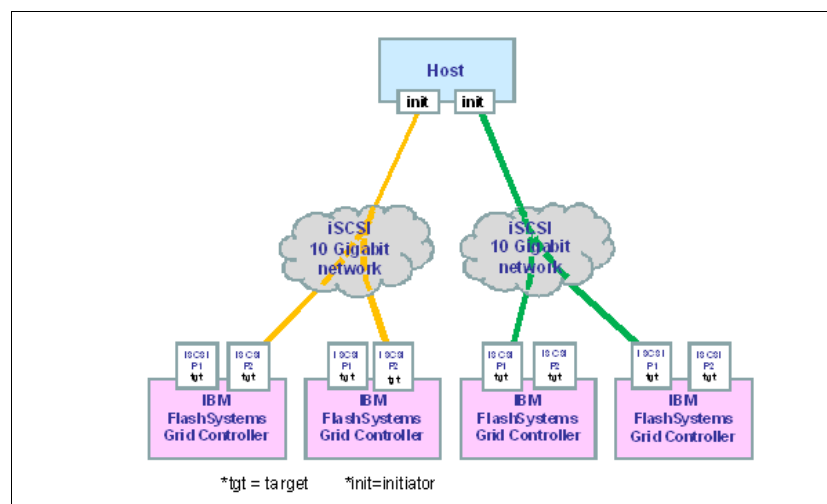


Figure 3-7 IBM FlashSystem A9000 and IBM FlashSystem A9000R iSCSI connection

► IBM Spectrum Accelerate

- The storage system runs only on iSCSI. The number of modules dictates the number of iSCSI ports that are usable. Table 3-3 shows the specification.

Table 3-3 IBM Spectrum Accelerate iSCSI specification

Maximum number of iSCSI over Gigabit Ethernet ports	6 - 30
iSCSI rate	10 Gbps

Figure 3-8 shows a connectivity example.




Figure 3-8 Example of iSCSI connectivity for a four module IBM Spectrum Accelerate system

3.1.2 Installing host bus adapter drivers

VMware ESXi includes drivers for all the HBAs that it supports. VMware strictly controls the driver policy, and only drivers that are provided by VMware can be used. Any driver updates are normally included in service/update packs.

Supported FC HBAs are available from Brocade, Emulex, IBM, and QLogic. Further details about HBAs that are supported by IBM are available from the System Storage Interoperation Center (SSIC) website:

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

Unless otherwise noted in the SSIC, use the firmware and driver versions that are promoted by VMware in association with the relevant hardware vendor. You can find supported VMware driver versions at the following website:

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

3.1.3 Identifying the ESXi host port WWN

Identify the host port WWN for FC adapters that are installed in the ESXi Servers before you start defining the ESXi cluster and its host members by completing the following steps:

1. Run the VMWare vSphere Web Client.
2. Connect to the vCenter Server.
3. In the VMWare vSphere Web Client, click **Hosts and Clusters**, click the host, and click **Manage** → **Storage** → **Storage Adapters**. Figure 3-9 shows the Device name, Type, and WWNs for the installed FC adapters.




Figure 3-9 ESXi host port WWNs

Figure 3-10 shows the Device Name and IQN for the iSCSI adapter for the host.

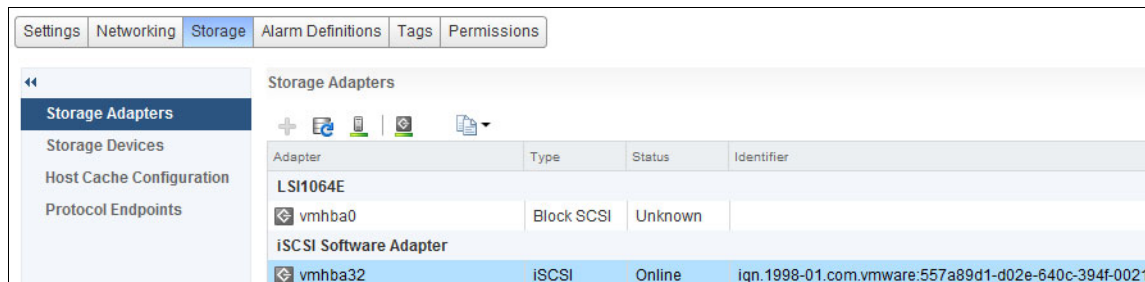


Figure 3-10 ESXi iSCSI Adapter

4. Repeat this process for all ESXi hosts that you are planning to connect to each storage system by using FC and or iSCSI.

After identifying the ESXi host port WWNs, you can define hosts and clusters for the ESXi servers. Create LUNs, and map them to defined ESXi clusters and hosts on the storage system.

Tip: Group the ESXi hosts that access the same LUNs in a cluster on the storage system and assign the LUNs to that cluster.

Considerations for the size and quantity of volumes

For volumes being mapped to an ESXi server, the maximum volume size that you can create on the storage system is 64 TB.

The following configuration maximums for FC are documented for vSphere 5.5 and 6.0:

- ▶ The maximum number of LUNs per server is 256.
- ▶ The maximum number of paths per LUN is 32 for FC attachment.

- ▶ The maximum number of paths per LUN is 8 for iSCSI attachment.
- ▶ The maximum number of paths per server is 1024.

The following configuration maximum for FC is documented for vSphere 6.5:

- ▶ The maximum number of LUNs per server is 512.
- ▶ The maximum number of paths per LUN is 32 for FC attachment.
- ▶ The maximum number of paths per LUN is 8 for iSCSI attachment.
- ▶ The maximum number of paths per server is 2048.

For more information, see the following websites:

- ▶ VMware vSphere 5.5
<https://www.vmware.com/pdf/vsphere5/r55/vsphere-55-configuration-maximums.pdf>
- ▶ VMware vSphere 6.0
<https://www.vmware.com/pdf/vsphere6/r60/vsphere-60-configuration-maximums.pdf>
- ▶ VMware vSphere 6.5
<https://www.vmware.com/pdf/vsphere6/r65/vsphere-65-configuration-maximums.pdf>

3.1.4 Creating a host and volume, and mapping new LUNs

Before scanning for a new LUN on vSphere, you must create a host on HyperScale Manager (GUI) on a selected storage system. Complete the following steps:

1. Log in to the GUI, go to the upper right corner, click **New**, and click **Host**, as shown in Figure 3-11.




Figure 3-11 Create a host

2. To add the new host, complete the appropriate fields and click **Create**, as shown in Figure 3-12:
 - Name: Name of the host.
 - Type: Select either **Default** (all other systems that are not HPUX or ZVM), HPUX, or ZVM.
 - System: The storage system that HyperScale Manager is monitoring.
 - Cluster: Select the cluster if it is created, otherwise leave the field blank.
 - Domains: Select the domain if it is created, otherwise leave the field blank.
 - Ports: Select if a host is connecting either by FC or iSCSI. If FC, find and select **WWPN**; if iSCSI, select **Port Address**.

The screenshot shows a web-based form for adding a host. At the top, there are two input fields: 'Name' with the value 'ITSO_ESXi_1' and 'Type' with a dropdown menu set to 'Default'. Below these is a section header 'SYSTEMS / CLUSTERS / DOMAINS' with a '+' icon. This section contains two sub-sections: 'System (Derived from System Sele...)' with a dropdown menu showing 'A9000', and 'Cluster' with a dropdown menu. Below these is a 'DOMAINS' section with a '+' icon and a 'Domain' dropdown menu showing '/Global Space/'. At the bottom of this section is a 'PORTS' section with a '+' icon. This section contains two identical blocks, each with a radio button for 'FC' (selected) and 'iSCSI', and a 'Port Address' dropdown menu showing '21000024FF28C150'. At the very bottom of the form are two buttons: 'Cancel' and 'Create', with a hand cursor pointing at the 'Create' button.

Figure 3-12 Add a host

3. Click **New** and select **Volumes**, as shown in Figure 3-13.




Figure 3-13 Create a volume

4. Select the **Name**, **Quantity**, **Size**, **System**, and **Pool**, and click **Create**, as shown in Figure 3-14.

Name: ITSO_ESXi_1_Vol

Quantity: 1

Volume Size (GB): 103

System (Derived from System Selector): A9000

Domain: /Global Space/

Pool: ITSO

Buttons: Cancel, Create

Figure 3-14 Add a volume

5. Right-click the volume and select **Mapping** → **View/Modify Mapping**, as shown in Figure 3-15.

1 selected out of 41 V		Volume ITSO_ESXi_1_Vol	
Volume ^	Properties		
	Snapshots		
ITSO_ESXi_1_Vol	Mirror		
	HyperSwap		
ITSO_Mirror	Mapping	View/Modify Mapping	
SVC_SW_1_001	Consistency Group	Unmap All	
SVC_SW_1_002			

Figure 3-15 Modify the mapping

6. Click the + icon to add the mapping, as shown in Figure 3-16.




Figure 3-16 Add a host mapping

7. Select the host and click **Apply**, as shown in Figure 3-17.

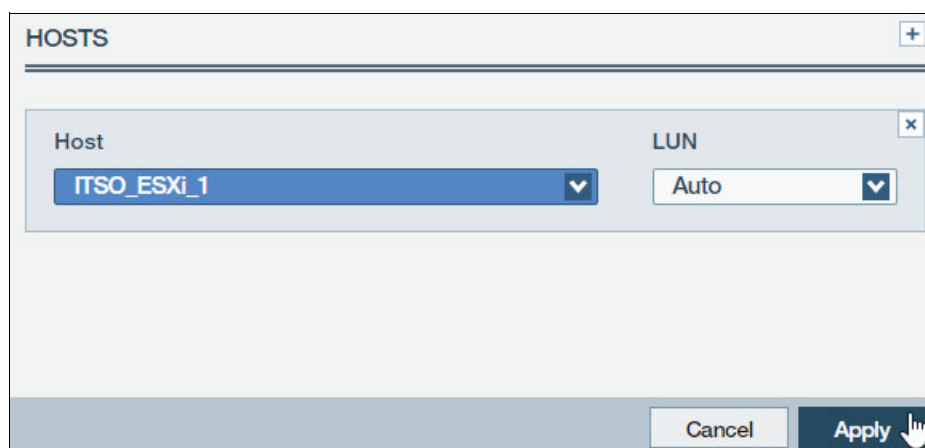


Figure 3-17 Create a host mapping

8. To scan and configure new LUNs in vSphere, complete the following steps:
- Select your ESXi host and click **Manage** → **Storage** → **Storage Adapters**. Click the rescan button, as shown in Figure 3-18.

Here you can see vmhba32 highlighted, but a rescan searches across all adapters. The adapter numbers might be enumerated differently on the different hosts, but this is not an issue.




Figure 3-18 Select FC storage adapters

- Click **OK** to scan for new storage devices, as shown in Figure 3-19.



Figure 3-19 Rescan storage

9. The new LUN is displayed in the Devices pane, as shown in Figure 3-20.

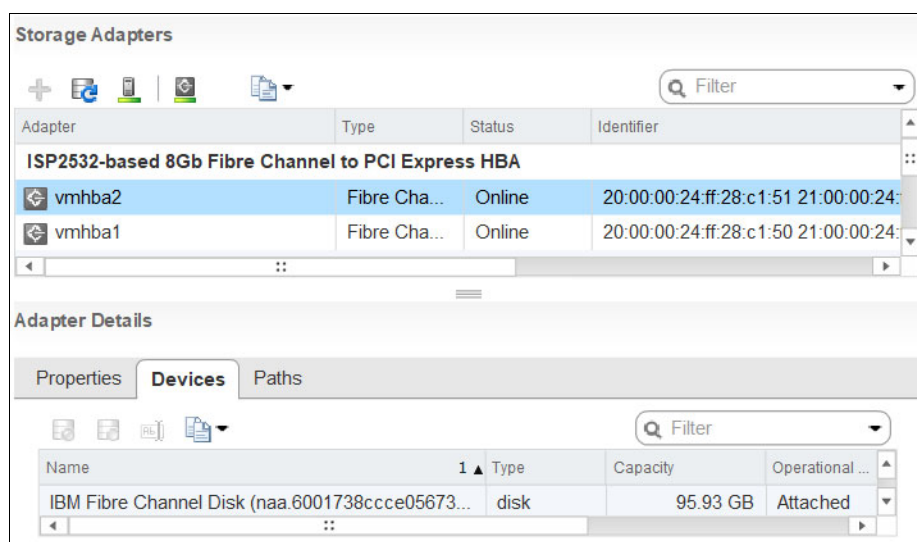


Figure 3-20 Discovered LUN

3.1.5 Attaching an ESXi host

This section describes the attachment of ESXi hosts to the storage system. It provides specific instructions for FC and iSCSI connections. All the information in this section relates to ESXi 5.5, 6.0, and 6.5 (and not other versions of ESXi) unless otherwise specified.

The procedures and instructions that are given here are based on code that was available at the time of writing. For the latest support information, see the SSIC at:

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

By default, ESXi supports the following types of storage arrays:

- **Active/active storage systems:** Systems that allow access to the LUN simultaneously through all storage ports. Although all the paths are active all the time, the access is further defined as either symmetric or asymmetric depending on the architecture of the storage system. In a symmetric storage system, a given LUN can be owned by more than one storage controller at a time, meaning that there is no need to use another storage controller as a proxy to access the LUN. In contrast, an asymmetric storage subsystem designates a particular storage controller as the exclusive LUN owner on a LUN-by-LUN basis, and accessing the LUN through the non-owning controller (by proxy) potentially incurs the additional impact that is associated with involving more than one controller in the I/O path for traditional monolithic storage systems.

The storage systems have a grid architecture that harnesses all resources in tandem, which represents a notable exception to this phenomenon. Finally, a key attribute of the active/active design is that if one or more ports fail, all of the other available ports continue allowing access from servers to the storage system for all LUNs.

- **Active/passive storage systems:** Systems where a LUN is accessible over a single storage port. The other storage ports act as backup for the active storage port.

- Asymmetrical storage systems (VMW_SATP_DEFAULT_ALUA): Systems that support asymmetrical logical unit access (ALUA). ALUA-compliant storage systems provide different levels of access on a per-port basis. This configuration allows the SCSI Initiator port to make intelligent decisions in terms of internal bandwidth usage and processing efficiency on the storage subsystem. The host uses some of the active paths as primary and others as secondary. Accessing a LUN by using a path that exclusively incorporates the managing controller or processor of a traditional monolithic storage subsystem is referred to as an *active-optimized path selection*, whereas accessing a LUN that involves a non-owning controller or processor is referred to as *active-non-optimized*.

VMware introduced the concept of a Pluggable Storage Architecture (PSA) in ESX 3.5. PSA in turn introduced additional concepts to its Native Multipathing Plug-in (NMP).

ESXi provides default Storage Array Type Plug-Ins (SATPs) that support non-specific active-active (VMW_SATP_DEFAULT_AA) and ALUA storage systems (VMW_SATP_DEFAULT_ALUA). Each SATP accommodates special characteristics of a certain class of storage systems. It can perform the storage system-specific operations that are required to detect the path state and activate an inactive path.

ESXi automatically selects the appropriate SATP plug-in for the storage system.

PSPs run with the VMware NMP, and are responsible for choosing a physical path for I/O requests. The VMware NMP assigns a default PSP to each logical device based on the SATP that is associated with the physical paths for that device.

VMware ESXi 5.5, 6.0, and 6.5 support the following PSP types:

- Fixed (VMW_PSP_FIXED): Always use the preferred path to the disk if it is available. If the preferred path is not available, a random alternative path to the disk is chosen. When the preferred path is restored, an automatic failback to the preferred path occurs. This policy is not ALUA-aware, precluding exploitation of the distinction between active-optimized and active non-optimized paths, and so on, in the path selection policy. This type can result in a phenomenon that is known as *path thrashing* when implemented with Active/Passive or asymmetric Active/Active arrays that are based on a traditional monolithic storage subsystem architecture.
- Most Recently Used (VMW_PSP_MRU): This PSP selects the first working path, discovered at boot time, and continues this path (the most recently used) while the path remains available. Whenever a path failure occurs, an alternative path is chosen that uses ALUA-awareness for compatible storage subsystems, meaning that whenever possible paths are chosen to incorporate the managing controller or processor for a given LUN. There is no automatic failback to the original path, and manual intervention is necessary to restore the original state of access after a path failure/repair.
- Round-Robin (VMW_PSP_RR): The Round-Robin policy employs a path selection algorithm that is ALUA-aware and implements load balancing by rotating the physical access to the LUN through all active-optimized paths by default (and also uses active non-optimized paths, if configured to do so). The criteria for migrating to the next available *active* path to a given LUN, and thus to optimize the distribution of workload across paths, relies upon either a path-centric I/O counter or a path-centric byte-wise counter exceeding a pre-set threshold (depending on configuration settings). Paths that are designated as standby, unavailable, or transitioning status are not included in the rotation until they are reactivated or reestablished.

Note: In vSphere 5.5 and later, Round-Robin PSP is now supported for disks that are engaged in Microsoft Cluster Services (MSCS) clustering.

ESXi has built-in rules defining relationships between SATP and PSP for the storage system.

3.1.6 Configuring the ESXi host for multipathing

With ESXi, VMWare supports a round-robin multipathing policy for production environments. The round-robin multipathing policy is always preferred over other choices when attaching to the storage system.

Before proceeding with the multipathing configuration, complete the tasks that are described in the following sections:

- ▶ 3.1.2, “Installing host bus adapter drivers” on page 32
- ▶ 3.1.3, “Identifying the ESXi host port WWN” on page 33
- ▶ “Considerations for the size and quantity of volumes” on page 33

To add a data store, complete the following steps:

1. Start the VMware vSphere Web Client, and connect to your vCenter server.
2. In the vSphere Web Client, click **Home** → **Storage**, as shown in Figure 3-21. Here you can see the data store that is defined for the ESXi host.

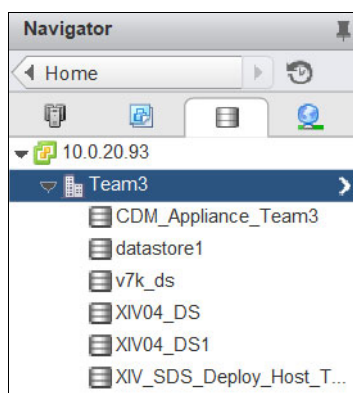


Figure 3-21 ESXi defined data stores

3. Click **Home** → **Hosts and Clusters**, right-click the host, select **Storage**, and click **New Datastore**, as shown in Figure 3-22.

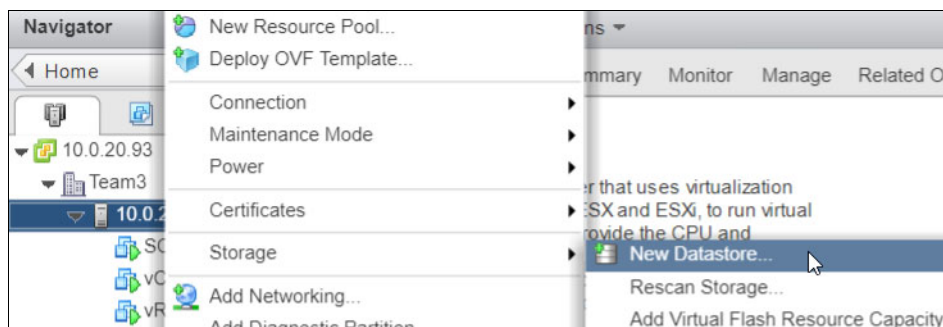


Figure 3-22 ESXi New Datastore

4. Select VMFS as type and click **Next**, as shown in Figure 3-23.

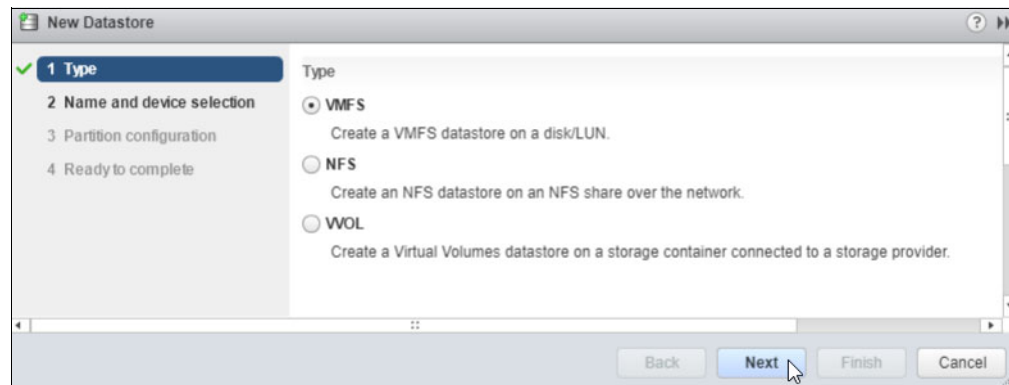


Figure 3-23 ESXi New Datastore: Type

5. Name the data store, select the device, and click **Next**, as shown in Figure 3-24.

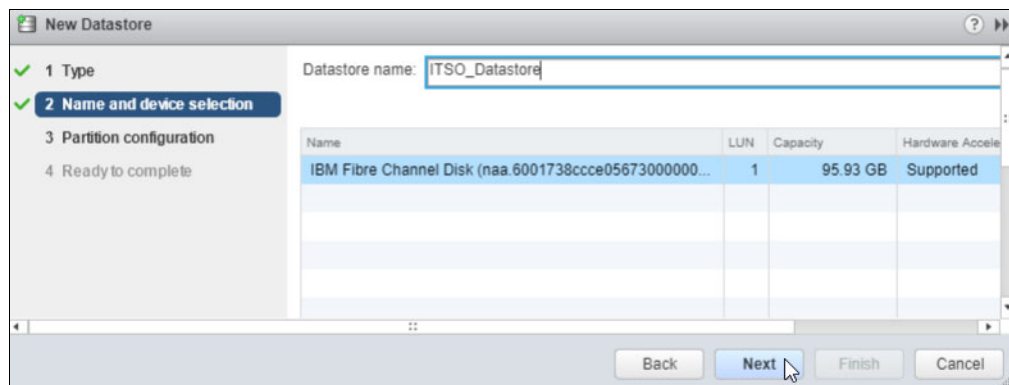


Figure 3-24 ESXi New Datastore: Name and device selection

6. Select the data store size and click **Next**, as shown in Figure 3-25.

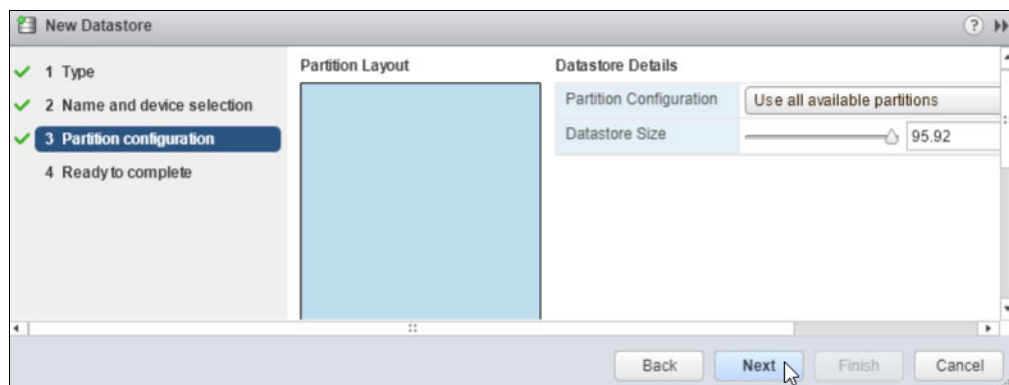


Figure 3-25 ESXi New Datastore: Partition configuration

7. Review the summary and click **Finish**, as shown in Figure 3-26.

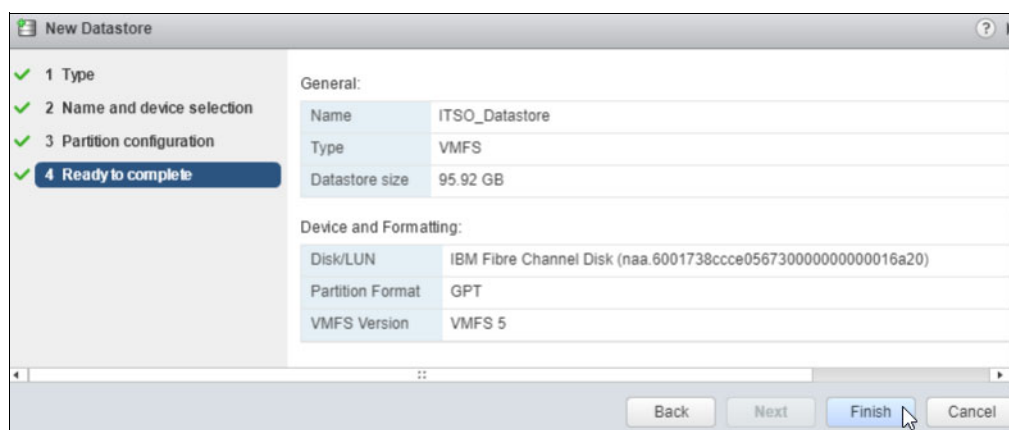


Figure 3-26 ESXi New Datastore: Summary

8. The new data store is displayed, as shown in Figure 3-27.

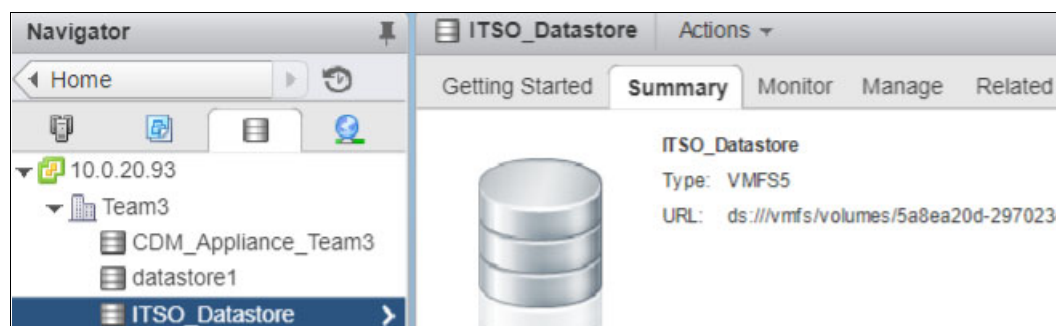


Figure 3-27 ESXi New Datastore created

3.1.7 Performance tuning tips for ESXi hosts

Review settings in ESXi to see whether they affect performance in your environment and with your applications. Settings that you might consider changing are:

- ▶ Using larger LUNs rather than LVM extents.
- ▶ Using a smaller number of large LUNs instead of many small LUNs.
- ▶ Increase the queue size for outstanding I/O on HBA and device levels, only if necessary.
- ▶ Using all available paths for round-robin up to a maximum of 12 paths for FC attachment and up to eight paths for iSCSI attachment. For additional considerations about the number of paths, see 3.1.3, “Identifying the ESXi host port WWN” on page 33.
- ▶ You do not need to manually align your VMFS partitions.

Queue size for outstanding I/O

In general, you do not need to change the HBA queue depths and the device queue depths. Only if you see queue depth used close to 100% for a long period should you increase both queue depths.

To check the device queue depth on VMware ESXi, run **esxtop** and press u. The queue depth is listed under LQLEN. If the %USD (percentage of queue depth used) is constantly close to 100%, increase the device queue depth. Make sure that the HBA queue depths can cope with all device queue depths.

For more information about VMware ESXi queues, see the following website:

<http://blogs.vmware.com/apps/2015/07/queues-queues-queues-2.html>

To change the queue depth on HBA level, see:

https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKnowledgeExternalId=1267

Tip: Commands that use **esxcli** require either the vSphere CLI installed on a management workstation or the Tech Support Mode enabled on the ESXi server. Enabling Tech Support Mode also allows remote SSH shell access. If **esxcli** is run from a command prompt without any form of configuration file, the command uses the following syntax:

```
esxcli --server 9.155.113.135 --username root --password passw0rd <command>
```

If **esxcli** is run from a Tech Support Mode shell or on a host with UNIX utilities, commands such as **grep** and **egrep** can be used. For more information, see the following websites:

- *Using Tech Support Mode in ESXi 4.1, ESXi 5.x, and ESXi 6.x*

<http://kb.vmware.com/kb/1017910>

- *Using ESXi Shell in ESXi 5.x and 6.x*

<http://kb.vmware.com/kb/2004746>

You can also change the queue depth parameters on your HBA by using the tools or utilities that are provided by your HBA vendor.

To change the corresponding parameter on a device level in the VMWare kernel after changing the HBA queue depth, you must use the **esxcli** commands, as shown in the following steps. Since vSphere 6, the parameter *Disk.SchedNumReqOutstanding* is not available in the vSphere Client anymore.

1. Log in to your ESXi host and check the current device queue depth, as shown in Example 3-1.

Example 3-1 List device queue depth

```
# esxcli storage core device list -d eui.001738009c48151b | grep "No of  
outstanding IOs with competing worlds"  
No of outstanding IOs with competing worlds: 32
```

2. Change the device queue depth, as shown in Example 3-2.

Example 3-2 Change device queue depth

```
# esxcli storage core device set -d eui.001738009c48151b -O 64
```

3.2 VMware ESXi 5.0 and 5.1

This section describes attaching ESXi 5.0 and 5.1 hosts.

3.2.1 ESXi 5.0 and 5.1 Fibre Channel configuration

The steps that are required to attach a storage system to vSphere 5.0 and 5.1 servers are similar to the steps in 3.1, “VMware ESXi 5.5 and 6.0” on page 26, but instead of using the vSphere Web Client, you must use the vSphere Client.

To attach a IBM Spectrum Accelerate family system to a vSphere 5.0 or 5.1 server, complete the following steps, which are described in the following example for a XIV Storage System and vSphere 5.0:

1. Identify your ESXi host ports, as shown in 3.1.3, “Identifying the ESXi host port WWN” on page 33.
2. Scan for new LUNs, as shown in “Considerations for the size and quantity of volumes” on page 33.
3. Create your data store, as shown in 3.1.6, “Configuring the ESXi host for multipathing” on page 40. However, when adding a data store, there are three variations from the windows in ESXi 4.1.
4. You are prompted to create either a VMFS-5 or VMFS-3 file system, as shown in Figure 3-28. If you do not use VMFS-5, you cannot create a data store larger than 2 TiB.

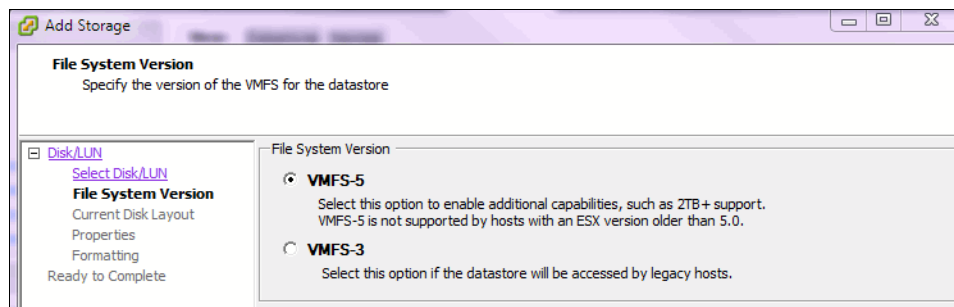


Figure 3-28 The data store file system prompt in vSphere 5.0

5. If you use VMFS-5, you are not prompted to define a maximum block size. You are given the option to use a custom space setting, limiting the size of the data store on the volume. You can expand the data store at a later time to use the remaining space on the volume. However, you cannot use that space for a different data store.

There is no need to manage the paths to the XIV Storage System because round-robin must already be in use by default.

Considerations for the size and quantity of volumes

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 for FC attachment.
- ▶ The maximum number of paths per LUN is eight for iSCSI attachment.

These facts have some important design considerations. If each volume can be accessed through 12 fabric paths (which is a large number of paths), the maximum number of volumes is 85. Dropping the paths to a more reasonable count of six increases the maximum LUN count to 170. For installations with large numbers of raw device mappings, these limits can become a major constraint.

For more information, see the following websites:

- ▶ VMware vSphere 5.0
<http://www.vmware.com/pdf/vsphere5/r50/vsphere-50-configuration-maximums.pdf>
- ▶ VMware vSphere 5.1
<http://www.vmware.com/pdf/vsphere5/r51/vsphere-51-configuration-maximums.pdf>

3.2.2 Performance tuning tips for ESXi 5.0 and 5.1 hosts

Performance tips for ESXi 5.0 and 5.1 hosts are similar to those in 3.1.7, “Performance tuning tips for ESXi hosts” on page 42. However, the syntax of some commands changed, so they are documented here.

Queue size for outstanding I/O

In general, you do not need to change the HBA queue depths and the corresponding `Disk.SchedNumReqOutstanding` VMWare kernel parameter. Only if you see queue depth used close to 100% for a longer period do you need to increase both queue depths.

To check the device queue depth on VMware ESXi, run **esxtop** and press **u**. The queue depth is listed under `LQLEN`. If the `%USD` (percentage of queue depth used) is constantly close to 100%, increase the device queue depth. Make sure that the HBA queue depths can cope with all device queue depths.

For more information about VMware ESXi queues, see the following website:

<http://blogs.vmware.com/apps/2015/07/queues-queues-queues-2.html>

Tip: Commands that use **esxcli** require either the vSphere CLI that is installed on a management workstation or the Tech Support Mode that is enabled on the ESXi server. Enabling Tech Support Mode also allows remote SSH shell access. If **esxcli** is run from a command line without any form of configuration file, the command uses the following syntax:

```
esxcli --server 9.155.113.135 --username root --password passw0rd <command>
```

If **esxcli** is run from a Tech Support Mode shell or on a host with UNIX utilities, commands such as **grep** and **egrep** can be used. For more information, see the following websites:

- ▶ *Using Tech Support Mode in ESXi 4.1, ESXi 5.x, and ESXi 6.x*
<http://kb.vmware.com/kb/1017910>
- ▶ *Using ESXi Shell in ESXi 5.x and 6.x*
<http://kb.vmware.com/kb/2004746>

To set the queue size, complete the following steps:

1. Run the **esxcli system module list** command to determine which HBA type you have (Example 3-3).

Example 3-3 Use the module list command

```
# esxcli system module list | egrep "qla|lpfc"
Name                Is Loaded  Is Enabled
-----
qla2xxx              true       true
or
lpfc820              true       true
```

2. Set the queue depth for the relevant HBA type. In both Example 3-4 and Example 3-5, the queue depth is changed to 64. In Example 3-4, the queue depth is set for an Emulex HBA.

Example 3-4 Set a new value for the queue_depth parameter on the Emulex FC HBA

```
# esxcli system module parameters set -p lpfc0_lun_queue_depth=64 lpfc820
```

In Example 3-5, the queue depth is set for a QLogic HBA.

Example 3-5 Set a new value for the queue_depth parameter on QLogic FC HBA

```
# esxcli system module parameters set -p ql2xmaxqdepth=64 -m qla2xxx
```

3. Restart your ESXi server. After the restart, confirm that the new settings are applied by running the command that is shown in Example 3-6. The example shows only one of many parameters. You must change the syntax if you have an Emulex HBA.

Example 3-6 Check the queue depth setting for a QLogic HBA

```
# esxcli system module parameters list -m qla2xxx | grep qdepth
Name                Type  Value  Description
-----
ql2xmaxqdepth       int   64     Maximum queue depth
```

After changing the HBA queue depth, change the `Disk.SchedNumReqOutstanding` parameter in the VMWare kernel. To change the parameter, complete the following steps:

1. Start the VMWare vSphere Client.
2. Select the server for which you plan to change the settings.
3. Click the **Configuration** tab under the Software section, and click **Advanced Settings** to display the Advanced Settings window.

4. Select **Disk** (circled in green in Figure 3-29) and set the new value for Disk.SchedNumReqOutstanding (circled in red on Figure 3-29).

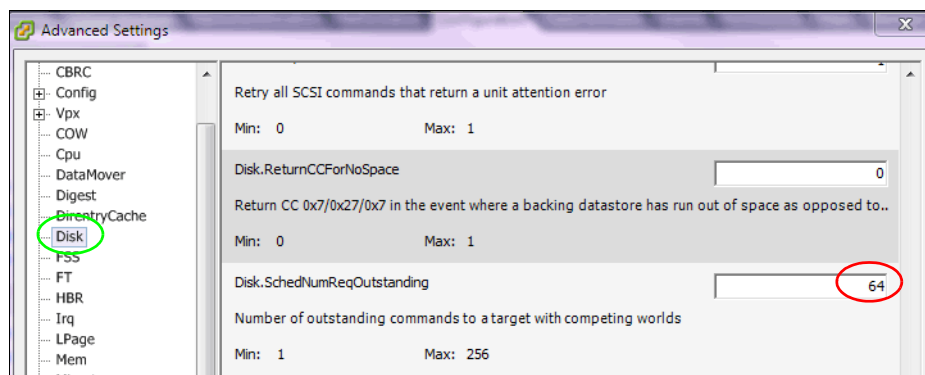


Figure 3-29 Change the Disk.SchedNumReqOutstanding parameter in VMware ESXi 5

5. Click **OK** to save your changes.

Tuning multipathing settings for round-robin

Important: The default ESXi VMware settings for round-robin are adequate for most workloads and normally must not be changed.

If you must change the default settings, enable the non-optimal use for round-robin and decrease the amount of I/O going over each path. This configuration can help the ESXi host use more resources on the XIV Storage System.

If you determine that a change is required, complete the following steps:

1. Start the VMware vSphere Client, and connect to the vCenter server.
2. From the vSphere Client, select your server, click the **Configuration** tab, and select **Storage** in the Hardware section, as shown in Figure 3-30.

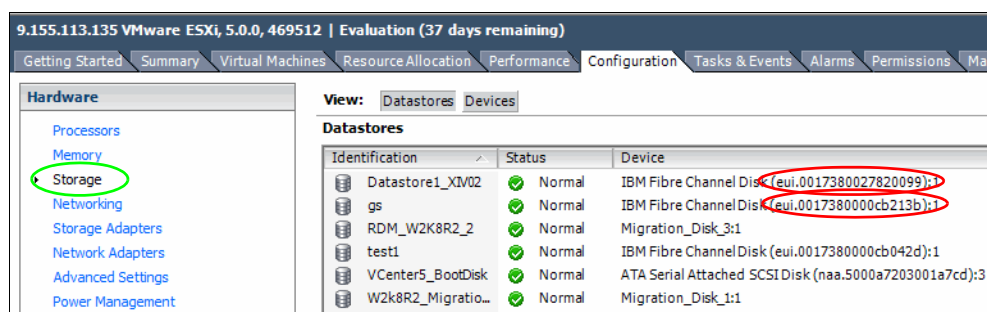


Figure 3-30 Identification of a device identifier for your data store

Here you can view the device identifier for your data store (circled in red). You can also get this information by using **esxcli**, as shown in Example 3-7.

Example 3-7 List storage devices

```
# esxcli storage nmp device list | grep "IBM Fibre Channel Disk (eui.001738"
Device Display Name: IBM Fibre Channel Disk (eui.0017380000cb11a1)
Device Display Name: IBM Fibre Channel Disk (eui.0017380027820099)
Device Display Name: IBM Fibre Channel Disk (eui.00173800278203f4)
```

3. Change the amount of I/O run over each path, as shown in Example 3-8. This example uses a value of 10 for a heavy workload. Leave the default (1000) for normal workloads.

Example 3-8 Change the amount of I/O run over one path for the round-robin algorithm

```
# esxcli storage nmp psp roundrobin deviceconfig set --iops=10 --type "iops"
--device eui.0017380000cb11a1
```

4. Check that your settings are applied, as illustrated in Example 3-9.

Example 3-9 Checking the round-robin options on the datastore

```
#esxcli storage nmp device list --device eui.0017380000cb11a1
eui.0017380000cb11a1
Device Display Name: IBM Fibre Channel Disk (eui.0017380000cb11a1)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: {implicit_support=on;explicit_support=off;
explicit_allow=on;alua_followover=on;{TPG_id=0,TPG_state=A0}}
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config:
{policy=iops,iops=10,bytes=10485760,useA
NO=0,lastPathIndex=1: NumIOsPending=0,numBytesPending=0}
Path Selection Policy Device Custom Config:
Working Paths: vmhba1:C0:T6:L1, vmhba1:C0:T5:L1, vmhba2:C0:T6:L1,
vmhba2:C0:T
5:L1
```

If you need to apply the same settings to multiple data stores, you can also use scripts similar to the ones that are shown in Example 3-10.

Example 3-10 Setting round-robin tweaks for all IBM XIV Storage System devices

```
# script to display round robin settings
for i in `ls /vmfs/devices/disks/ | grep eui.001738*|grep -v \:~ ; \
do echo "*** Current settings for device" $i ; \
esxcli storage nmp device list --device $i
done

# script to change round robin settings
for i in `ls /vmfs/devices/disks/ | grep eui.001738*|grep -v \:~ ; \
do echo "Update settings for device" $i ; \
esxcli storage nmp psp roundrobin deviceconfig set --device $i --iops=1000 --type
"iops";\
done
```

3.2.3 Creating data stores that are larger than 2 TiB

With VMFS-3, the largest possible data store is 2 TiB. With VMFS-5 (introduced in vSphere 5.0), this limit is raised to 64 TiB. Combined with Atomic Test and Set (ATS) and the VAAI primitive that the storage system software levels support, you can use much larger data stores.

ATS locks only the blocks containing the relevant metadata when acquiring the on-disk locks that are necessary to perform metadata updates, rather than implementing SCSI2 reservations to serialize host access with a minimum scope of an entire LUN backing the data store. This procedure improves performance and eliminates the risk of SCSI reservation conflicts.

Do not create a single giant data store rather than multiple smaller ones for the following reasons:

- ▶ Each storage system volume is assigned a SCSI queue depth by ESXi. More volumes mean more SCSI queues, which means more commands can be issued at any one time.
- ▶ The maximum number of concurrent storage vMotions per data store is still limited to eight.

Presenting a volume larger than 64 TiB

If a volume larger than 64 TiB is mapped to an ESXi 5 server, a data store that is formatted with VMFS-5 uses only the first 64 TiB. In Figure 3-31, a 68.36 TiB volume is presented to ESXi 5, but only the first 64 TiB are used.

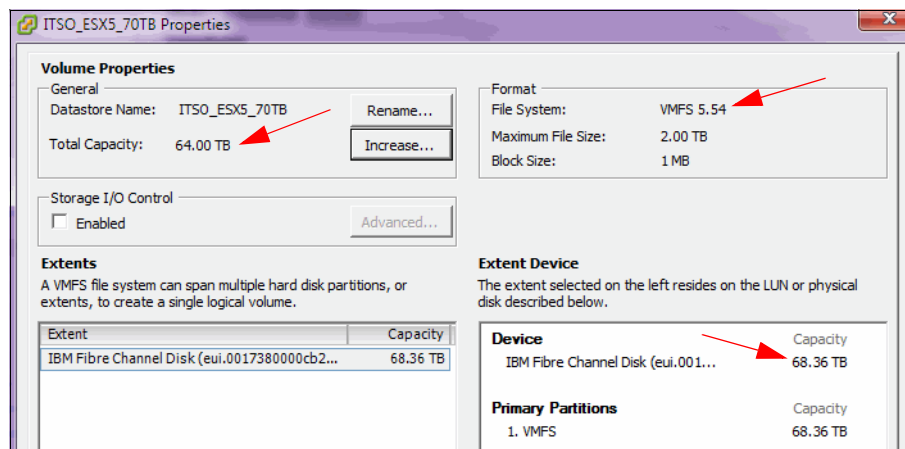


Figure 3-31 vSphere 5.0 volume larger than 64 TiB

If you want to create a data store that approaches the maximum size, limit the maximum storage system volume size as follows:

70364 GB (65531 GiB or 137428467712 Blocks)

Example 3-11 shows the largest possible data store, which is exactly 64 TiB. The **df** command was run on the ESXi server by using the tech support mode shell.

Example 3-11 Largest possible VMFS-5 datastore

```
~ # df
file system      Bytes      Used      Available Use% Mounted on
VMFS-5           44560285696 37578866688      6981419008 84% /vmfs/volumes/Boot
VMFS-5           70368744177664 1361051648 70367383126016 0% /vmfs/volumes/Giant
```



IBM Spectrum Connect software

This chapter introduces the IBM Spectrum Connect Version 3.4 software and covers the following topics:

- ▶ IBM Spectrum Connect overview
- ▶ IBM Spectrum Connect first-time configuration

4.1 IBM Spectrum Connect overview

IBM Spectrum Connect (formerly known as Spectrum Control Base) is a centralized server system that consolidates a range of storage provisioning, automation, and monitoring solutions through a unified server platform. It provides a single-server, back-end location and enables centralized management of storage resources for different virtualization and cloud platforms.

Starting with Spectrum Control Base Version 3.0, IBM Spectrum Connect supports all members of the IBM Spectrum Accelerate family, including IBM FlashSystem A9000 and IBM FlashSystem A9000R systems. Some specific functions, such as virtual volumes (VVoLs), are not supported by IBM FlashSystem A9000, IBM FlashSystem A9000R, or IBM Spectrum Accelerate systems. For more information, see the IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/STWMS9/landing/vmware_compatibility_matrix.html

IBM Spectrum Connect facilitates the integration of IBM Spectrum Accelerate family system resources by using options with supported independent software vendor (ISV) platforms and frameworks. It provides a foundation for integration with IBM systems and ISV solutions.

Figure 4-1 shows a conceptual diagram of the consolidation that IBM Spectrum Connect enables. This is a conceptual diagram, with VMware vSphere as the only currently supported target environment.

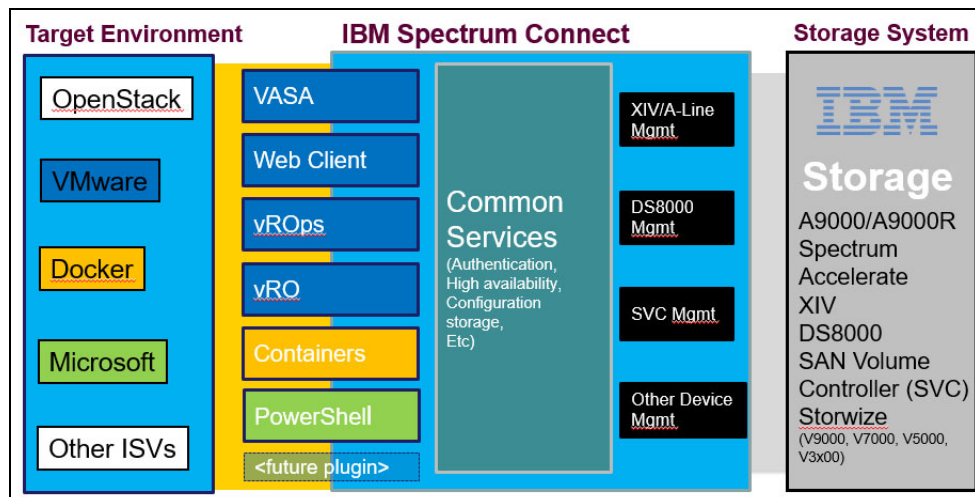


Figure 4-1 IBM Spectrum Connect concept

4.1.1 IBM Spectrum Connect management

IBM Spectrum Connect can be managed through a standard web browser, graphical user interface (GUI), or through a terminal and command-line interface (CLI).

The storage administrator uses IBM Spectrum Connect resources to control preferred IBM storage systems for use in the VMware environment. The administrator also controls the specific vCenter servers that can use the IBM storage resources.

Within the VMware vSphere Web Client, administrators can create, map, and have complete control of storage volumes on the available storage systems and storage pools, as defined by the storage administrator.

In parallel, IBM Spectrum Connect allows registered VMware vCenter servers to connect and use its vSphere APIs for Storage Awareness (VASA), VMware vRealize Operations Manager (vROps), and VMware vRealize Orchestrator (vRO) API functions.

4.1.2 IBM Spectrum Connect advantages

Figure 4-2 shows how IBM Spectrum Connect acts as a middle layer between the VMware environment and the IBM Spectrum Accelerate family, consolidating and reducing the VMware components that generate requests against the storage system.

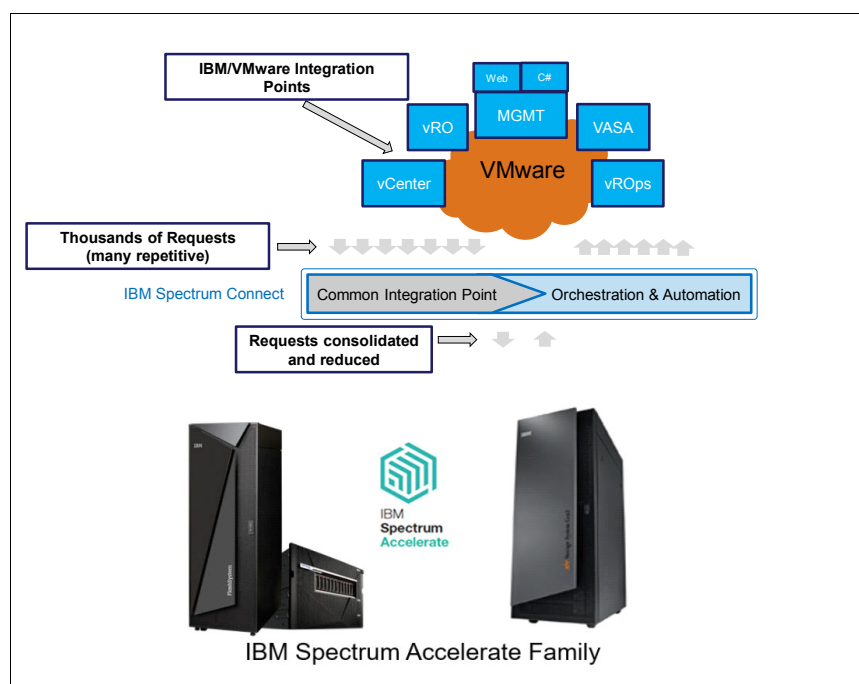


Figure 4-2 IBM Spectrum Accelerate family and VMware Integration with IBM Spectrum Connect

There are many integration points in VMware, including vCenter, vRO, management capabilities in both traditional vSphere Client and vSphere Web Client, VASA, and vROps. Without an integration layer, those various VMware entities generate thousands of requests against the storage system for gathering the information that they need. Moreover, these various entities are requesting the same information again and again from the storage system.

This approach, with many redundant requests against the storage system, does not scale well in large environments with multiple virtual centers. In large environments (several thousands of ESXi hosts), those many requests cannot be handled in a timely fashion. For example, when using the IBM Storage Enhancements for vSphere Web Client in such large environments, the time that is spent to refresh that agent (plug-in) can be 30 minutes easily, or even up to 1 hour for large sites.

With IBM Spectrum Connect, you have a control layer that fetches the information from the storage and cache it. All the repetitive requests from the VMware components can now be served from that middle tier. All of the orchestration and integration is happening through that cached server, and it can provide the requested information to the layer above without hammering the storage. This capability is important for vROps, which send many thousands of commands for information to an IBM storage system, monitoring the health of the system and volumes and how they relate to VMware, and also gathering performance data.

IBM Spectrum Connect runs as a host application under Linux with minimum requirements and a straightforward installation and configuration process. The configuration process is described in 4.2, “IBM Spectrum Connect first-time configuration” on page 54. VMware vCenter Service Appliance (vCSA) 5.5 or later support VASA, so you do not need to use a Windows-based vCenter server.

For more information about compatibility, requirements, extraction, and installation of IBM Spectrum Connect, see IBM Knowledge Center:

<https://www.ibm.com/support/knowledgecenter/en>

4.2 IBM Spectrum Connect first-time configuration

The following steps are necessary for a first-time configuration:

1. Initial Setup Wizard:
 - a. Define the IBM Spectrum Connect fully qualified domain name and high availability group.
 - b. Generate a server certificate.
 - c. Set up storage credentials.
 - d. Change default admin password.
2. Set up the VASA credentials.
3. Add the IBM storage system as a storage array.
4. Add the vCenter server to IBM Spectrum Connect.
5. Control IBM Storage Enhancements for vSphere Web Client plug-in on vCenter.
6. Add the IBM Spectrum Connect server as a storage provider on vCenter.

To start the configuration, complete the following steps:

1. Log in to the IBM Spectrum Connect web interface:
`https://IBM_Spectrum_Connect_IP_address:8440`
2. Enter the default login credentials of user admin and password admin1!, and click **Login**, as shown in Figure 4-3.

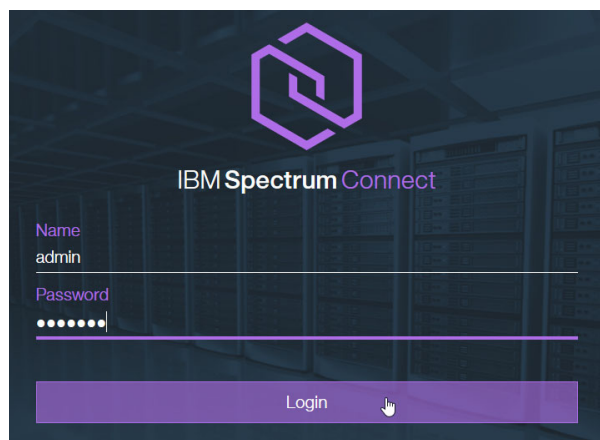


Figure 4-3 Log in to the IBM Spectrum Connect user interface

The IBM Spectrum Connect GUI consists of these four panes:

Interfaces	Integration with vCenter, vRO, PowerShell, and Enabler for Containers
Spaces/Services	Handling storage spaces and services
Storage Systems	Management of storage systems and storage resources
Monitoring	Integration with vROps server

After a successful login, the Spaces/Services and Storage Systems panes are displayed. In the upper left and upper right corners are the Navigation arrows (< >) to navigate to the Applications or Monitoring panes.

The sections that follow guide you through the necessary steps.

4.2.1 Initial Setup Wizard

If the IBM Spectrum Connect server is new, after the first logon an Initial Setup Wizard opens. To perform the setup, complete the following steps:

1. Complete the General settings form with the following parameters and click **Next**, as shown in Figure 4-4:

FDQN	Fully domain qualified name of the IBM Spectrum Connect server
HA GROUP	High availability group containing this IBM Spectrum Connect server

The first defined IBM Spectrum Connect server within a HA group is the active server. The second one is the standby.

For more information about the high availability feature, see the *IBM Spectrum Connect User Guide*, which is provided when you download IBM Spectrum Connect from IBM Fix Central:

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

IBM Spectrum Connect - Initial Setup

☒ General Settings
☐ SSL Certificate
☐ Storage System Credentials
☐ Spectrum Connect Credentials

General Settings

VASA high-availability group ⓘ

FDQN

Back Next

Figure 4-4 IBM Spectrum Connect Initial Setup: General Settings

2. A certificate normally already exists and is valid from the date of your configuration. Regardless, it is required to regenerate the certificate so that the FQDN previously generated is integrated in the certificate. Complete the fields on the form and click **Next**, as shown in Figure 4-5.

The screenshot shows the 'IBM Spectrum Connect - Initial Setup' window. On the left, a sidebar lists four steps: 'General Settings' (checked), 'SSL Certificate' (checked), 'Storage System Credentials' (unchecked), and 'Spectrum Connect Credentials' (unchecked). The main area is titled 'SSL Certificate' with a sub-note: 'The new certificate will be applied during next login.' Below this, there are two radio buttons for 'Method': 'Generate' (selected) and 'Upload'. There are three input fields: 'Common Name' with the value '10.0.20.44', 'Hostname/IP address' with the value '10.0.20.44', and 'Validity' with a value of '3' and the unit 'years'. At the bottom right, there are 'Back' and 'Next' buttons, with a mouse cursor pointing at the 'Next' button.

Figure 4-5 IBM Spectrum Connect Initial Setup: SSL Certificate

3. Set the storage credentials, which must be common to all of the storage devices that are connected to your IBM Spectrum Connect server and are already created on IBM Spectrum Accelerate family systems. Click **Next**, as shown in Figure 4-6.

The screenshot shows the 'IBM Spectrum Connect - Initial Setup' window. On the left, a sidebar lists four steps: 'General Settings' (checked), 'SSL Certificate' (checked), 'Storage System Credentials' (checked), and 'Spectrum Connect Credentials' (unchecked). The main area is titled 'Storage System Credentials' with a sub-note: 'The credentials must be the same for all connected storage systems.' Below this, there are two input fields: 'User name' with the value 'ITSO_SCB' and 'Password' with a masked value '.....'. There is a checkbox labeled 'Directory account' which is currently unchecked. At the bottom right, there are 'Back' and 'Next' buttons, with a mouse cursor pointing at the 'Next' button.

Figure 4-6 IBM Spectrum Connect Initial Setup: Storage System Credentials

4. Change the password for the admin user and click **Finish**, as shown in Figure 4-7.

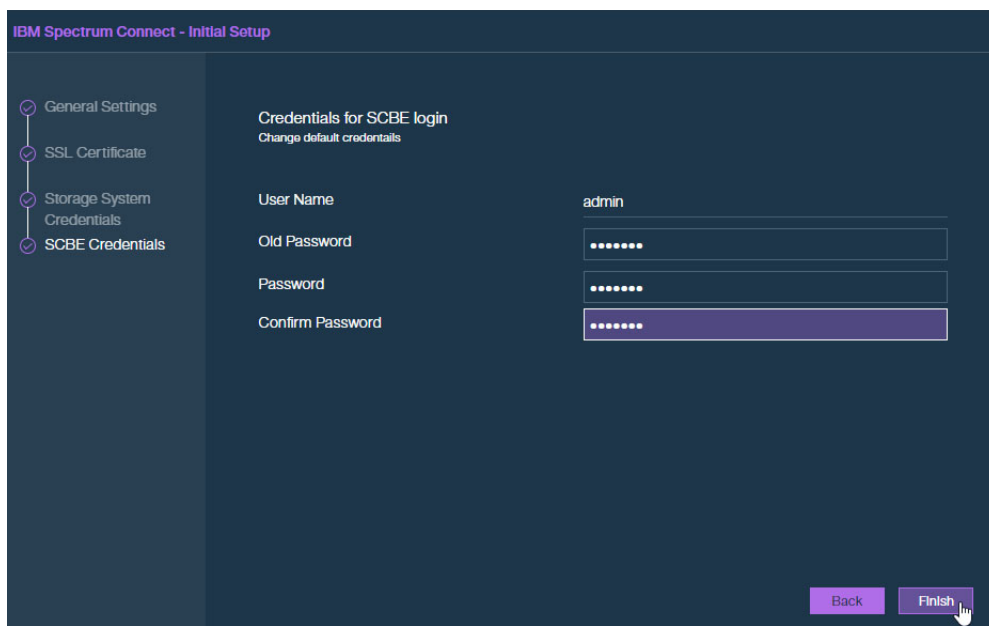
The screenshot shows the 'IBM Spectrum Connect - Initial Setup' window. On the left, a sidebar lists four steps: 'General Settings', 'SSL Certificate', 'Storage System Credentials', and 'SCBE Credentials', each with a checkmark. The main area is titled 'Credentials for SCBE login' with a subtitle 'Change default credentials'. It contains three input fields: 'User Name' (pre-filled with 'admin'), 'Old Password' (masked with dots), and 'Password' (masked with dots). Below the 'Password' field is a 'Confirm Password' field, also masked with dots. At the bottom right, there are two buttons: 'Back' and 'Finish'. A mouse cursor is pointing at the 'Finish' button.

Figure 4-7 BM Spectrum Connect Initial Setup: Change admin password

5. After the wizard finishes and the certificate generates, click **OK** to reload the web page again, as shown in Figure 4-8. Confirm the security exceptions in your browser afterward.

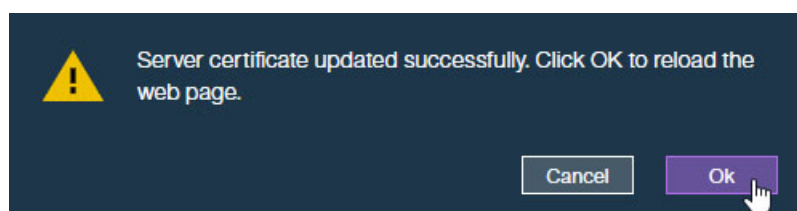


Figure 4-8 Reload the web page

Important information:

- ▶ The storage credentials must be common to all the storage devices that will be connected to your IBM Spectrum Connect server.
- ▶ By using only IBM Storage Enhancements for vSphere Web Client, it is sufficient to run with a Storage Administrator user role and without domains.
- ▶ If VVoLs are required, you must use Storage Integration Administrator user role and domains.

4.2.2 Setting up the VASA credentials

To configure VASA credentials in IBM Spectrum Connect, complete the following steps:

1. Click the **Settings** icon, and then select **VASA Provider credentials** to set up the VASA credentials in IBM Spectrum Connect, as shown in Figure 4-9.




Figure 4-9 Open the VASA credentials form

2. When the VASA credentials form opens, complete the credentials fields and click **Apply**, as shown in Figure 4-10. The credentials are used to register IBM Spectrum Connect as a storage provider at vCenter.

A screenshot of the 'VASA Provider Settings' form. The form has a title bar with a close button. Below the title bar, there is a section titled 'VASA Provider Credentials' with an information icon. It contains three input fields: 'Username' with the value 'vasa', 'Password' with masked characters, and 'Confirm Password' with masked characters. Below this is a section titled 'Storage Systems priority' with an information icon. It contains a 'Priority' dropdown menu with the value '1'. At the bottom right, there are 'Cancel' and 'Apply' buttons. A mouse cursor is pointing at the 'Apply' button.

Figure 4-10 Create VASA credentials

4.2.3 Adding an IBM storage system as a storage array

To add an IBM storage system to IBM Spectrum Connect, complete the following steps:

1. Click the + (plus sign) icon next to the Storage Systems pane to add the storage system to IBM Spectrum Connect, as shown in Figure 4-11.




Figure 4-11 Open the Add New Array form

2. Complete the Add New IBM Storage System form, specifying the IBM storage system Internet Protocol (IP) address or host name, as shown in Figure 4-12. Then, click **Add**.

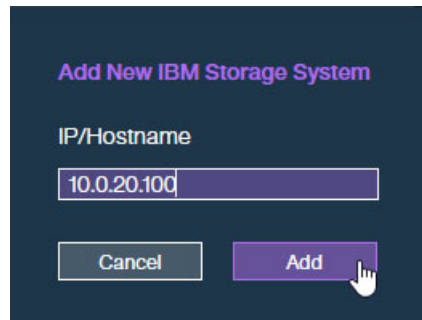


Figure 4-12 Add an IBM storage system as a storage array

3. The storage systems that are added to IBM Spectrum Connect are displayed in the Storage System pane that is shown in Figure 4-13.




Figure 4-13 View of IBM storage systems successfully added to the IBM Spectrum Connect server as a storage array



VMware Virtual Volumes

This chapter describes the tasks that are performed by a storage and VMware administrator to configure virtual volumes (VVoLs).

It covers the following topics:

- ▶ Introduction to VMware vSphere Virtual Volumes
- ▶ Defining Virtual Volumes in XIV
- ▶ Select the appropriate storage when creating a VMware virtual machine

5.1 Introduction to VMware vSphere Virtual Volumes

With the announcement of vSphere 6.0, VMware officially released support for the vSphere Virtual Volumes (VVoLs) architecture. VVoLs allow more efficient operations and control of external storage resources, such as the IBM XIV Storage System running XIV Software Version 11.5.1 or later.

At the time of writing, VVoLs are not supported by IBM FlashSystem A9000 Version 12.2.1, IBM FlashSystem A9000R Version 12.2.1, or IBM Spectrum Accelerate Version 11.5.4. This chapter only applies to IBM XIV Storage System.

There are many benefits for storage administrators and virtualization administrators when adopting VVoLs storage management. Advantages include enhanced efficiencies through automation and thorough integration with VVoLs. An XIV Storage System provides an excellent level of storage abstraction to the virtual machine (VM) through the following features:

- ▶ Easy automation provisioning
- ▶ Policy-compliant service levels
- ▶ Snapshots
- ▶ Cloning
- ▶ Offloading
- ▶ Instant space reclamation
- ▶ Hotspot-free performance predictability
- ▶ Extreme capacity use

This section provides a short overview of the VVoLs architecture implementation in an XIV Storage System. The integration of VVoLs with an XIV Storage System is based on the VMware APIs for Storage Awareness (VASA). IBM support for VASA is part of IBM Spectrum Connect.

This section describes the prerequisites and shows a step-by-step illustration of how to set up an XIV Storage System to use VVoLs.

5.1.1 VMware vSphere Virtual Volumes with IBM XIV

Before the availability of VVoLs, a VM in a VMware environment would be presented a disk in the form of a file that is called a *VMware disk* (VMDK). This file represents a physical disk to the VM and can be accessed by the operating system that is installed on the VM in the same way as a physical volume on a regular server. The VMDK file was then placed onto a file system called VMware File System (*VMFS*), which is hosted by a standard volume (LUN), for example, implemented on external storage system, such as an XIV Storage System.

Although this design has the advantage of simplicity, it also imposes constraints and limitations on the management of the VM data. Indeed, the Storage Administrator and the VMware Administrator must agree about the size and placement of volumes in the storage array before the deployment of VMs. This approach presents scalability and granularity issues and cannot respond to business needs dynamically. It also inhibits using advanced storage system functions such as instant snapshots and replication, and complicates backup solutions.

With the availability of the VVoLs technology, each VM disk can now be mapped to an external storage volume (for example, an XIV volume).

Tip: With VVoLs, the XIV Storage System becomes aware of individual VMDK files. Data operations such as snapshot and replication can be performed directly by the XIV Storage System at the VMDK level rather than the entire VMFS data store.

Figure 5-1 shows how VVoLs change the landscape of storage in a virtualized environment.

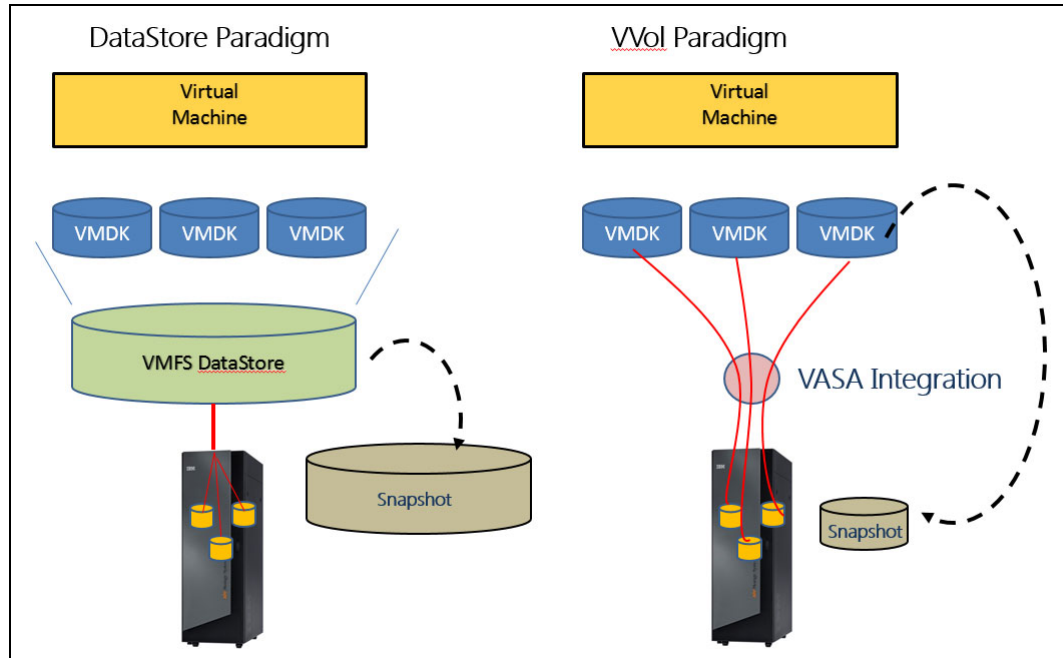


Figure 5-1 VMFS data store and VVoLs paradigms

5.1.2 Implementing VMware vSphere Virtual Volumes on an XIV Storage System

The VVoLs architecture maintains the concept of VMDK files and remains compatible with data storage implementations that are already in place. However, under the VVoLs technology that an XIV Storage System supports, each VMware VM disk can correspond to an XIV volume and can use the storage functions that apply to an XIV volume, such as encryption, snapshot, and replication.

An XIV Storage System uses VASA to present VVoLs to the ESXi host and inform the VMware vCenter of the availability of VVoLs-aware storage, as shown in Figure 5-2.

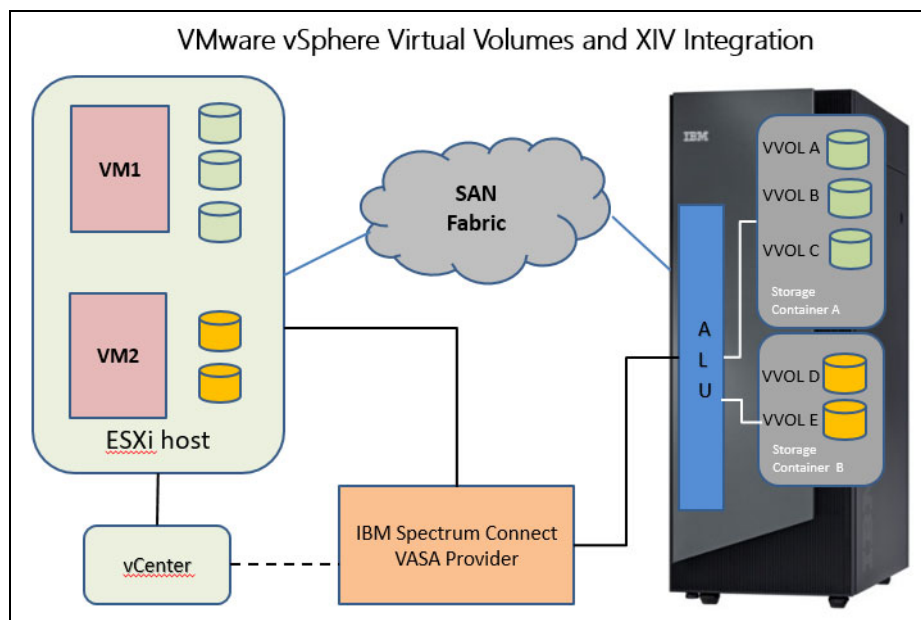


Figure 5-2 VVoLs and XIV integration

Storage containers are configured on VASA by the storage administrator and are used to manage VVoLs and storage resources. Storage containers represent a grouping of VVoLs that are attached to a VM. In the XIV implementation, VASA associates a storage container with a single XIV pool and uses it to present the pool resources to vSphere. The storage containers are characterized by storage services and combine storage capacity with a set of attributes, such as encryption or provisioning type. The storage container is used as a virtual data store to match the requirements of a specific VM and constitutes the basis of a Storage Policy Based Management.

VASA uses the concept of an Administrative Logical Unit (ALU), which is the SCSI object, essentially appearing as a LUN that redirects the SCSI stream to its underlying VVoLs. As such, VVoLs are not mapped directly to a host like regular volumes. Rather, they are bound to a host through the ALU. The ALU is also known as the Protocol Endpoint LUN. The Protocol Endpoint (PE) represents the access point from VM hosts to the storage system, and allows the storage system to carry on storage-related tasks on behalf of the ESXi hypervisor.

To separate the management of regular storage pools in an XIV Storage System from those managed through VASA, they are grouped into separate XIV domains. The VASA provider must be assigned to and control a single domain. Because that domain is not directly managed by the storage administrator, it is marked as an externally managed domain. A new user role, Storage Integration Administrator, is introduced in the XIV Storage System, and is required to perform specific operations on a managed domain.

The VASA implementation that is provided by IBM is packaged with IBM Spectrum Connect (Figure 5-2).

5.1.3 VVoLs concepts mapping in IBM Spectrum Connect

Figure 5-3 shows a detailed mapping of VVoLs concepts in IBM Spectrum Connect and an XIV Storage System:

- ▶ A VVoL maps to an XIV *volume* (or *LUN*).
- ▶ A VMware *storage container* maps to an IBM Spectrum Connect *storage resource*.
- ▶ VMware *VVoLs data store capabilities* map to an IBM Spectrum Connect *VVoL service*.

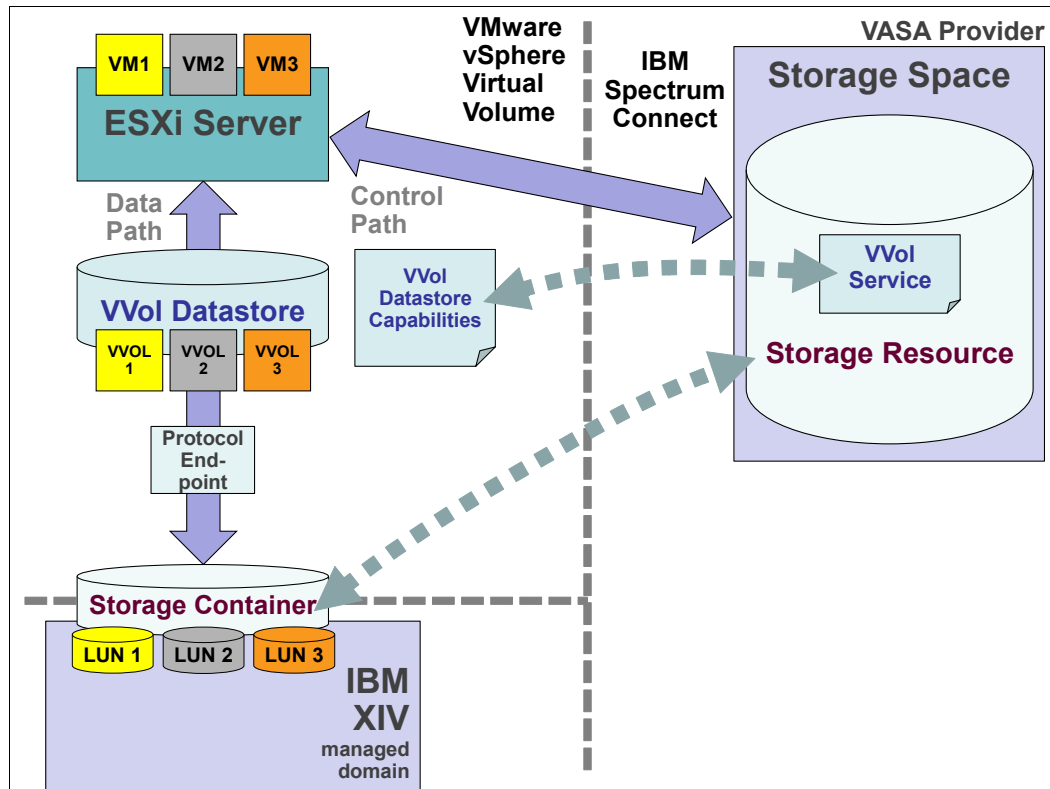


Figure 5-3 VVoLs concepts mapping with IBM Spectrum Connect

5.2 Defining Virtual Volumes in XIV

This section provides a quick overview of the requirements and practical tasks to perform for defining and using VVoLs with an XIV Storage System.

5.2.1 Prerequisites and configuration

The following software and specific versions are required:

- ▶ VMware vCenter 6.0 or later installed.
- ▶ ESXi 6.0 or later installed.
- ▶ IBM XIV Software v11.5.1 or later installed. Consult with your IBM Technical Advisor for details.

- IBM Spectrum Control Base 3.0.1 or later, or Spectrum Connect 3.4 or later installed. The Linux installation package can be downloaded from IBM Fix Central:
<http://www.ibm.com/support/fixcentral/>
Specify **IBM Spectrum Connect** as the product.
- ESXi managed by vCenter.

The following configurations must be completed to connect the VMware infrastructure with an XIV Storage System through the VASA interface. To enable VVoLs for this XIV Storage System, complete these steps:

1. Configure the XIV Storage System to create a managed domain for all VVoLs resources.
2. Configure IBM Spectrum Connect to connect to the XIV storage.
3. Configure VASA on both vCenter and IBM Spectrum Connect.

5.2.2 XIV configuration

The following configuration steps must be completed on the XIV Storage System:

1. Create a domain that includes all XIV elements that are involved in the VVoLs infrastructure. These tasks can be accomplished either through Hyper-Scale Manager (GUI) or by using the XCLI. This example shows these configuration steps with the GUI:
 - a. Create a managed domain to host VVoLs-related XIV elements.
 - b. Define the VMware ESXi hosts in this domain.
2. Create a user and make the created domain an externally managed domain. These required steps must be completed in the IBM XCLI:
 - a. Create a user in this domain with the Storage Integration Administrator role.
 - b. Enable metadata service for the XIV Storage System.
 - c. Create a XIV ALU for each VMware ESXi host.

Creating an XIV domain for all VVoLs components

To create an XIV domain, complete the following steps:

1. From the GUI, log on as the storage administrator.
2. The storage components that are used as VVoLs must be in an XIV managed domain, so an XIV domain must be created:
 - a. From the top navigation pane, click **CREATE NEW** → **Domain**, as shown in Figure 5-4.



Figure 5-4 Create a domain

- b. When the Create Domain wizard window opens, specify the domain hard size and soft size, making sure that the soft size is two times larger than hard size (if you are going to use thin provisioning), and enter a domain name, as shown in Figure 5-5.

The 'Create Domain' wizard window is shown with the 'Capacity' tab selected. The 'Name' field contains 'ITSO_Domain'. The 'System (Derived from System Selector)' dropdown is set to 'XIV_04_1340008'. The 'Physical Size (GB)' and 'Domain Size (GB)' fields both contain '2000'. The 'Max Pools' field contains '3'. The 'Max Volumes and Snapshots' field contains '10'. The 'Max CGs' field contains '3'. The 'Max mirrors and HyperSwaps' field contains '10'. The 'Max Data Migrations' field contains '3'. The 'LDAP Domain ID' field contains 'ITSO_Domain'. The 'Create' button is highlighted with a mouse cursor.

Figure 5-5 Capacity tab when creating a domain

- c. In the Properties tab, specify Pools as at least 3, as shown in Figure 5-5. This minimum is required because whenever a VVoL is created by IBM Spectrum Connect, three pools are created:

Meta pool	Holds VMware VM-related management metadata
Thick pool	For thick provisioning
Thin pool	For thin provisioning

Specify the other parameters according to your needs, and then click **Create**.

3. The VMware ESX hosts featuring VVoLs must be defined in the newly created domain. Complete the following steps for each of these hosts:
 - a. In Top of navigation menu, click **CREATE NEW** → **Host**, as shown in Figure 5-6.




Figure 5-6 Select Host from the CREATE NEW menu

- b. When the Add Host window opens, specify the previously created domain name, as shown in Figure 5-7.

Figure 5-7 Complete the fields in the Add Host wizard

4. The newly created hosts now must be mapped to the corresponding Fibre Channel (FC) ports. For this purpose, complete the following steps for each ESX host:
 - a. In the bottom of the pane, click in **Port Address** and select the appropriate FC, as shown in Figure 5-8.

Figure 5-8 Add a port to a host

- b. Select the ESXi host port name in the drop-down list, as shown in Figure 5-9. Repeat this step for each port, and then click **Create**.

USE THIS AS ADDRESS	
WWPN	Used
21000024FF28C151	Yes
21000024FF28C151	

Figure 5-9 Add a port to a host

Tip: To retrieve the ESX port names, run the following command from an SSH session to your host:

```
[root@localhost:~] esxcli storage san fc list|grep "Port Name"
Port Name: 21:00:00:24:ff:28:c1:50
Port Name: 21:00:00:24:ff:28:c1:51
```

You can also retrieve the ESX FC ports from the vSphere Web Client.

- c. Your added ports then appear in the Host configuration in the right pane of the selected host, as shown in Figure 5-10.

PORTS	
<input checked="" type="radio"/> FC <input type="radio"/> iSCSI	Port Address 21000024FF28C150
<input checked="" type="radio"/> FC <input type="radio"/> iSCSI	Port Address 21000024FF28C151

Figure 5-10 Fibre Channel port that is defined on a host

Creating a user and a protocol endpoint

To create the user and PE, complete the following steps:

1. A user with the Storage Integration Administrator role must be created and associated with this domain. As the Storage Integration Administrator role is not available in the GUI, you must use the XCLI. Complete the following steps:
 - a. Log on to your XCLI as user admin.
 - b. Create a user with the Storage Integration Administrator role, as shown in Example 5-1.

Example 5-1 Create Storage Integration Administrator

```
XIV_04_1340008>>user_define user=ITS0_SCB category=storageintegrationadmin  
domain=ITS0_Domain password=Test1234a password_verify=Test1234a  
Command executed successfully.
```

2. Enable the metadata service by running the command that is shown in Example 5-2.

Example 5-2 Enable the metadata service

```
XIV_04_1340008>>metadata_service_enable  
Command executed successfully.
```

3. Log on to the XCLI with the user ID (Storage Integration Administrator role) that you created and create an ALU for each VMware ESXi host by running the command that is shown in Example 5-3 with the following parameters:

alu	A name for this ALU
host	The ESXi host as defined in the XIV Storage System
lun	Any unique number 512 - 755

Example 5-3 Create an administrative logical unit

```
XIV_04_1340008>>alu_create alu=ITS0_VVOL_ALU host=ITS0_ESXi3 lun=603  
Command executed successfully.
```

IBM Spectrum Connect configuration

After the XIV Storage System is configured, the IBM Spectrum Connect must be made aware of it. Complete these high-level steps to accomplish this task:

1. Complete the first-time IBM Spectrum Connect configuration, which is described in Chapter 4, “IBM Spectrum Connect software” on page 51:
 - a. Define the IBM Spectrum Connect fully qualified domain name and high availability group.
 - b. Generate a server certificate.
 - c. Set up VASA credentials.
 - d. Set up XIV credentials.
2. Add the XIV Storage System.
3. Create a storage space for this XIV Storage System in IBM Spectrum Connect.
4. Add a VVOL-enabled service for this storage space.
5. Define a storage resource for this VVOLs-enabled service.
6. Delegate the service to the vCenter.

Detailed description of the necessary steps

Complete the following steps:

1. Log in to IBM Spectrum Connect Web Interface:
`http://IBM_Spectrum_Connect_IP_address:8440`
2. Enter the login credentials, as shown in Figure 5-11.

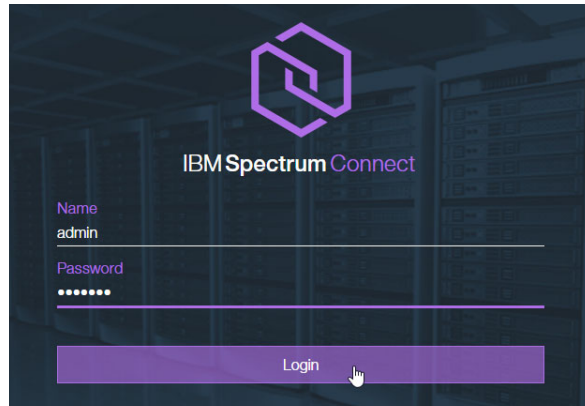


Figure 5-11 Log in to the IBM Spectrum Connect user interface

3. If needed, complete the first-time configuration steps that are described in 4.2, “IBM Spectrum Connect first-time configuration” on page 54. Use a Storage Integration Administrator user ID and password as storage credentials.
4. Complete the following steps to add your XIV Storage System to IBM Spectrum Connect:
 - a. Click the + (plus sign) icon next to the Storage Systems pane to add your XIV Storage System to IBM Spectrum Connect, as shown in Figure 5-12.



Figure 5-12 Add a storage system

- b. Complete the Add New IBM Storage System form with your XIV Internet Protocol (IP) address or host name, as shown in Figure 5-13.




Figure 5-13 Add your XIV Storage System

- c. After the XIV address is given, the XIV Storage System is displayed as shown in Figure 5-14.

The Free size that is shown in Figure 5-14 corresponds to the size that was defined in step b on page 67.




Figure 5-14 Display your XIV Storage System

- d. **Optional:** You can click the icon and select **Modify**, as shown in Figure 5-15, to display the XIV properties.




Figure 5-15 Open XIV properties

The XIV properties window opens, as shown in Figure 5-16. Notice that no service is defined yet.

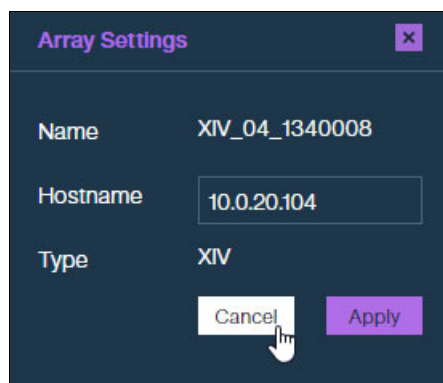


Figure 5-16 Show XIV properties

Defining a storage space and configuring a storage service

Complete the following steps:

1. Within IBM Spectrum Connect server, virtual storage is defined with a *storage service* and a *storage space*. Therefore, you must first define a storage space. For that purpose, complete the following steps:
 - a. Click the **Configuration** icon, and then click **Add New Storage Space** to add a storage space to your IBM Spectrum Connect server, as shown in Figure 5-17.




Figure 5-17 Call the Add New Space form

- b. Provide a name for your new storage space and click **Apply**, as shown in Figure 5-18.

A screenshot of the 'New Storage Space' form in the IBM Spectrum Connect web interface. The form has a title bar with 'New Storage Space', an information icon, and a close button. It contains two input fields: 'Storage Space Name' with the value 'ITSO_VVOL_Space' and 'Description' with the value 'Space for VVOL'. At the bottom right, there are two buttons: 'Cancel' and 'Apply'. A mouse cursor is clicking the 'Apply' button.

Figure 5-18 Add a storage space

IBM Spectrum Connect GUI automatically brings you to this newly created storage space.

2. Now that storage space is defined, a storage service must be configured. A storage service is the combination of storage resources and associated user-defined policies, such as encryption and mirroring. To add a VVoL-enabled service to the newly created storage space, complete the following steps:
 - a. From your newly created storage space, click the **+** icon next to the Services pane, as shown in Figure 5-19.




Figure 5-19 Call the Add New Service form

- b. When the New Service form is displayed, complete it as shown in Figure 5-20. Specify the features that are fulfilled by this service according to your needs. Do not forget to select **VVOLs Service**.

ITSO_VVOL_Space ▼ Storage Services +

New Storage Service ⓘ

Name ITSO_VVOL_Service

Description Service for VVOL

☒ VVol Service ⓘ

Related VVol Containers:

Capabilities ⓘ

Space Efficiency ▼ Thin ▼ x

Select Capability ▼

Cancel Create

Figure 5-20 Add a VVOL-enabled service in the New Service form

Use the **Select Capability** drop-down box to define the new service attributes and their values.

The following options are available in IBM Spectrum Connect Version 3.1:

- Encryption: Enables encryption for the service. If enabled, you can attach only encrypted storage resource to the service.
- Space Efficiency: Enables storage space efficiency features for the service. When selected, you can configure the service to be attached to a thick- or thin-provisioned storage resource.
- Flash: Enables utilization of a storage resource on a flash-based storage resource. This can be one of the following storage systems: IBM FlashSystem 9000, IBM FlashSystem V9000, and the IBM Storwize® family of systems.
- QoS: Enables the use of the Quality of Service (QoS) feature for the service. QoS is applicable to volumes (Max Independent Performance) or storage resources (Max Shared Performance). Set the IOPS and bandwidth limits within the following ranges:
 - IOPS: 0 -100000
 - BW (bandwidth): 0 -10000 MBps

Restriction: The QoS capability is not available for the Spectrum Accelerate Family of products, in the context of the VVol service creation window shown in Figure 5-20.

- Availability: Defines the resiliency that is expected for an application. It can be stretched or regular. It applies to SAN Volume Controller.
 - Data Reduction: Enables the use of IBM Real-time Compression™.
 - Replication: Enables synchronous mirroring for the service. Not available for VVOL-enabled services. It applies as IBM HyperSwap® to Spectrum Virtualize family and as synchronous replication to Spectrum Accelerate family.
- c. The service now appears in your newly created space, as shown in Figure 5-21.




Figure 5-21 New VVOLs service added

3. To add a storage resource for this newly created VVOL-enabled service, complete the following steps:
- Right-click the service and select **Manage Resources**, as shown in Figure 5-22.




Figure 5-22 Service Manage Resources

- Click the + button on the Storage System (IBM XIV Storage System attached), as shown in Figure 5-23.

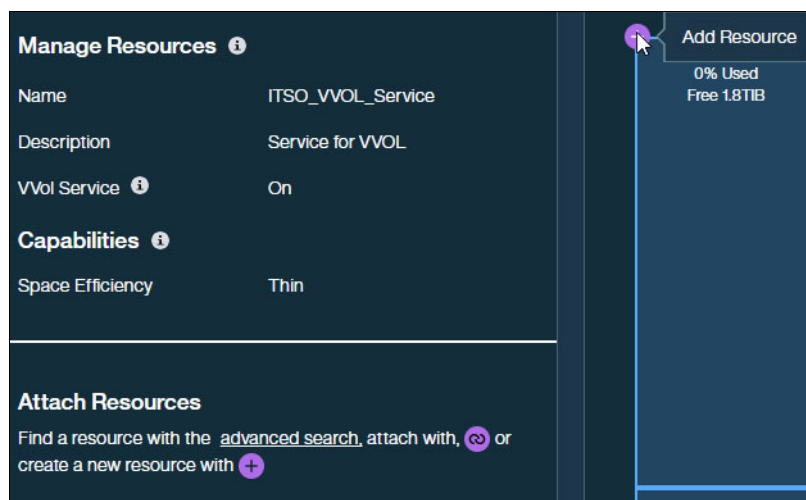
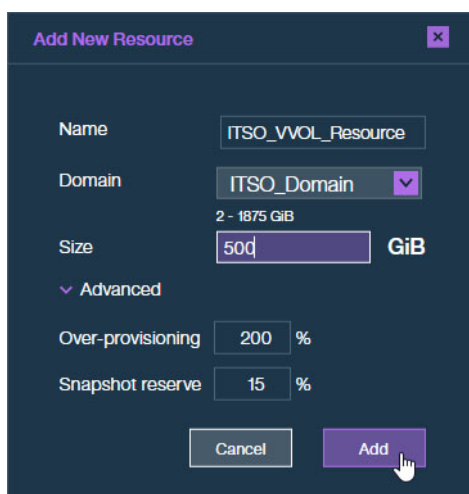


Figure 5-23 Service Add Resource

- c. Enter the appropriate details to add a storage resource and click **Add**, as shown in Figure 5-24.

Note: Wait a few seconds for the VVOLs service to be activated on the storage array.

A dark-themed dialog box titled "Add New Resource" with a close button (X) in the top right corner. It contains several input fields: "Name" with the value "ITSO_VVOL_Resource", "Domain" with a dropdown menu showing "ITSO_Domain", and "Size" with a text input "500" and a unit selector "GiB". Below these is an "Advanced" section with a downward arrow. Inside the "Advanced" section, there are two rows: "Over-provisioning" with a value of "200" and a percentage sign, and "Snapshot reserve" with a value of "15" and a percentage sign. At the bottom are two buttons: "Cancel" and "Add". A mouse cursor is pointing at the "Add" button.

Add New Resource

Name ITSO_VVOL_Resource

Domain ITSO_Domain

Size 500 GiB

Advanced

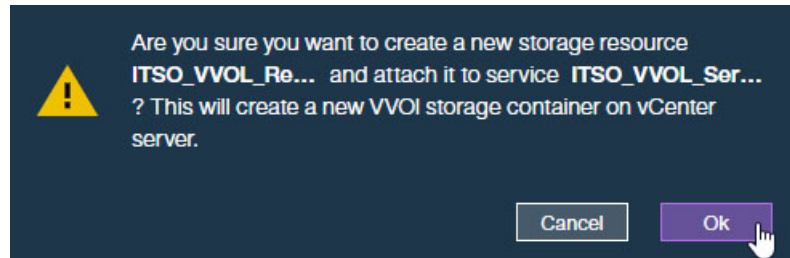
Over-provisioning 200 %

Snapshot reserve 15 %

Cancel Add

Figure 5-24 Add a resource

- d. Click **Ok** to confirm, as shown in Figure 5-25.

A dark-themed warning dialog box with a yellow triangle icon containing an exclamation mark on the left. The text inside reads: "Are you sure you want to create a new storage resource ITSO_VVOL_Re... and attach it to service ITSO_VVOL_Ser... ? This will create a new VVOL storage container on vCenter server." At the bottom right are two buttons: "Cancel" and "Ok". A mouse cursor is pointing at the "Ok" button.

Are you sure you want to create a new storage resource ITSO_VVOL_Re... and attach it to service ITSO_VVOL_Ser... ? This will create a new VVOL storage container on vCenter server.

Cancel Ok

Figure 5-25 Confirm warning

The created storage resource now appears in the Storage System, as shown in Figure 5-26.



Figure 5-26 Created resource

The associated service also shows that it is now used, as shown in Figure 5-27.




Figure 5-27 Service used

4. In the GUI, as admin user, you can now see the associated pools for this storage domain by completing the following steps:
 - a. Click the Domain and select **Navigate to Object(s)**, as shown in Figure 5-28.

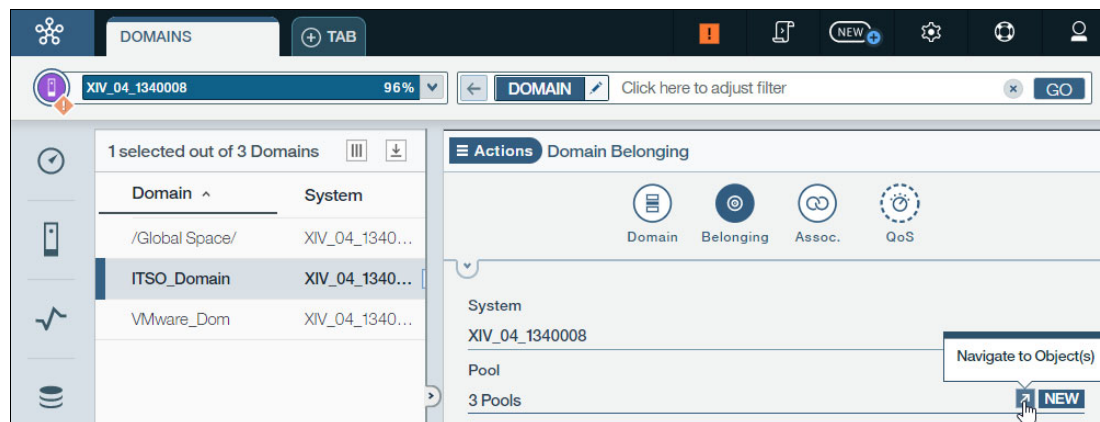


Figure 5-28 Open the Storage Pools window

- b. Three pools were created for this storage resource, pertaining to the same pool group, as shown in Figure 5-29.

Pool ^	System	Domain	Pool Size
ITSO_VVOL_...	XIV_04_1340...	ITSO_Domain	0 GB
ITSO_VVOL_...	XIV_04_1340...	ITSO_Domain	52 GB
ITSO_VVOL_...	XIV_04_1340...	ITSO_Domain	1,136 GB

Figure 5-29 Pools that are associated to created storage domain

These three pools are within the same pool group:

Meta Pool	Holds VMware VM-related management metadata
Thick Pool	For thick provisioning
Thin Pool	For thin provisioning

Important: Upon creation of VMware VVoLs, metadata pools might show 100% utilization in the GUI. Metadata pools look 100% used because every time a new space is added it is fully allocated (both hard and soft space) by the XIV Storage System.

The first time that you define a VM by using VVoLs, a 68 GB volume is allocated from the soft size, and 17 GB is allocated from the hard size. This space is used to serve the newly added VM and three more to come. On the addition of the fifth VM, another 68 GB of soft and 17 GB of physical space is allocated automatically, and so on.

5. To add a vCenter server in IBM Spectrum Connect, complete the steps that are described in 6.2.1, “Adding the vCenter server in IBM Spectrum Connect” on page 89. These steps are only necessary if monitoring in IBM Storage Enhancements for vSphere is needed for the VVoL spaces and services. Click the left-arrow navigation pointer, as shown in Figure 5-30.




Figure 5-30 Go to Interfaces

- a. Click the vCenter and select **Delegate to 9.155.117.25**, as shown in Figure 5-31.



Figure 5-31 Delegate a service to vCenter

- b. A warning opens. Click **Ok** to continue, as illustrated in Figure 5-32.

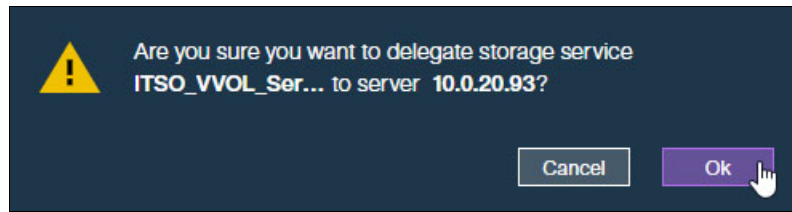


Figure 5-32 Delegate a service to vCenter warning

- c. The service is now delegated to the vCenter, as shown in Figure 5-33.




Figure 5-33 Service delegated to vCenter

5.2.3 VMware vCenter and IBM Spectrum Connect configuration

After both the XIV Storage System and IBM Spectrum Connect are configured, IBM Spectrum Connect can make vCenter aware of the XIV Storage System. To do so, complete the following high-level steps:

1. Register IBM Spectrum Connect as a VASA storage provider in your VMware vCenter.
2. Attach your IBM Spectrum Connect configured services to your vCenter.
3. After all of the above steps are complete, the storage resources that are configured by IBM Spectrum Connect in step 2 on page 73 are seen by VMware vCenter as *storage containers*. Therefore, you can now configure VVoLs data stores in vCenter.
4. You can then define VVoLs data stores with associated VMware storage policies.

These steps are detailed in the next subsections.

Registering IBM Spectrum Connect as a VASA storage provider

To register IBM Spectrum Connect as a storage provider in vCenter, complete the following steps, these steps are only necessary for VVols:

1. Click **Global Inventory Lists** → **vCenter Servers**, and then click the vCenter server and select the Configure tab and the Storage Providers tab. The storage provider window opens, as shown in Figure 5-34.



Figure 5-34 Display the storage provider vCenter window

2. Click the + icon. The New Storage Provider window in Figure 5-35 opens. Complete the fields with the following information:

Name	Any name describing your IBM Spectrum Connect server
URL	<code>https://IBM_Spectrum_Connect_IP_address:8440/services/vasa</code>
User name	The VASA user name, as defined in 4.2.2, “Setting up the VASA credentials” on page 58
Password	The VASA password, as defined in 4.2.2, “Setting up the VASA credentials” on page 58

A screenshot of the 'New Storage Provider' dialog box in the vSphere Web Client. The dialog has a title bar with the IP address '10.0.20.93' and a question mark icon. It contains several input fields: 'Name' with the value 'SC', 'URL' with the value 'https://10.0.20.44:8440/services/vasa', 'User name' with the value 'vasa', and 'Password' with masked characters '*****'. Below these fields is a checkbox labeled 'Use storage provider certificate' which is currently unchecked. To the right of this checkbox is a 'Certificate location' field and a 'Browse...' button. At the bottom right of the dialog are 'OK' and 'Cancel' buttons.

Figure 5-35 Add IBM Spectrum Connect as a storage provider in the vSphere Web Client

- Next, vCenter prompts you to verify the IBM Spectrum Connect security certificate, as shown in Figure 5-36. Click **Yes** to proceed.

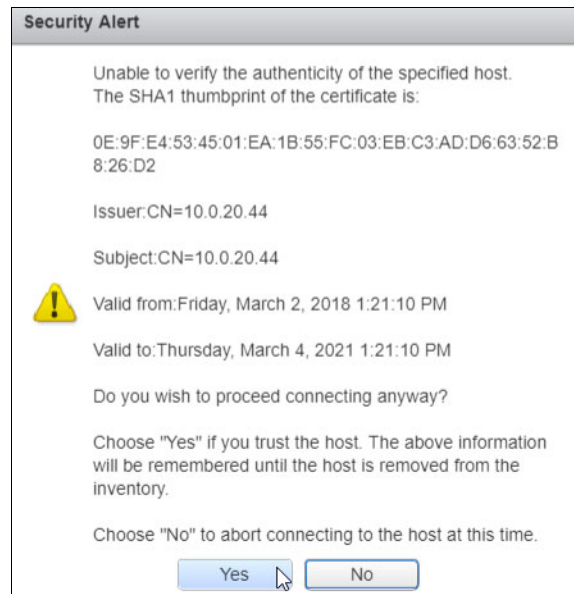


Figure 5-36 Configure IBM Spectrum Connect as new storage provider in vSphere Web Client

- The IBM Spectrum Connect storage provider is displayed in vSphere Web Client, as shown in Figure 5-37.

Storage Providers						
<div> + ✖ 📄 🔍 🔄 </div> <div>Group by: Storage provider</div>						
Storage Provider/Storage System	Status	Active/Standby	Priority	URL	Last Rescan Time	VASA API Version
<div> <div>▼</div> <div>IOFILTER Provider 10.0.20.23</div> </div>	Online	--	--	https://10.0.20.23:908...	3/1/2018 7:0...	1.5
<div> <div>5582c4fb-61a0-06dc-a93e-...</div> </div>		Active	1			
<div> <div>▼</div> <div>SC</div> </div>	Online	--	--	https://10.0.20.44:844...	3/1/2018 6:4...	2.0
<div> <div>XIV_04_1340008 (1/1 online)</div> </div>		Active	1			

Figure 5-37 Registered IBM Spectrum Connect server

- You can verify that IBM Spectrum Connect indeed accepted vCenter registry by going to the IBM Spectrum Connect GUI settings and clicking **VASA trusted certificates**, as shown in Figure 5-38.




Figure 5-38 Calling VASA trusted certificate window

- The window shows the VASA trusted certificate, as shown in Figure 5-39.

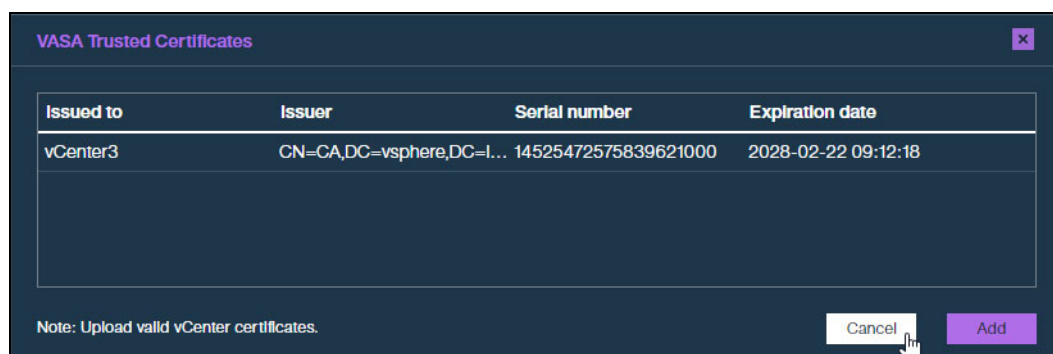


Figure 5-39 VASA trusted certificate

Configuring a VMware VVoLs data store

To configure a VMware VVoLs data store, complete the following steps:

- In the vCenter Web Interface, click **Home** → **Storage**, right-click the data center, and then select **Storage** → **New Datastore**, as shown in Figure 5-40.



Figure 5-40 Start the VMware New Datastore wizard

- On the **Location** tab, select your location and click **Next**.
- On the **Type** tab, select **VVoL** and then click **Next**, as shown in Figure 5-41.



Figure 5-41 Select the VVoL data store type

- On the Name and container selection tab, select the storage container of your choice for this VVOL data store, as shown in Figure 5-42, and click **Next**. You can recognize the storage space that you defined in IBM Spectrum Connect, in step b on page 73. This storage space can contain different services with different capabilities.

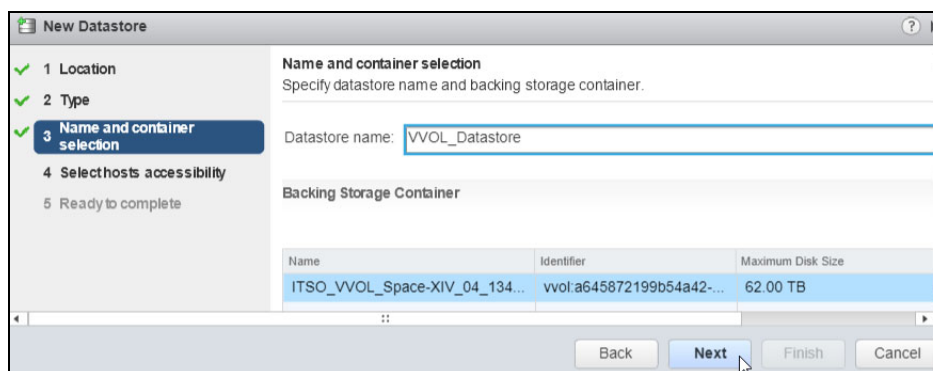


Figure 5-42 Select the storage container

- On the Select hosts accessibility tab, select the host of your choice and click **Next**.
- The Ready to complete tab is displayed. Click **Finish** to begin creating the data store.
- Click the created data store, click **Configure** and then **General**. The properties are shown in Figure 5-43.

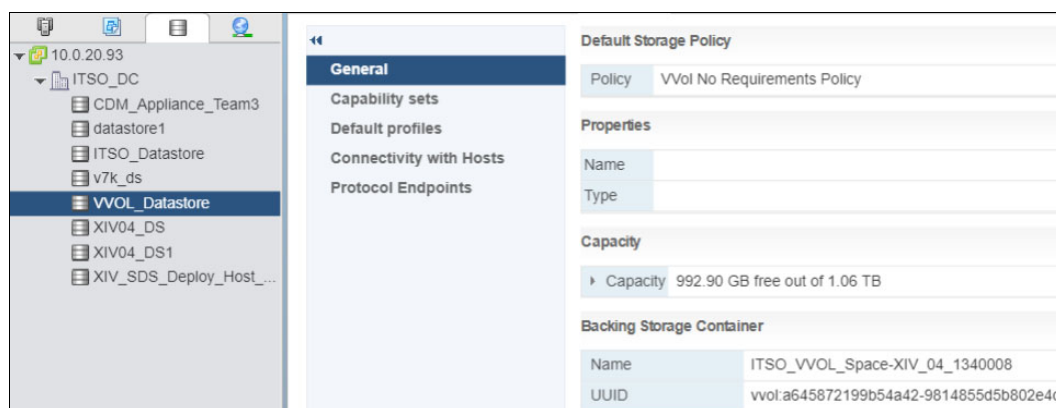


Figure 5-43 Create VVOL data store settings

Defining VMware storage policies

To define the VMware storage policies, complete the following steps:

1. From the vSphere Web Client Home, click **Policies and Profiles** → **VM Storage Policies**.
2. Click the **Create a new VM storage policy** icon, as shown in Figure 5-44.

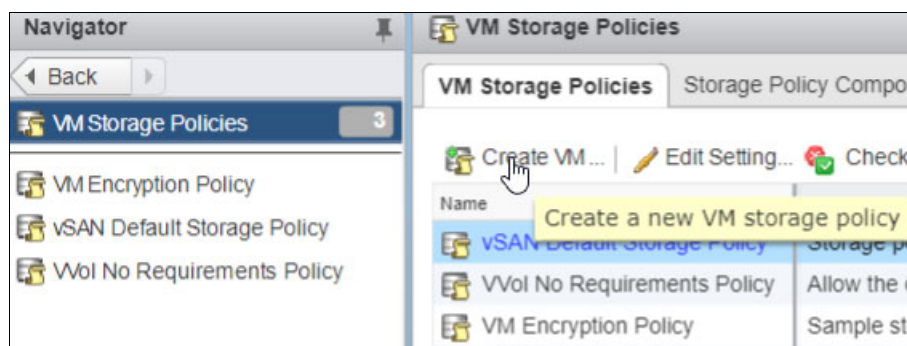


Figure 5-44 Create VM storage policy icon

3. Provide a name for your policy, and click **Next** twice.
4. On the 2b Rule set 1 tab, complete the wizard, as shown in Figure 5-45.

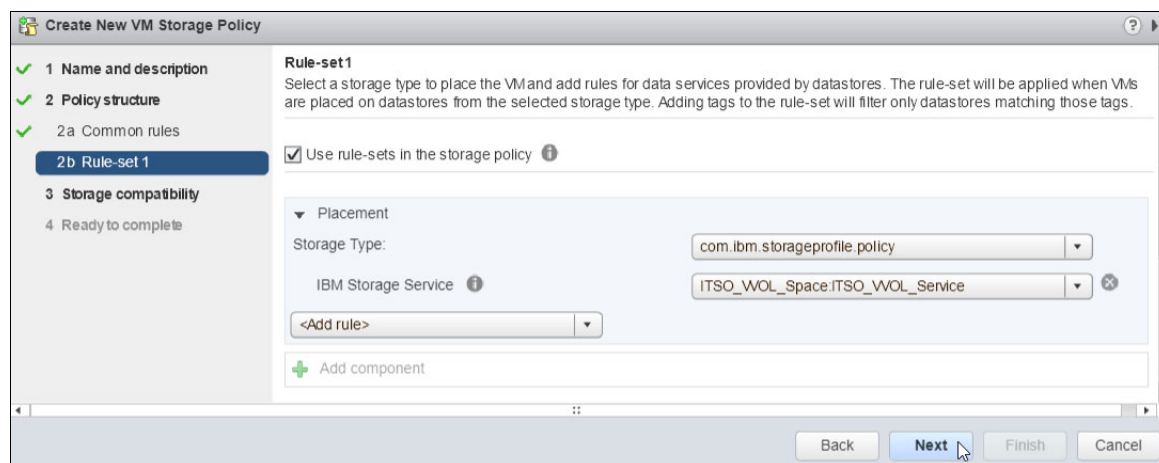


Figure 5-45 Define one rule

5. Repeat the previous steps for all of the rules that you need.

- Verify in the storage compatibility that the defined rule set corresponds to the VVoLs data store that you expect, as shown in Figure 5-46.

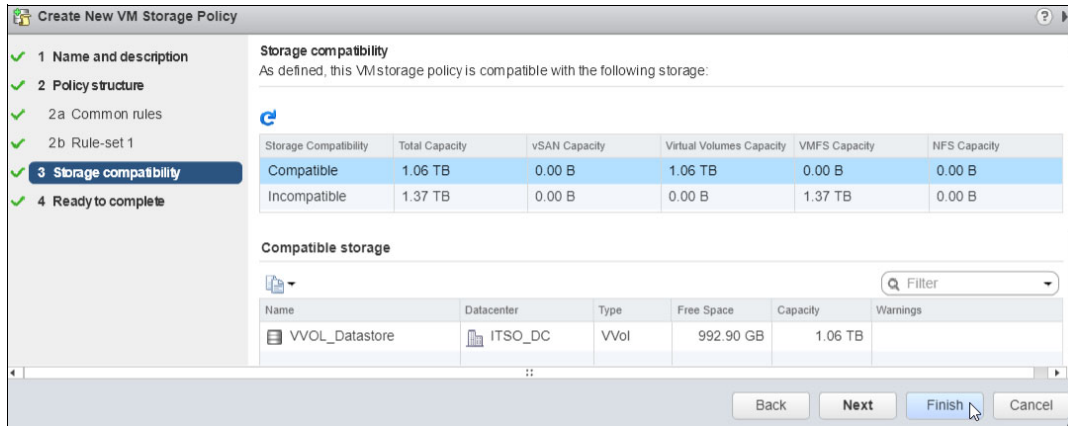


Figure 5-46 Verifying storage compatibility

- Click **Next** and then **Finish**. Your policy is created, as shown in Figure 5-47.




Figure 5-47 Created VM storage policy

- Whenever you are ready to create a VMware VM, you now can select a VM storage policy that takes you to the VVoLs data store that fulfills that policy, as shown in Figure 5-48.

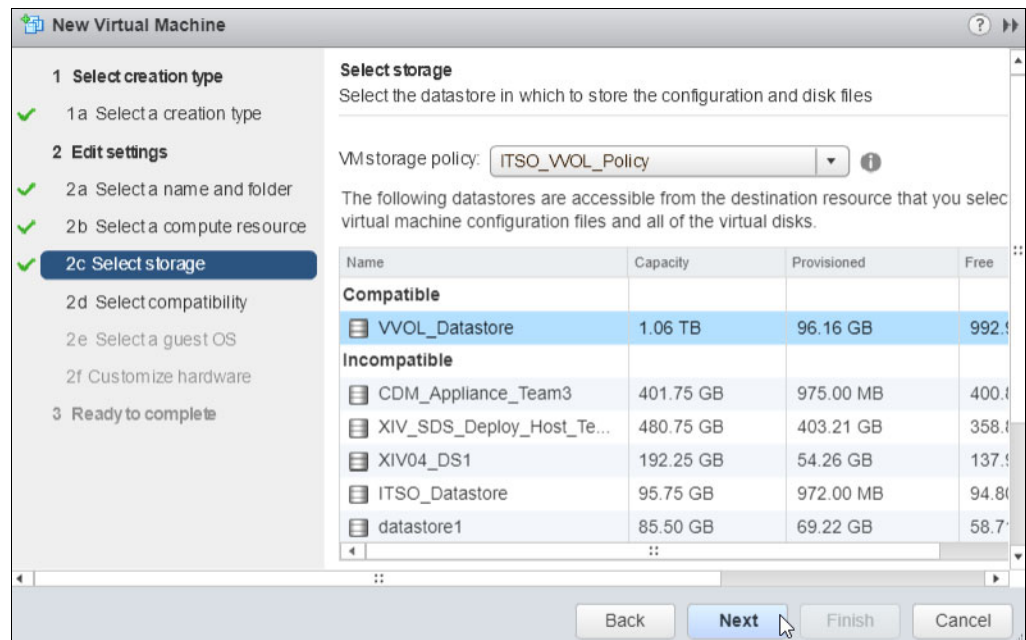


Figure 5-48 Select the appropriate storage when creating a VMware virtual machine



vSphere Web Client

This chapter describes the tasks that are performed by a storage and VMware administrator to configure and use IBM Storage Enhancements for vSphere Web Client.

It covers the following topics:

- ▶ vSphere Web Client illustration
- ▶ IBM Spectrum Connect configuration for IBM Storage Enhancements for vSphere Web Client
- ▶ Using IBM Storage Enhancements for vSphere Web Client

6.1 vSphere Web Client illustration

The IBM Storage Enhancements for VMware vSphere Web Client integrates into the VMware vSphere Web Client platform. It enables VMware administrators to independently and centrally manage their storage resources on IBM storage systems.

Note: In this chapter, when referring to information that relates to IBM XIV Storage System, IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM Spectrum Accelerate systems, the generic *storage system* is used unless specified otherwise.

When VMware is used with the IBM XIV Storage System, IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM Spectrum Accelerate systems, VMware administrators can self-provision volumes (logical unit numbers (LUNs)) in selected storage pools that are predefined by the storage administrators. The volumes are mapped to ESXi hosts, clusters, or data centers as logical drives that can be used for storing VMware data stores (virtual machine (VM) data containers).

The Storage Enhancements for VMware vSphere Web Client are automatically deployed and enabled for each vCenter server that is registered for vSphere Web Client services on the connected IBM Spectrum Connect Server.

The following steps are necessary for using Storage Enhancements for vSphere Web Client:

1. If not done so yet, perform the first-time IBM Spectrum Connect configuration as described in 4.2, “IBM Spectrum Connect first-time configuration” on page 54:
 - a. Define the IBM Spectrum Connect fully qualified domain name and high availability group, as shown in 4.2.1, “Initial Setup Wizard” on page 55.
 - b. Generate a server certificate, as shown in step 2 on page 56.
 - c. Set up VASA credentials, as shown in 4.2.2, “Setting up the VASA credentials” on page 58.
 - d. Set up storage system credentials, as shown in 4.2.3, “Adding an IBM storage system as a storage array” on page 58.

Important: If you must use IBM Storage Enhancements for vSphere Web Client in parallel with VVoLs, you must create an additional non-externally managed domain on the storage system and use the same Storage Integration Administrator user as for VVoLs.

- e. Add the ESXi hosts with ports, as described in 3.1.4, “Creating a host and volume, and mapping new LUNs” on page 34.
 - f. If domains are used, add the ESXi hosts with ports, as described in 5.2.2, “XIV configuration” on page 66 and add a user to the domain with storage admin role, except that you run VVoL in parallel, then add a user with Storage Integration Administrator role to the domain, as shown in “Creating a user and a protocol endpoint” on page 70.
 - g. Add the storage system, as shown in 4.2.3, “Adding an IBM storage system as a storage array” on page 58.
2. Perform the IBM Spectrum Connect Configuration for IBM Storage Enhancements for vSphere Web Client.
3. Use IBM Storage Enhancements for vSphere Web Client.

6.2 IBM Spectrum Connect configuration for IBM Storage Enhancements for vSphere Web Client

Complete the steps in the subsections that follow.

6.2.1 Adding the vCenter server in IBM Spectrum Connect

To add the vCenter Server to IBM Spectrum Connect, complete the following steps:

1. Click the left-pointing navigation arrow, as shown in Figure 6-1.




Figure 6-1 Open the IBM Spectrum Connect Interfaces pane

2. In the Interfaces pane, select **Add vCenter**, as shown in Figure 6-2.




Figure 6-2 Add a vCenter to IBM Spectrum Connect

3. Enter the appropriate IP address and credentials, as shown in Figure 6-3, and then click **Apply**.




Figure 6-3 Add a vCenter

4. The vCenter now appears in Applications, as shown in Figure 6-4.




Figure 6-4 Created vCenter

When a vCenter Server is added to IBM Spectrum Connect, it automatically deploys the IBM Storage Enhancements for VMware vSphere Web Client plug-in to the vSphere Web Client. The warning symbol disappears when a storage service is attached.

6.2.2 Controlling the IBM Storage Enhancements for vSphere Web Client plug-in on vCenter

To confirm that the plug-in was successfully deployed and enabled in vSphere, complete the following steps:

1. To log in to vSphere Web Client, open a browser and enter the following address:

`https://vCenter_IP_address:9443`

Enter the credentials and click **Login**, as shown in Figure 6-5.




Figure 6-5 Log in to vSphere Web Client

2. Select **Administration** → **Solutions** → **Client Plug-ins**, as shown in Figure 6-6.

Client Plug-ins				
Check for New Plug-ins				
Name	Vendor	Version	Description	State
Virtual Infrastructure	VMware, Inc.	6.6.1	Virtual SAN Web Client (build 58...	Enabled
IBM Storage	IBM, Corp.	3.4.0.9947	IBM Storage Enhancements for V...	Enabled
Hybrid Cloud Mgr Prev...	VMware	6.5.0.10000	VMware vCloud Air Hybrid Cloud ...	Enabled
SR File Upload Plugin	VMware	6.5.0.10000	Uploads files as attachments to e...	Enabled
vCenter Orchestrator p...	VMware	6.5.0.10000	vCenter Orchestrator plugin	Enabled
SSO Admin UI plugin	VMware	6.5.0.10000	SSO Admin UI plugin	Enabled

Figure 6-6 Verify plug-in

6.2.3 Defining a storage space

Within the IBM Spectrum Connect server, virtual storage is defined with a *storage service* and a *storage space*. Therefore, you must first complete the following steps to define a storage space:

1. Click **Allocation** on the left. Click **ITSO_VVOL_Space** from the drop-down box and select **Add New Storage Space** to add a storage space to IBM Spectrum Connect, as shown in Figure 6-7.



Figure 6-7 Add New Storage Space

2. Provide a name and description for your new storage space and click **Apply**, as shown in Figure 6-8.

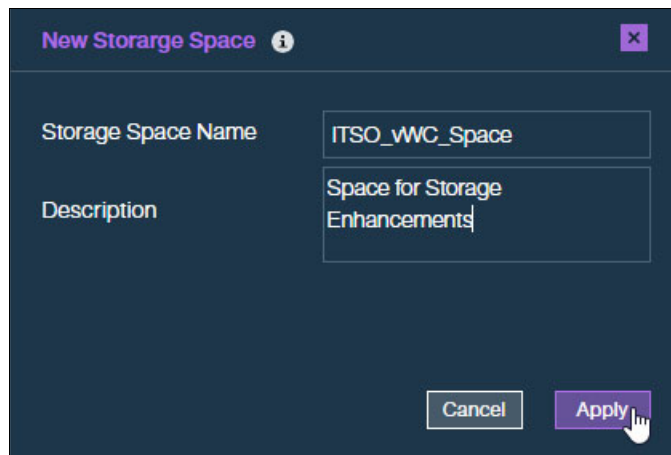


Figure 6-8 Add a storage space

The IBM Spectrum Connect GUI automatically brings you to this newly created storage space.

6.2.4 Configuring a storage service

Now that storage space is defined, a storage service must be configured. A storage service is the combination of storage resources and associated user-defined policies, such as encryption and thin provisioning. To add a service to the newly created storage space, complete the following steps:

1. From the newly created storage space, click the **Storage Services +** icon next to the **Default_Space** drop down box.
2. When the New Service form opens, complete it as shown in Figure 6-9 on page 93. Specify features that are fulfilled by this service according to your needs. Table 6-1 explains the parameters and description for capabilities.

Table 6-1 Parameters and description for capabilities

Parameters	Description
Encryption	Enables encryption for the service. If enabled, you can attach only encrypted storage resources to the service.
Flash	Enables utilization of a storage resource on a flash-based storage resource, which can be one of the following storage systems: IBM FlashSystem 900, IBM FlashSystem V9000, or any of the Storwize family systems.
Space Efficiency	Enables storage space efficiency features for the service. When selected, you can configure the service to be attached to a thick- or thin-provisioned storage resource.
QoS	Enables the use of the Quality of Service (QoS) feature for the service. QoS is applicable to volumes (Max Independent Performance) or storage resources (Max Shared Performance), setting the IOPS and bandwidth limits within the following ranges: <ul style="list-style-type: none">► IOPS: 0 - 100000► BW (bandwidth): 0 - 10000 MBps
Availability	Enables the use of IBM HA technology for highly available storage deployments of IBM SAN Volume controller (Stretched Cluster).
Data Reduction	Enables the use IBM Real-time Compression with or without data deduplication.
Replication	Enables synchronous mirroring for the service. Not available for VVoL-enabled services. It applies as HyperSwap to Spectrum Virtualize family and as synchronous replication to Spectrum Accelerate family.

Figure 6-9 shows completing the form.

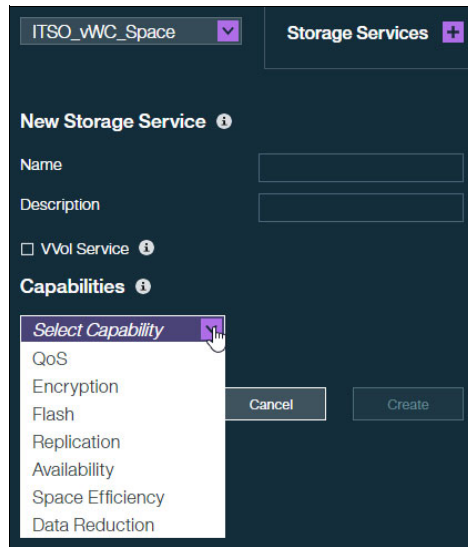


Figure 6-9 Add a service

The service now appears in the newly created space, as shown in Figure 6-10.




Figure 6-10 New service added

6.2.5 Adding a storage resource

To add a storage resource for this service, complete the following steps:

1. Right-click the service and select **Manage Resources**, as shown in Figure 6-11.

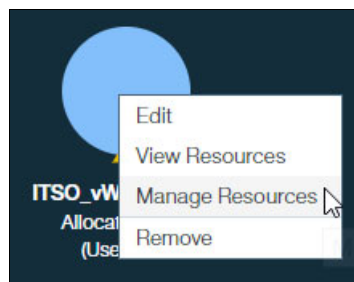


Figure 6-11 Service Manage Resources

2. Click the + icon to add a resource, as shown in Figure 6-12.




Figure 6-12 Service Add Resource

3. Enter the appropriate details for the new storage resource and click **Add**, as shown in Figure 6-13. Confirm the operation by clicking **Ok**.

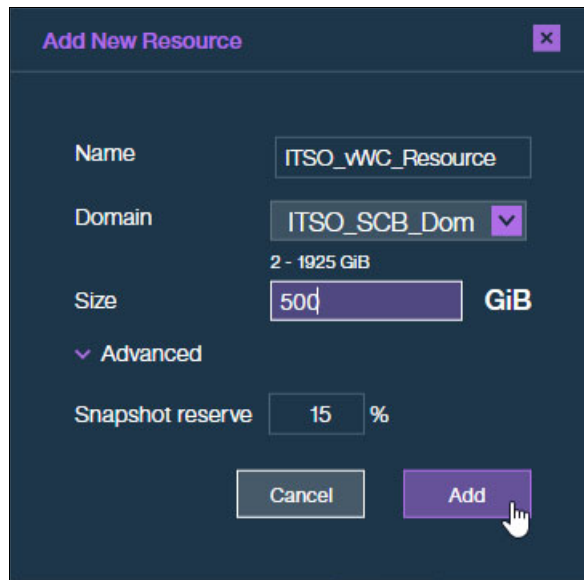


Figure 6-13 Service Add New Resource

4. To add a vCenter server in IBM Spectrum Connect, complete the steps that are described in 6.2.1, “Adding the vCenter server in IBM Spectrum Connect” on page 89. Click the left-arrow navigation pointer, as shown in Figure 6-14.

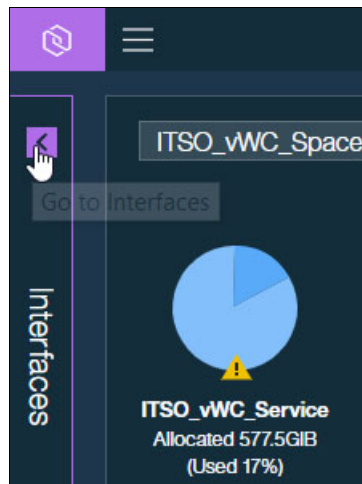


Figure 6-14 Interface pane

5. Delegate the service to the vCenter by clicking the vCenter and then the **Delegate** icon, as shown in Figure 6-15.




Figure 6-15 Delegate Service to vCenter

6. Click **Ok** to confirm, as shown in Figure 6-16.




Figure 6-16 Delegate Service to vCenter confirmation

6.3 Using IBM Storage Enhancements for vSphere Web Client

This section explains how to use IBM Storage Enhancements for vSphere Web Client.

6.3.1 Reviewing the available storage enhancements

Complete the following steps:

1. To create a volume on an IBM Spectrum Accelerate family system, open vSphere Web Client, click **Home** → **Storage**, right-click your data center, and select **All IBM Storage Actions** → **Create New Volume**, as shown in Figure 6-17.

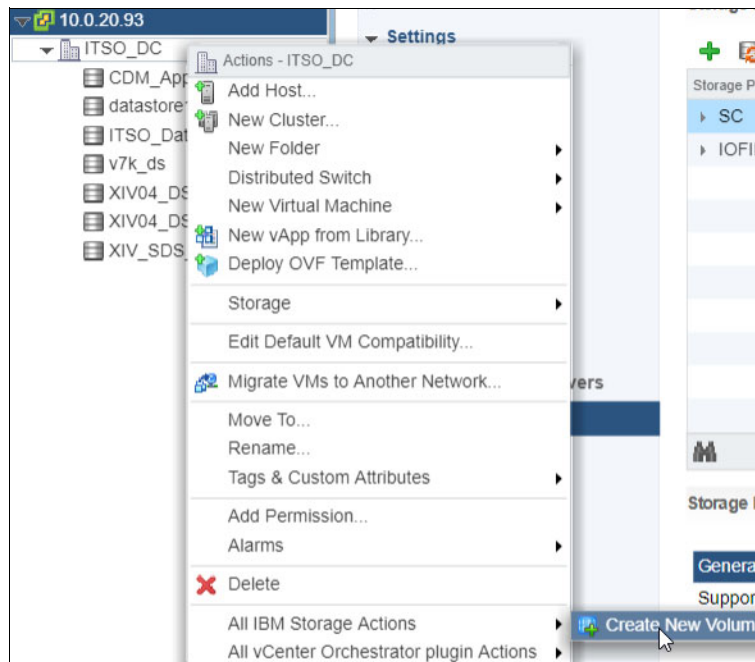


Figure 6-17 IBM Storage Enhancements for VMware vSphere Web Client plug-in

2. When the Create New Volume window opens, complete the appropriate information as shown in Figure 6-18 on page 97 and click **OK**:
 - a. For Storage Services, select the storage resource that you created in 6.2.4, “Configuring a storage service” on page 92.
 - b. For Suggested LUN, by default, the client uses with the next available LUN ID, but you can select a different ID if needed.

According to selected hosts mapping, 1 volume(s) with a size of 100 GiB will be created on the storage service named ITSO_WVC_Service.

Host Mapping: 10.0.20.23 Custom

Volume Size: 100 GiB Max Provision Size: 480 GiB

Number of Volumes: 1

Volume Name: ITSO_WVC_Vol

Storage Service: ITSO_WVC_Service (ITSO_WVC_Space)

Suggested LUN: 2

Consistency Group: N/A

Service Capabilities

Space Efficiency: Thin

OK Cancel

Figure 6-18 Create a storage system volume with the Web Client plug-in

3. Click the vCenter and click **More Objects** → **IBM Storage Volumes** to see the new volume, as shown in Figure 6-19.

Storage Device Na...	Volume Identifier	Storage Array	Volume Name	Storage Type	Volume Size (GiB)	Usage	Serial	Path S
ITSO_WVC_Vol	naa.6001738ccc	A9000	ITSO_WVC_Vol	FlashSystem A9	100.0	Not in use	1305673000000	Rou

Figure 6-19 A storage system volume that is created and mapped from the vSphere Web Client

Important: This volume was created and mapped directly from the vSphere Web Client without needing to access Hyper-Scale Manager (GUI) or XCLI directly.

6.3.2 Additional storage control and monitoring options

The IBM Storage Enhancements for VMware vSphere Web Client plug-in provides additional rich storage control and monitoring capabilities beyond simple volume creation and mapping. These capabilities are accessible through the IBM Storage Volumes section in the vSphere Web Client.

To see a summary of a volume, click **Home** → **Global Inventory List** → **vCenter Servers**. Double-click the vCenter server, select **IBM Storage Volumes** and the volume, and then click the Summary tab. Figure 6-20 shows an example.

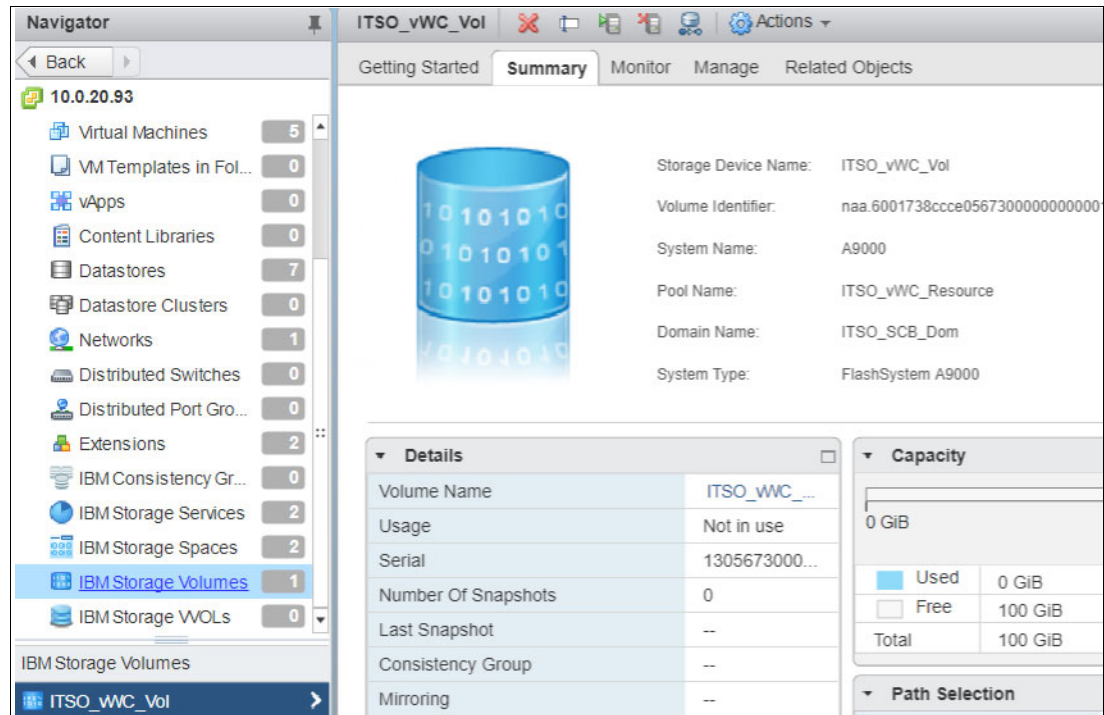


Figure 6-20 The Summary window for a storage volume in the vSphere Web Client

Figure 6-21 illustrates the extra actions that the plug-in provides for storage system volumes.

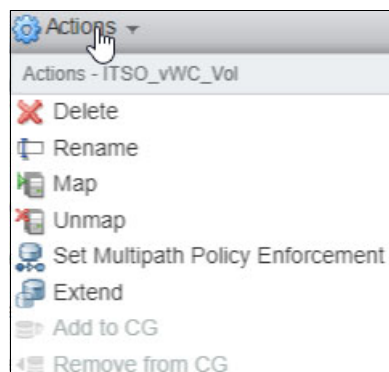


Figure 6-21 The available actions for an IBM Storage Volume in the vSphere Web Client

The IBM Storage Enhancements for VMware vSphere Web Client plug-in also provides a volume management window where you can do the following tasks for a storage system volume:

- ▶ Adjust volume settings, such as volume name, multipath policy enforcement, and volume size.
- ▶ View and modify volume host mappings.

Figure 6-22 shows Settings tab of the Manage window for an IBM Storage Volume in the vSphere Web Client.




Figure 6-22 Volume settings tab in the vSphere Web Client

Figure 6-23 shows the Host Mapping section of the Manage page for an IBM Storage Volume in the vSphere Web Client.



Figure 6-23 The Host Mapping section of the Manage page in the vSphere Web Client

If you want to see a summary for a service, click **IBM Storage Services** and the service, and then click the **Summary** tab, as shown in Figure 6-24.

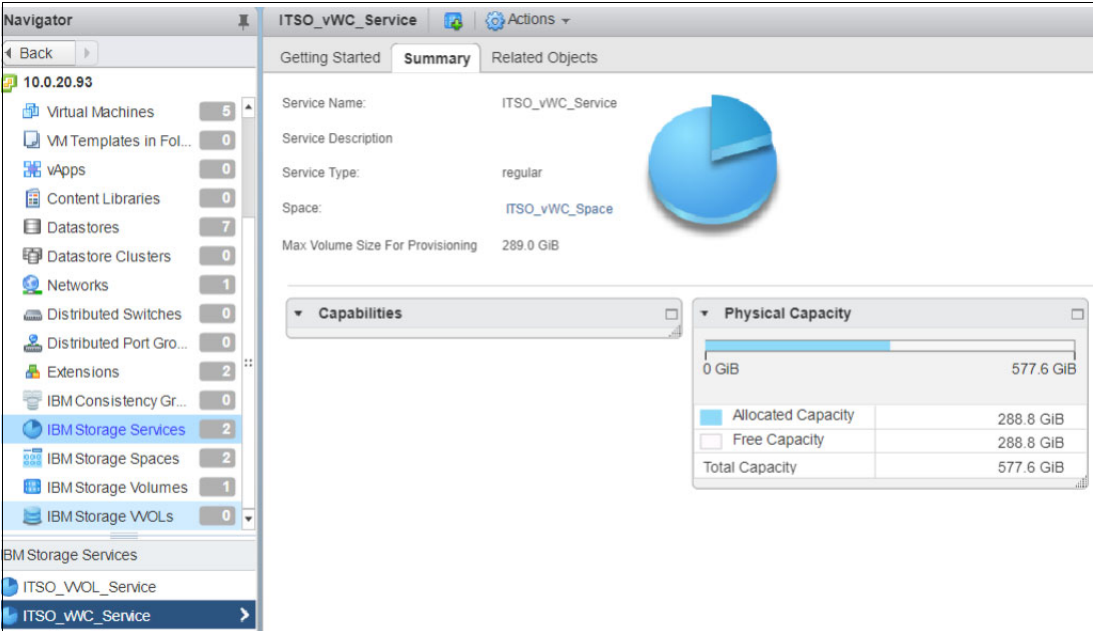


Figure 6-24 Service summary tab

To see the volume that was created before, click **Related Objects**, as shown in Figure 6-25.

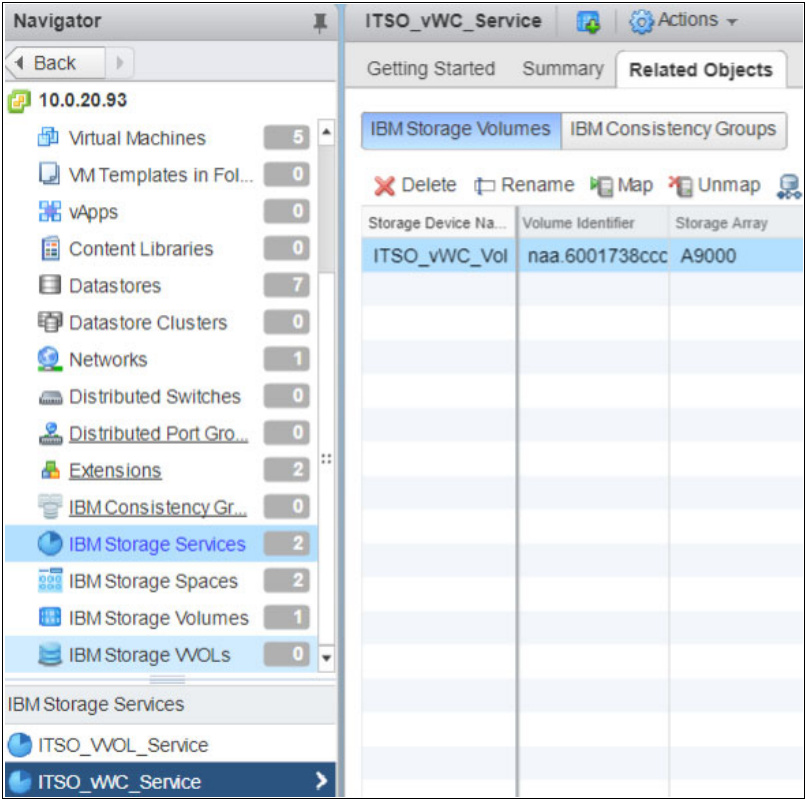


Figure 6-25 Service Related Objects tab



VMware vRealize Operations Manager

This chapter describes the tasks that are performed by storage and VMware administrators to configure IBM Spectrum Connect for vRealize Operations Manager (vROps). It covers the following topics:

- ▶ Configuration of IBM Spectrum Connect server for vROps
- ▶ Install the storage management package onto vROps
- ▶ IBM XIV dashboards in vROps

7.1 Configuration of IBM Spectrum Connect server for vROps

The integration of IBM Spectrum Connect with vROps allows monitoring and analysis of the storage system health, performance, and capacity. It can dynamically cope with policy-governed workflows to maintain SLAs.

To start the configuration, you must download a PAK file by completing the following steps:

1. To configure IBM Spectrum Connect to work with vROps, log on to IBM Spectrum Connect and click the right-arrow navigation pointer, as shown in Figure 7-1.

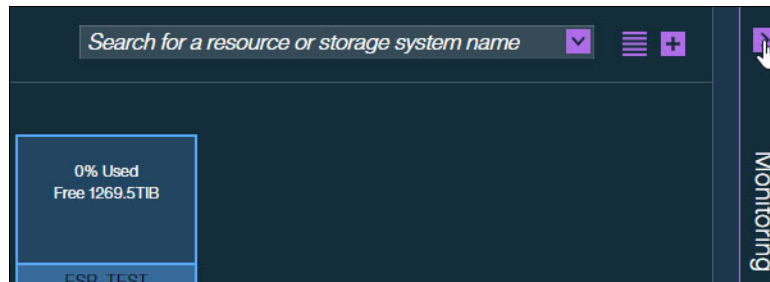


Figure 7-1 Open Monitoring window

2. Click **Download PAK file**, as shown in Figure 7-2. Your browser prompts you to save the PAK file locally on your workstation.




Figure 7-2 Download PAK file

7.2 Install the storage management package onto vROps

Now, you must install the storage management package on to vROps by completing the following steps:

1. In your browser, log in to the vROps administration GUI of your vROps server:

`https://vROps_UI_IP_address/ui`

Click **Administration** → **Solutions** and click the + icon, as depicted in Figure 7-3.

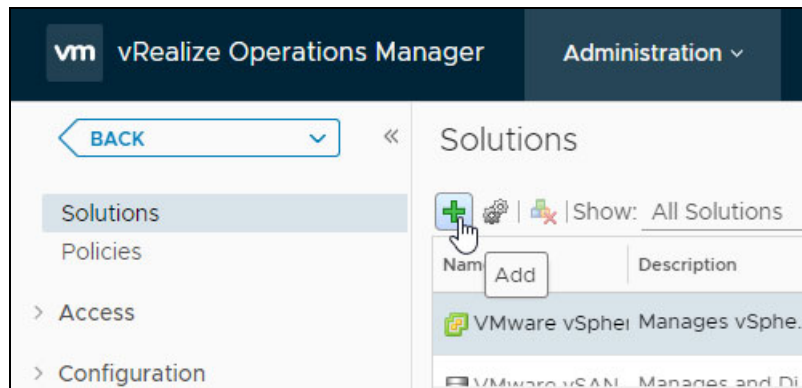


Figure 7-3 Add solution

2. Click **Browse**, locate the PAK file that is on your workstation, and click **UPLOAD**, as shown in Figure 7-4.

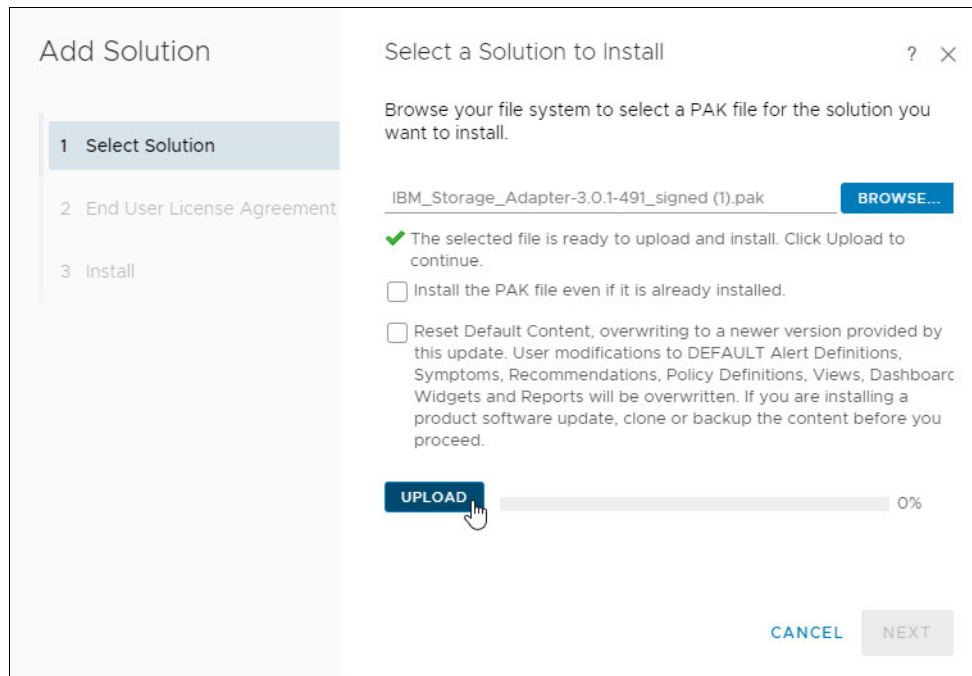


Figure 7-4 Find the PAK file

3. When the file is uploaded and ready to install, click **NEXT**, as shown in Figure 7-5.

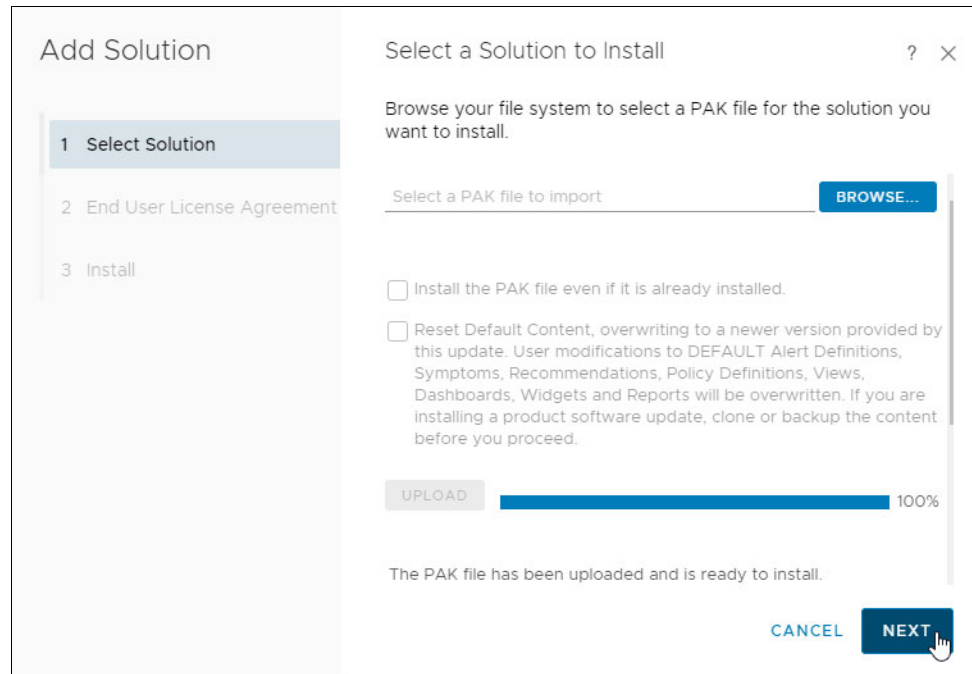


Figure 7-5 Ready to install

4. Accept the License Agreement and click **Next**.
5. When the installation completes, click **FINISH**, as shown in Figure 7-6.

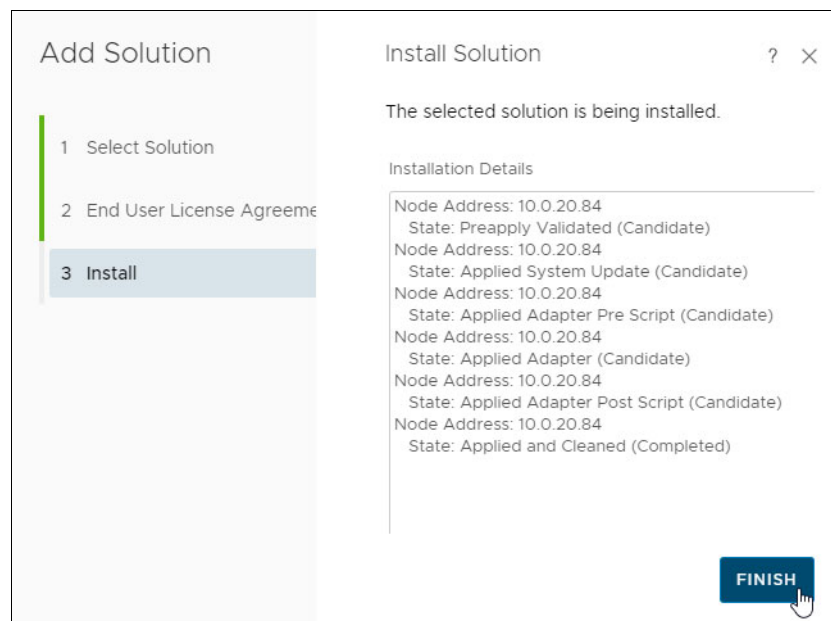
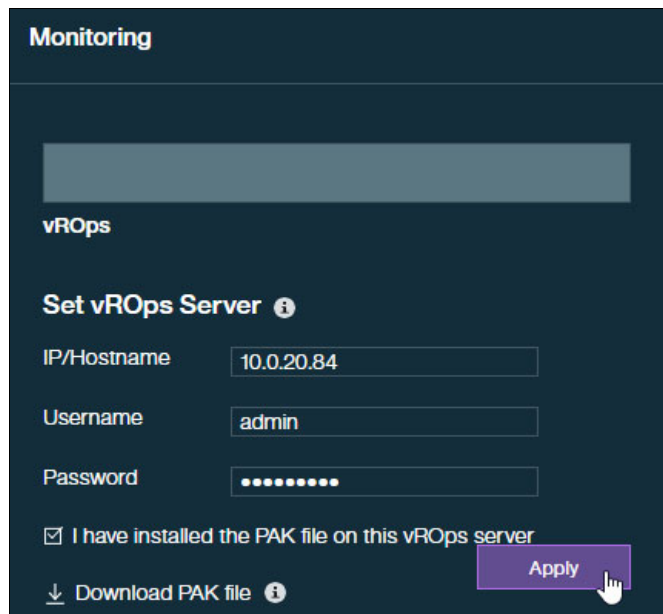


Figure 7-6 Installation complete

6. Proceed back to IBM Spectrum Connect GUI Monitoring window and complete the Set vROps Server fields. Add the Internet Protocol (IP) address, user name, and password, select **I have installed the PAK file on this vROps server**, and click **Apply**, as shown in Figure 7-7.



The screenshot shows the 'Monitoring' section of the IBM Spectrum Connect GUI. Under the 'vROps' heading, there is a 'Set vROps Server' form. The form includes fields for 'IP/Hostname' (10.0.20.84), 'Username' (admin), and 'Password' (masked with dots). A checkbox labeled 'I have installed the PAK file on this vROps server' is checked. Below the checkbox is a link 'Download PAK file' with a download icon and an information icon. An 'Apply' button is located at the bottom right of the form, with a mouse cursor hovering over it.

Figure 7-7 Configure monitoring

7. Go to the Storage Systems window and click the **Play** icon to allow VROps to start monitoring the storage system, as shown in Figure 7-8.

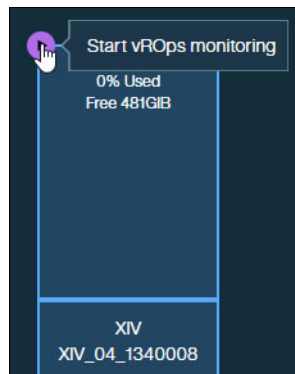


Figure 7-8 Start vROps monitoring

8. Click **OK** to confirm vROps to start monitoring the storage system, as shown in Figure 7-9.

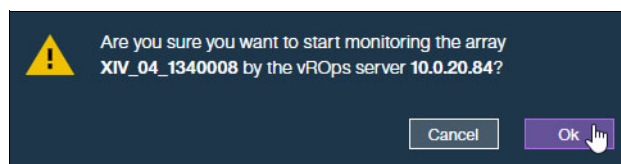


Figure 7-9 Confirmation to monitor the storage system

7.3 IBM XIV dashboards in vROps

After the PAK file is successfully installed on the vROps server and vROps monitoring is configured on the IBM Spectrum Connect, log in to the vROps GUI to view the storage systems:

`https://vROps_UI_IP_address/ui`

Then, complete the following steps:

1. In the vROps GUI, click **Dashboard** and select the storage system that you want to view, as shown in Figure 7-10.

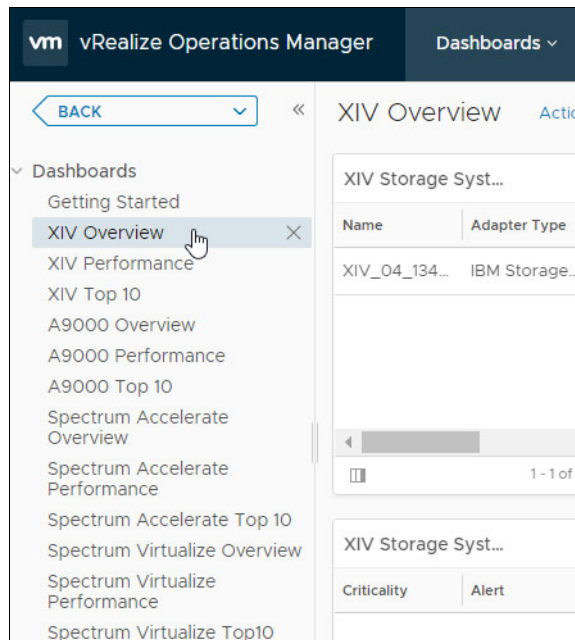


Figure 7-10 vROps storage view

Three main dashboards are available for each storage system in vROps:

- Overview

Provides clickable icons that represent the relationships between all the virtual elements and storage elements that are in use.

- Performance

Provides health and performance information for the storage resources. Performance information that is presented in the dashboard is collected for a defined period (the default period is 5 minutes).

- Top 10

Represents the top 10 IBM volumes and hosts in all storage systems that are monitored by the vROps Manager.

The top 10 dashboards for IBM XIV Storage System, IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM Spectrum Accelerate dashboards include the following information:

- Top 10 volumes by IOPS (in the last hour)
- Top 10 volumes by IOPS (in the 24 hours)
- Top 10 volumes by throughput (in the last hour)

- Top 10 volumes by throughput (in the 24 hours)
- Top 10 hosts by IOPS (in the last hour)
- Top 10 hosts by IOPS (in the 24 hours)
- Top 10 hosts by throughput (in the last hour)
- Top 10 hosts by throughput (in the 24 hours)

The top 10 dashboards for the storage systems that run IBM Spectrum Virtualize include the following information:

- Top 10 VDisks by read operations per second (last hour)
- Top 10 VDisks by write operations per second (last hour)
- Top 10 VDisks by read blocks per second (last hour)
- Top 10 VDisks by write blocks per second (last hour)
- Top 10 VDisks by worst read response in use since the last statistics collection
- Top 10 VDisks by worst write response in use since the last statistics collection
- Top 10 VDisks by average transfer response time in use (last hour)

2. Click **Dashboard List** → **IBM Storage** → **XIV Overview** and click one volume, as shown in Figure 7-11. The dark green color means that the resources are related.

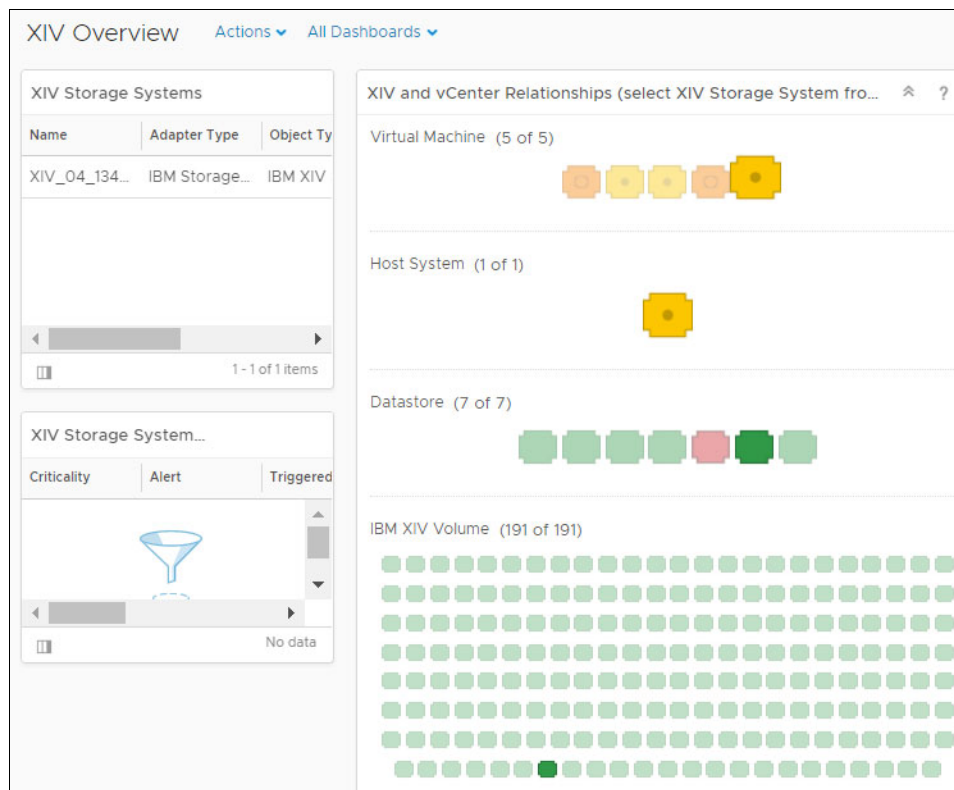


Figure 7-11 XIV Overview

3. Click **XIV Performance** and click the volume in the lower left, as shown in Figure 7-12. The volume and all the related objects are shown in the health tree.

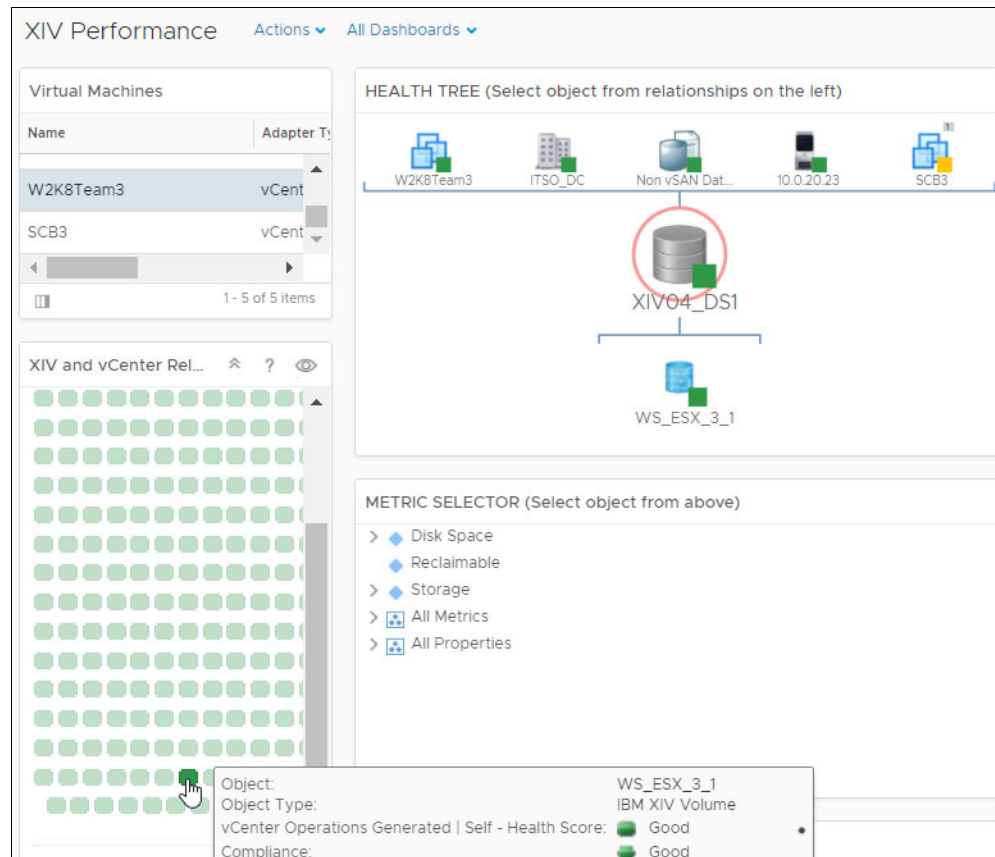


Figure 7-12 FlashSystem A9000 Performance

4. Click the volume in the health tree and scroll down. Various metrics are selectable. Click one and the metric graph is displayed, as shown in Figure 7-13.

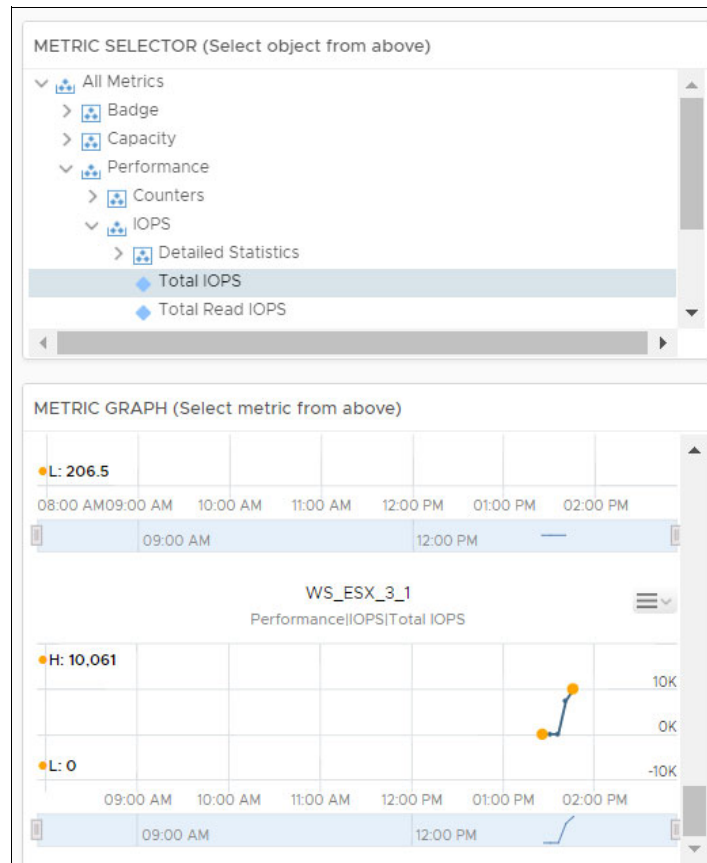


Figure 7-13 XIV performance metrics

5. Click **XIV Top 10** and the top 10 volumes and hosts are displayed, as shown in Figure 7-14.

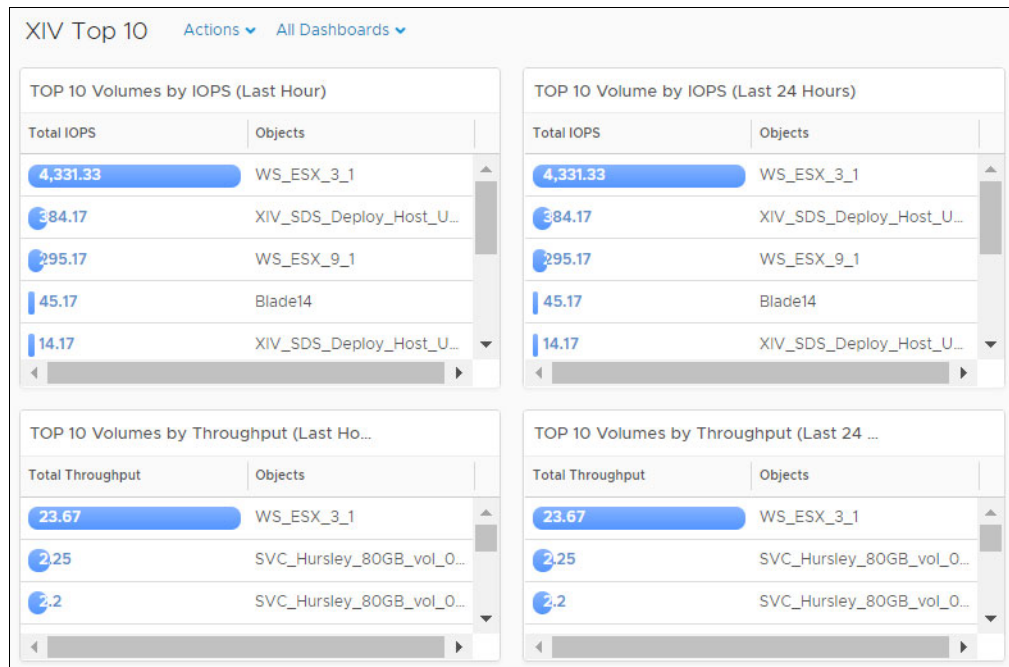


Figure 7-14 A9000 Top 10



IBM Spectrum Connect configuration for vRealize Orchestrator

This chapter describes the tasks that are performed by a storage and VMware administrator to configure IBM Spectrum Connect for VMware vRealize Orchestrator (vRO).

It covers the following topics:

- Configuration of IBM Spectrum Connect Server for vRO
- Running workflows in vRO

8.1 Configuration of IBM Spectrum Connect Server for vRO

You can use the vRO server to create workflows for VMware environments that further automate administrative actions and prevent inconsistent configurations, which allows for more self-service functions. vRO integration through IBM Spectrum Connect provides the ability to create, extend, map, unmap, and delete volumes on an IBM Spectrum Accelerate family system without any VMware or storage administrator actions.

At the time of writing, only vRO 5.5.x, 6.0.x, 7.0.x, 7.1.x, 7.2, 7.3 are supported. For more information about supported versions, see IBM Knowledge Center:

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

The following assumptions are made in this illustration:

- ▶ You have an installed, configured, and started the vRO server/application.
- ▶ You have an installed instance of IBM Spectrum Connect.

Important: Replace the self-signed certificate and key that come with the default IBM Spectrum Connect by either generating a new certificate and key file on the IBM Spectrum Connect itself or by using an externally generated certificate and key file.

To configure IBM Spectrum Connect Server for vRO, complete the following steps:

1. Click the left-arrow navigation pointer, as shown in Figure 8-1.




Figure 8-1 Opening IBM Spectrum Connect Configuration

2. Click **Add Interface** and select **Add vRO**, as shown in Figure 8-2.




Figure 8-2 Add vRO

3. Click **Ok** to confirm the creation of the vRO application, as shown in Figure 8-3.




Figure 8-3 Confirmation of the vRO application creation

4. Right-click **vRO** and select **Modify**, as shown in Figure 8-4.




Figure 8-4 Modify vRO

5. Click **Download plug-in package** from the vRO Settings window in IBM Spectrum Connect, as shown in Figure 8-5.

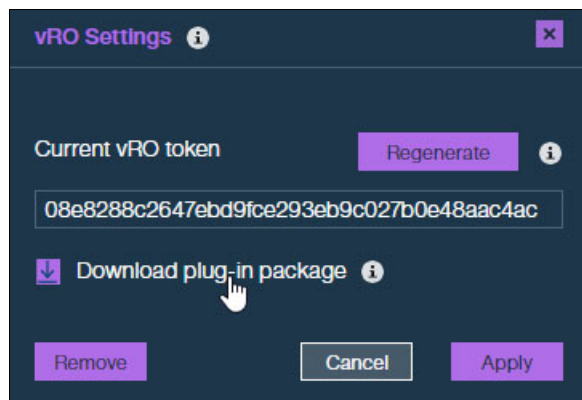
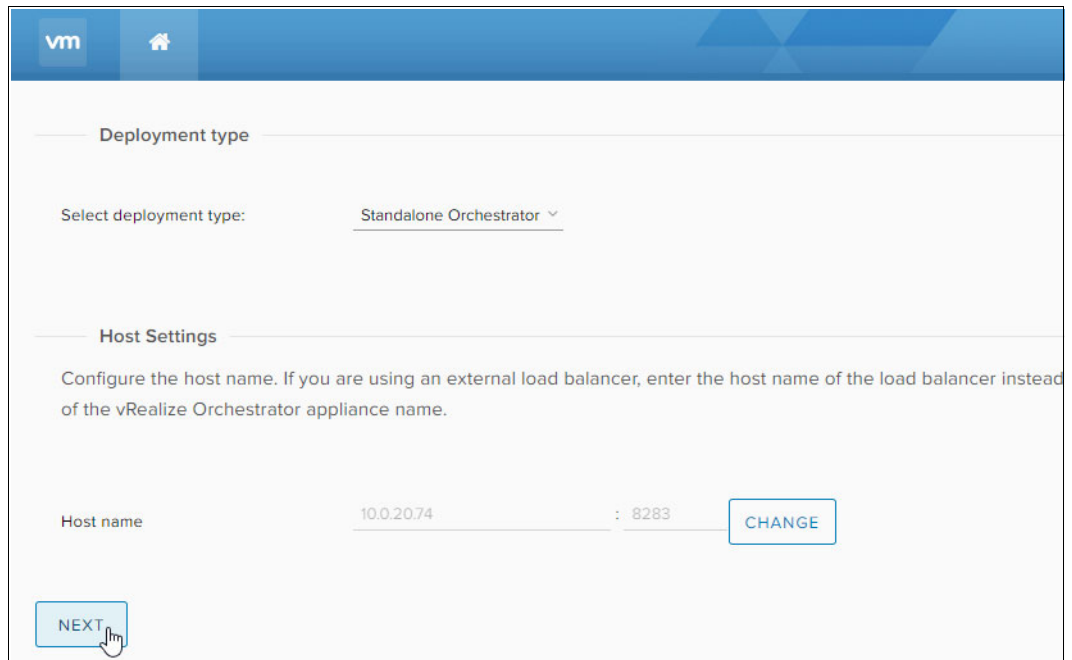


Figure 8-5 Download the plug-in package from IBM Spectrum Connect

6. Go to the vRO configuration web interface:
<https://Orchestrator-IP-address:8283/vco-controlcenter>

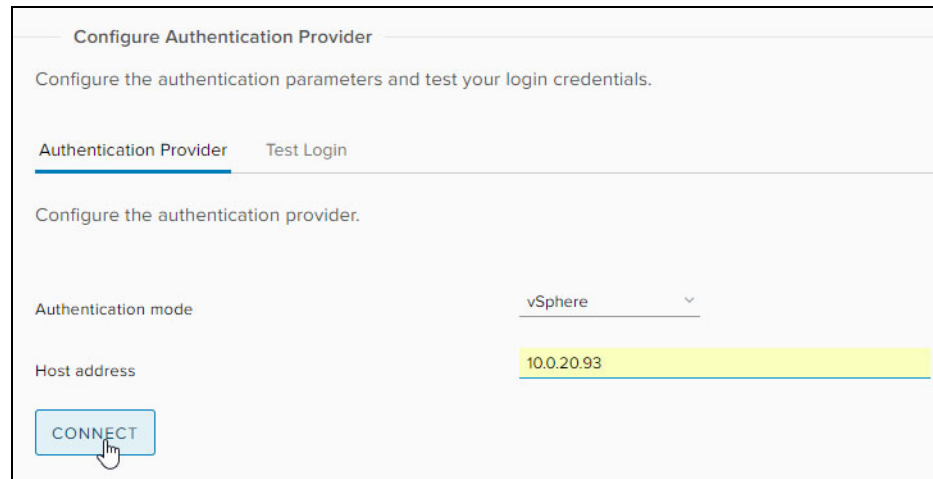
Select the **Standalone Orchestrator** and click **Next**, as shown in Figure 8-6.



The screenshot shows a web interface for configuring a Standalone Orchestrator. At the top, there is a blue header bar with a 'vm' logo and a home icon. Below the header, the 'Deployment type' section is active, showing 'Standalone Orchestrator' selected in a dropdown menu. The 'Host Settings' section follows, with a text prompt: 'Configure the host name. If you are using an external load balancer, enter the host name of the load balancer instead of the vRealize Orchestrator appliance name.' Below this, the 'Host name' field contains '10.0.20.74' and the port is set to '8283'. A 'CHANGE' button is located to the right of the port field. At the bottom left, a 'NEXT' button is highlighted with a mouse cursor.

Figure 8-6 Configuration: Standalone Orchestrator


Select the vCenter IP address and click **CONNECT**, as shown as in Figure 8-7.



The screenshot shows a web interface for configuring authentication. The 'Configure Authentication Provider' section is active, with a text prompt: 'Configure the authentication parameters and test your login credentials.' Below this, there are two tabs: 'Authentication Provider' and 'Test Login'. The 'Authentication Provider' tab is selected. The 'Authentication mode' dropdown menu is set to 'vSphere'. The 'Host address' field contains '10.0.20.93', which is highlighted in yellow. A 'CONNECT' button is located at the bottom left, with a mouse cursor pointing to it.

Figure 8-7 Configuration: Connect to vCenter

7. Click **ACCEPT CERTIFICATE**, as depicted in Figure 8-8.

 **Warning**

The certificate must be accepted to continue.

Authentication mode

vSphere

Host address

10.0.20.93

Common Name:10.0.20.93
Organization:
Serial number: 00:00:00:00:00:00:00:00:00:00:00:00:e5:6f:4d:87:9a:2d:e7:1c
Signature algorithm: SHA256withRSA
Fingerprint (MD5): 19:d5:b8:f1:19:85:ba:52:de:c8:36:03:a4:ae:0e:91
Fingerprint (SHA-1): fb:48:fd:97:e8:be:f9:45:25:35:06:88:45:01:db:1d:eb:0c:fd:a7
Valid from: Feb 27, 2018
Valid until: Feb 22, 2028

ACCEPT CERTIFICATE

SAVE CHANGES

RESET

CANCEL

Figure 8-8 Configuration: Accept vCenter certificate

8. Click **REGISTER** to register vRO with the vCenter, as depicted in Figure 8-9.

Authentication mode

vSphere

Host address

10.0.20.93

IDENTITY SERVICE

User name

vsphere.local/Administrator

Password

.....

Default tenant

vsphere.local

REGISTER

Figure 8-9 Configuration: Register vRO

9. Select the admin group and click **SAVE CHANGES**, as shown in Figure 8-10.

The screenshot shows a configuration page for vSphere. The 'Authentication mode' is set to 'vSphere'. The 'Host address' is '10.0.20.93' with an 'UNREGISTER' button. The 'Default tenant' is 'vsphere.local' with a 'CHANGE' button. The 'Admin group' is 'vsphere.local\Administrators' with a 'SEARCH' button. At the bottom, there are three buttons: 'SAVE CHANGES' (highlighted with a mouse cursor), 'RESET', and 'CANCEL'.

Figure 8-10 Configuration: Admin group and save changes

10. Select **Home** → **Manage** → **Certificates**, as shown in Figure 8-11.




Figure 8-11 Configuration: Certificates

11. Select **IMPORT** → **Import from URL**, as depicted in Figure 8-12.

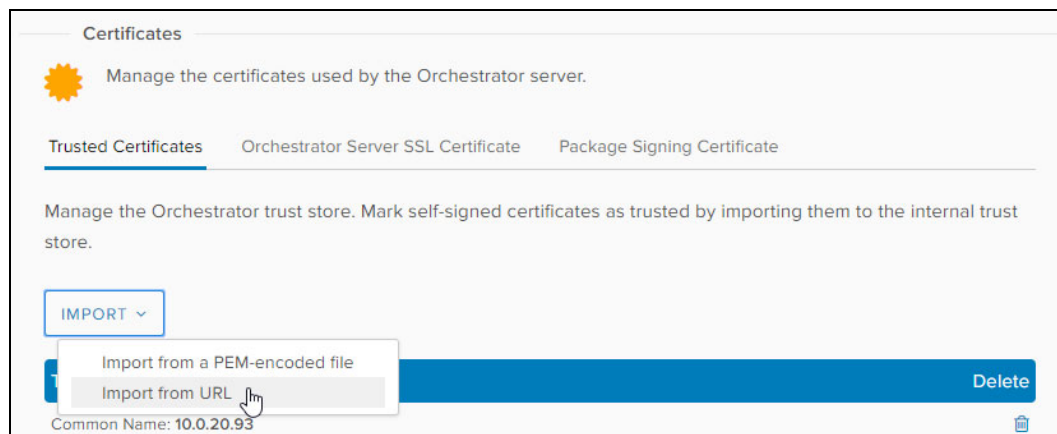


Figure 8-12 Configuration: Import from URL

12. Select the Spectrum Connect server IP address and port, and click **IMPORT**, as illustrated in Figure 8-13.

IMPORT FROM URL

URL

https://10.0.20.44:8440

Proxy URL

:

Port

Use Proxy

☐

CANCEL

IMPORT

Figure 8-13 Configuration: Import from URL from SC server step 1

13. Click **IMPORT**, as illustrated in Figure 8-14.

IMPORT THIS CERTIFICATE?

Common Name:

Organization:

Serial number: 00:00:00:00:00:00:00:00:00:00:00:00:00:00:da:42:44:c8:b0:07:b6

Signature algorithm: SHA256withRSA

Fingerprint (MD5): 20:50:14:9a:5f:be:77:ad:5b:5a:b3:82:03:5c:30:12

Fingerprint (SHA-1): 8e:5b:78:b9:31:fc:87:fb:69:c6:d8:3e:1d:d5:bc:59:f2:65:55:9a

Valid from: Feb 27, 2018

Valid until: Feb 25, 2028

CANCEL

IMPORT

Figure 8-14 Configuration: Import from URL from SC server step 2

14. Select **Home** → **Plug-Ins** → **Manage Plug-Ins**, as shown in Figure 8-15.




Figure 8-15 Configuration: Manage Plug-Ins

15. Select **Browse** to locate the plug-in that you downloaded in step 5, then click **INSTALL**, as shown in Figure 8-16.

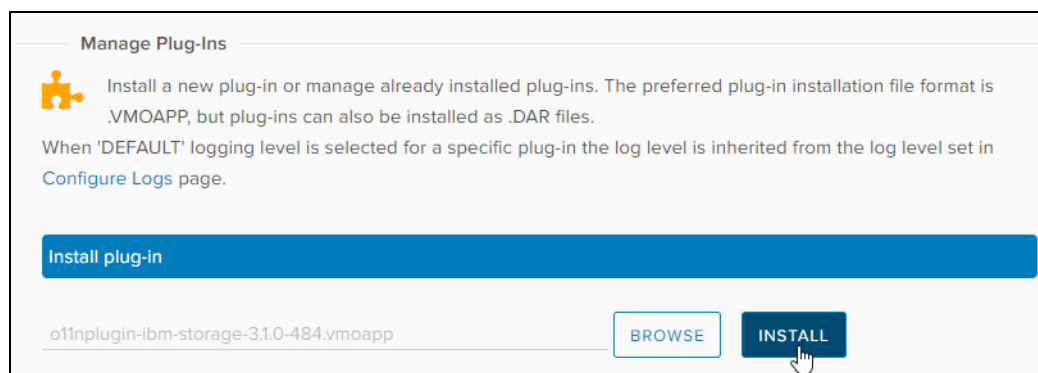


Figure 8-16 Configuration: Browse and install Plug-In

16. Select **Accept EULA** and click **INSTALL**, as illustrated in Figure 8-17.

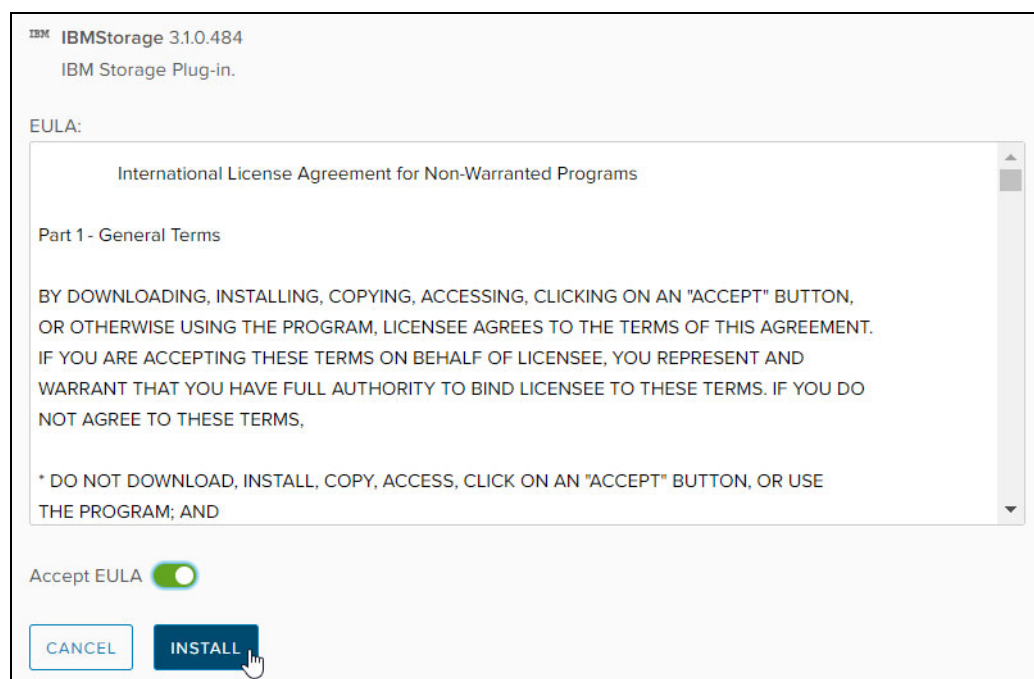


Figure 8-17 Configuration: Install Plug-In

17. Click **SAVE CHANGES** to save the plug-in, as shown in Figure 8-18.



Figure 8-18 Configuration: Save changes

18. Select **Home** → **Manage** → **Validate Configuration**, as depicted in Figure 8-19.

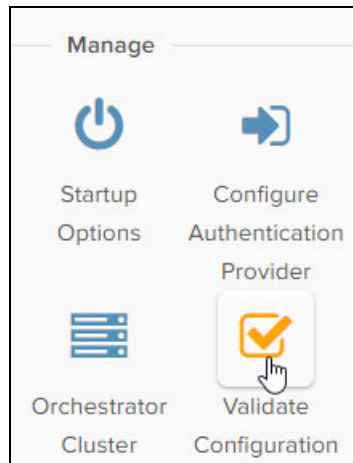


Figure 8-19 Configuration: Select Validate Configuration

19. The configuration validation is displayed, as illustrated in Figure 8-20.

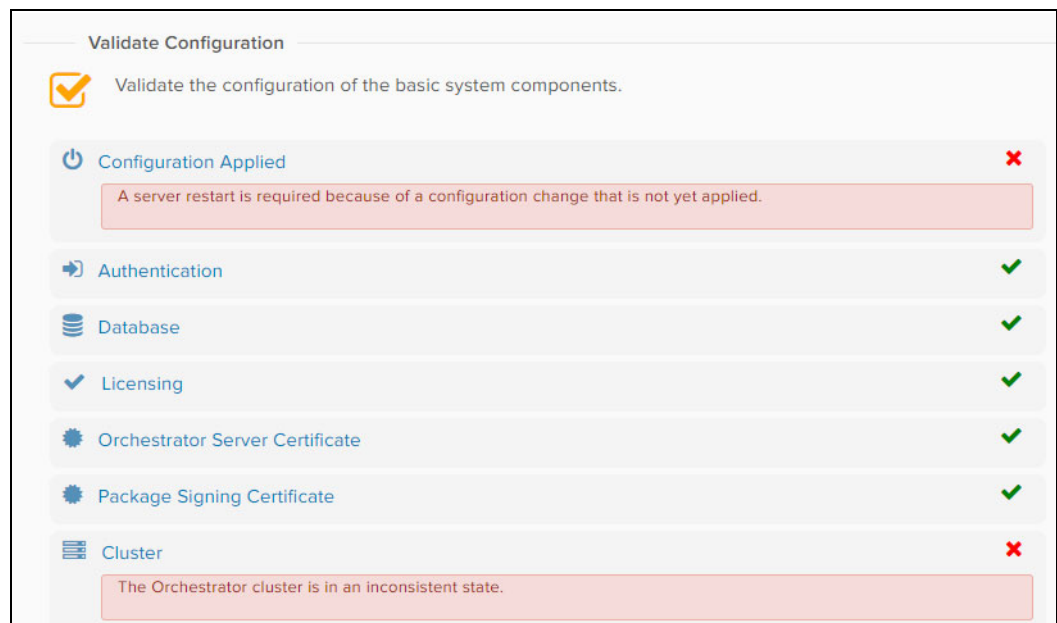


Figure 8-20 Configuration: Validate Configuration

20. Restart the vRO server VM to apply the changes.

21. Go to the vRO client interface:

<https://Orchestrator-IP-addresshttps:8281/vco/client/client.jnlp>

22. Select **Workflows** → **IBM** → **Configuration** → **Set Server and Token** → **Start Workflow**, as depicted in Figure 8-21.

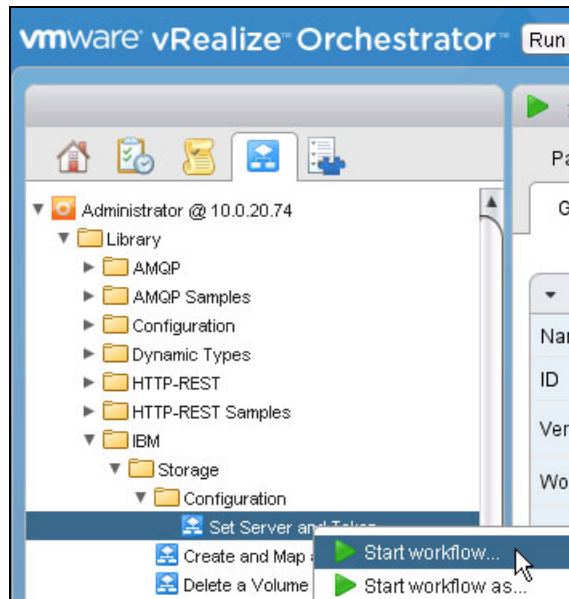


Figure 8-21 Select Set Server and Token

23. Define the IP-Address, Port, and use the Token from step 5 on page 113, and click **Submit**, as shown in Figure 8-22.

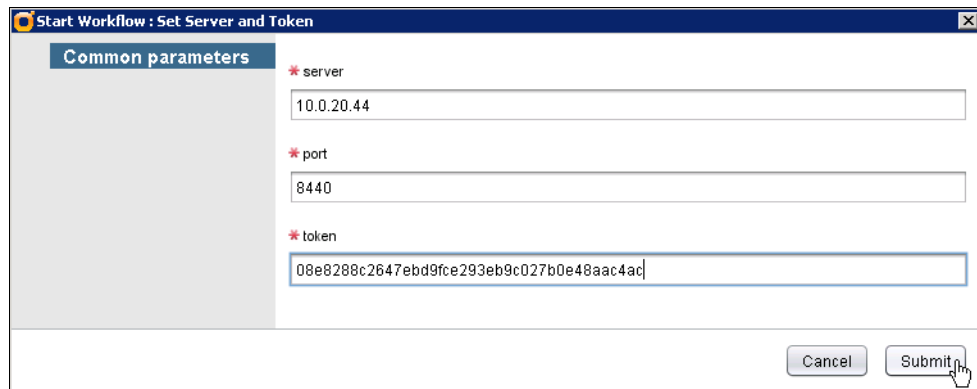


Figure 8-22 Set server and token

24. From the IBM Spectrum Connect web interface, you can now attach a service from the target storage system to be managed also by vRO or managed by vRO only. Click **vRO** and click the **Delegate** button on the service, as shown in Figure 8-23. For more details about the creation of services, see 6.1, “vSphere Web Client illustration” on page 88.



Figure 8-23 Attach service to vRO

25. Click **Ok** to confirm the delegation, as shown in Figure 8-24.




Figure 8-24 Confirmation of service delegation

26. After the installation and configuration finishes, verify the settings from the VMware vRO client, as shown in Figure 8-25.

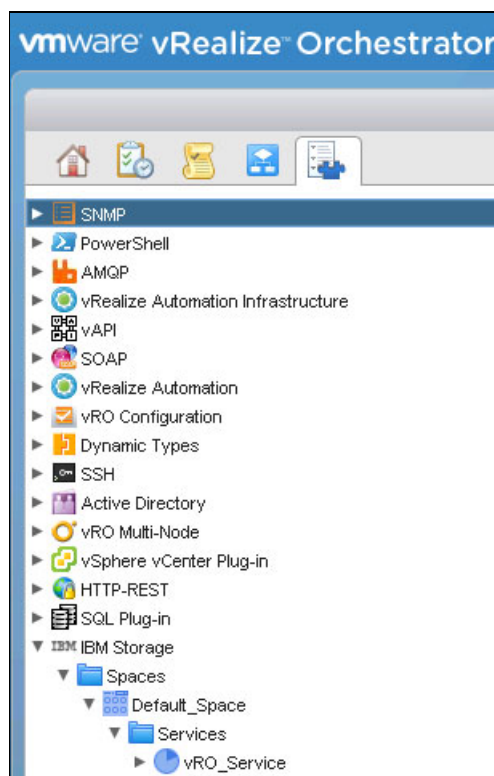


Figure 8-25 vRO client

8.2 Running workflows in vRO

There are five default workflows that come with IBM Spectrum Connect vRO integration:

- ▶ Create and Map a volume
- ▶ Delete a volume
- ▶ Extend a volume
- ▶ Map a volume
- ▶ Unmap a volume

Any of these workflows can be used to run the named operations. A more complex workflow can be created by using these workflows together, for example, “Create a volume” with “Map a volume” and “Add data store on iSCSI/FC/local SCSI” to allow creating, mapping, and preparing a volume for a VMware ESXi cluster.

Figure 8-26 shows a basic workflow example.

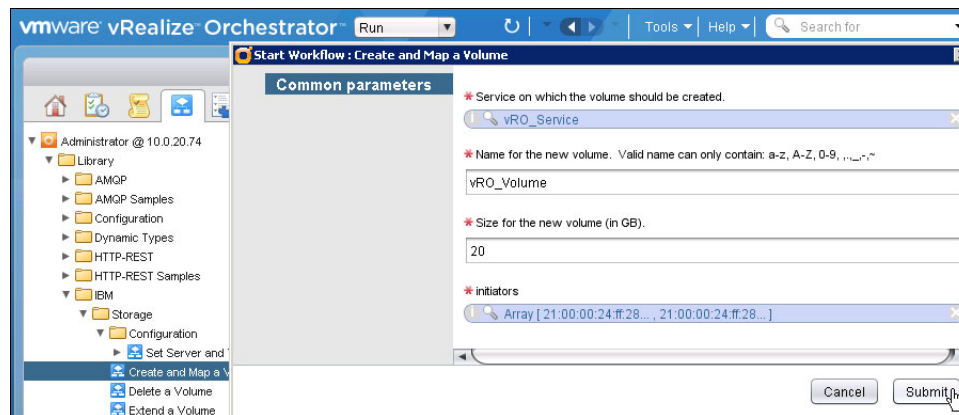


Figure 8-26 Running “Create and Map a volume workflow”



VMware vCenter Site Recovery Manager

This chapter presents the combined solution of the IBM Spectrum Accelerate family and VMware vCenter Site Recovery Manager (SRM) for business continuity and disaster recovery (DR). It describes how to deploy and operate the solution.

It covers the following topics:

- ▶ IBM Spectrum Accelerate family and VMware vCenter Site Recovery Manager
- ▶ Installing Spectrum Accelerate Family SRA for VMware vCenter SRM
- ▶ VMware vCenter Site Recovery Manager implementation and usage

9.1 IBM Spectrum Accelerate family and VMware vCenter Site Recovery Manager

Planning for all contingencies to meet essential business-driven objectives during events ranging from site-wide disasters to equipment outages constitutes one of the greatest challenges that are 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 spanning 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 in the absence of a comprehensive DR framework featuring powerful orchestration with turnkey execution.

With the combined solution of the IBM Spectrum Accelerate family and VMware vCenter SRM, the activities that are required for meeting business continuity objectives before, during, and after a critical event resulting in site-wide interruption of service are vastly simplified by using features such as:

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

In addition to streamlining classic DR operations, the combined solution can provide:

- ▶ Testing or data mining on existing replicated production data
- ▶ Backup at the DR site
- ▶ Planned data center migrations:
 - Disaster avoidance
 - Site maintenance

In summary, clients implementing VMware vCenter SRM with IBM Spectrum Accelerate family's advanced replication technology benefit from a comprehensive service-oriented framework jointly spanning DR and avoidance methodologies while realizing predictable, reliable, and efficient delivery of business continuity services and objectives.

The remainder of this section summarizes the solution components underpinning the DR and avoidance capabilities that are inherent to the joint IBM and VMware solutions. It assumes that the reader has a fundamental level of proficiency with IBM Spectrum Accelerate family, VMware, and general DR concepts and terminology.

Topics including solution design principles, planning considerations, and preferred practices are explored, in addition to a conceptual examination of several common usage cases. Reference materials covering the details of installation and configuration also follow.

9.1.1 Remote Mirroring overview

Fundamentally, the IBM Spectrum Accelerate family Remote Mirroring function maintains a real-time copy of consistent data. It does so by creating a persistent relationship between two XIV storage systems that are physically connected to the replication fabric. This relationship can be created by using either Fibre Channel (FC) or Internet Small Computer Systems Interface (iSCSI) links. These systems can be two IBM FlashSystem A9000 or IBM FlashSystem A9000R systems that are physically connected to the replication fabric by using either FC or iSCSI links, two IBM Spectrum Accelerate systems that use iSCSI links, or a XIV Storage System and a IBM Spectrum Accelerate system that use iSCSI links.

As an introduction to the rich features that are inherent to the remote mirroring technology, consider that the traditional core functions that are typical of remote mirroring solutions are augmented by the following unique capabilities with associated advantages:

- ▶ Both synchronous and asynchronous mirroring are supported on a single system.
- ▶ Mirroring is supported for consistency groups and individual volumes, and mirrored volumes might be dynamically moved into and out of mirrored consistency groups.
- ▶ Mirroring is data aware. Only actual data is replicated, except for IBM FlashSystem A9000 and IBM FlashSystem A9000R, where the replication layer is above the data reduction layer.
- ▶ Synchronous mirroring automatically resynchronizes couplings when a connection recovers after a network failure.
- ▶ Mirroring provides an option to automatically create subordinate volumes.
- ▶ IBM Spectrum Accelerate family allows user specification of initialization and resynchronization speed.

Furthermore, as an overview of the management capabilities that are facilitated through strong IBM Spectrum Accelerate family and VMware vCenter SRM integration, consider the comprehensive features that are available when performing the following DR common tasks:

- ▶ Create the configuration within the vSphere Web Client:
 - Add and remove storage systems from the VMware vCenter SRM configuration.
 - Create recovery plans for IBM Spectrum Accelerate family-based data stores.
 - Create and manage VMware vCenter SRM protection groups for IBM Spectrum Accelerate family-based data stores.
 - Enable, disable, and view the connectivity status.
 - Review mirroring status between volumes and consistency groups.
- ▶ Use VMware vCenter SRM to orchestrate the end-to-end workflow of failover and failback operations by harnessing storage system robust remote mirroring capabilities:
 - Fail the operation over to the recovery site by reversing mirroring, designating the Recovery LUNs as primary for updates, and mapping them to new Primary hosts. After restoring service at the original protected site, enable the reprotect capability to reinstate comprehensive protection of the data center that is running at the recovery site.
 - Failback by reverse mirroring from the recovery site storage systems to the protected site storage systems. Invoke reprotection again to restore the normal steady-state operation of VMware vCenter SRM with the protected and recovery sites fulfilling their standard roles.

- ▶ Test the DR plan (DRP):
 - Create and use snapshots of target mirror LUNs in failover testing without interrupting the replication.
 - Create backup LUN snapshot replication points before they become mirroring targets and are overwritten (applies to both failover and failback scenarios).
 - Perform cleanup (delete snapshots).
- ▶ Monitor and manage the storage systems:
 - Query storage system details and connectivity health status.
 - Detect and display paired arrays (mirrored relationships).

9.1.2 VMware vCenter Site Recovery Manager overview

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

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

VMware vCenter SRM uses the IBM Spectrum Accelerate family remote mirroring capabilities to create a copy of the data at a secondary location. Software for replicating data on the storage system is included with every system. Also included is the IBM Spectrum Accelerate family Storage Replication Adapter (SRA), which allows VMware vCenter SRM to suspend, snapshot, re-enable, and reverse replication on the IBM Spectrum Accelerate family systems.

VMware vCenter SRM 5.x disaster recovery and VMware vCenter SRM enhancements

VMware vCenter SRM 5.x enhances the ability to build, manage, and run reliable DRPs spanning bi-direction relocation of data center services, and thus provides unprecedented levels of protection. The sophistication of VMware vCenter SRM disaster recovery strategies have expanded through the addition of the following capabilities:

- ▶ Automated reprotection: The reprotection capability allows VMware vCenter SRM to take advantage of the advanced remote mirroring function to wrap the roles of the primary and DR sites following a failover, which continues data center protection while meeting RPOs and RTOs regardless of which site is operating in the role of the primary. In effect, the environment running at the recovery site uses the original primary data center to establish replication and protection of the environment back to the original protected site through a single click.
- ▶ Automated failback: After the reprotection process ensures that remote mirroring is reversed and data synchronization is then established at the original primary site, the automated failback capability can restore the roles of the two data centers to their original states, simultaneously restoring the state of the entire vSphere environment and maintaining full site-wide data protection operations. To accomplish this, failback runs the same workflow that was used to migrate the environment to the protected site, thus ensuring that the systems that are included in 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 (VMs) are restarted at the DR site by expanding the number of priority groups that are available to vSphere administrators and permitting VMware vCenter SRM to recognize VM dependencies within a priority group.
- ▶ Use Storage DRS and Storage vMotion on sites that VMware vCenter SRM protects in VMware vCenter SRM 5.1.
- ▶ Integration of the VMware vCenter SRM UI in vSphere Web Client in VMware vCenter SRM 5.8.
- ▶ An optional embedded vPostgreSQL database that you can use instead of a dedicated external database with minimal configuration in VMware vCenter SRM 5.8.

VMware vCenter SRM 6.0 enhancements

VMware vCenter SRM provides the following new features:

- ▶ Support for VMware vSphere 6.0, including integration with shared infrastructure components such as Platform Services Controller and vCenter Single Sign-On.
- ▶ Support for Storage vMotion and Storage DRS on both the protected and recovery sites.
- ▶ Protection and recovery of VMs in IPv6 environments.
- ▶ Internet Protocol (IP) customization enhancements to support dual-protocol IP configurations and independent IPv4 and IPv6 configurations.

VMware vCenter SRM 6.1 enhancements

VMware vCenter SRM provides the following new features:

- ▶ Support for VMware vSphere 6.0 update 1
- ▶ Storage policy based protection of virtual machines
- ▶ Support for stretched storage when using storage policy protection groups in enhanced linked mode
- ▶ Support for auto-mapping of stretched NSX networks
- ▶ Enhancements to mappings for test networks

VMware vCenter SRM 6.5 enhancements

VMware vCenter SRM provides the following new features:

- ▶ Integration with vRealize Operations Manager through a new management pack
- ▶ Support for VMware vSphere Virtual Volumes through vSphere Replication
- ▶ Support for silent installation, upgrade, and uninstallation
- ▶ Enhancements to Site Recovery Manager public API
- ▶ Enhancements and new workflows in the vRealize Orchestrator plug-in for Site Recovery Manager 6.5
- ▶ Support for the vCenter Server HA feature
- ▶ Support for migration of a vCenter Server installation on Windows to a vCenter Server Appliance installation during upgrade
- ▶ Support for the Virtual Machine Encryption feature
- ▶ Site Recovery Manager API support for Test Recovery operation when the protected and recovery sites are disconnected
- ▶ Participation in the VMware Customer Experience Improvement Program (CEIP)

Common disaster recovery planning terms and definitions

This section offers a brief review of the universal concepts that are inherent to DR and avoidance planning, which are presented by using terms and definitions that are specific to the joint VMware vCenter SRM and IBM Spectrum Accelerate family solution:

- ▶ **Site pairing:** Site pairing establishes the connection between two sites and ensures authentication. After this is done, the two sites can exchange information. This requires administrative privileges at both the sites.
- ▶ **Bidirectional operation:** You can use a single set of paired VMware vCenter SRM sites to provide protection in both directions. In this scenario, each site can simultaneously be a protected and a recovery site for different VMs and storage system volumes, also known as an Active/Active protection scheme. Bidirectional does not mean a single volume or VM is replicated in both directions at the same time; instead, bidirectional refers to the ability of the VMware vCenter SRM to fulfill the role of a protected site, a recovery site, or both simultaneously.
- ▶ **Protected and recovery sites:** In a typical installation, the protected site hosts the mission-critical services, and the recovery site is an alternative facility where these services can be migrated. Here again, Site A can be a protected site for some VMs and a recovery site for other VMs at Site B.
- ▶ **Protection groups:** A container for VMs and templates that use the same replicated data store group. Protection groups consist of pointers to the replicated vSphere data stores containing collections of VMs that are failed over from the protected site to the recovery site during actual DR or testing operations. Conceptually, VMware vCenter SRM protection groups specify the relationship of protected VMs in the same way that IBM Spectrum Accelerate family consistency groups specify recovery relationships among logical volumes.
- ▶ **Recovery plan:** A recovery plan specifies how the VMs in a specified set of protection groups are recovered.
- ▶ **SRA:** The storage system software that is required for VMware vCenter SRM to issue replication commands to the IBM Spectrum Accelerate family array. This software is included with the storage system and available for download on the VMware website. More details about which SRA version is supported for which VMware vCenter SRM version can be found at the following website:

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

- ▶ **Recovery point objective (RPO):** The RPO indicates how current the target replica needs to be relative to the source. The RPO 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.

The IBM Spectrum Accelerate family reports a mirror state as RPO OK or RPO Lagging for asynchronous mirroring. This status is determined by:

- The RPO parameter that is associated with the mirror.
- The time stamp of the master's current last_replicated snapshot.
- The current system time, which can be affected by time zone differences between the data centers.

The asynchronous mirror state is refreshed based on a system-defined schedule.

- ▶ **Recovery time objective (RTO):** Unlike RPO, which defines how much data is lost, RTO defines how much downtime is acceptable for a particular application or infrastructure. There are several options in the VMware vCenter SRM for assigning recovery priority and scheduling. This defines how VMs recover after a failover is initiated.

- ▶ Synchronous replication: Remote Mirroring can be a synchronous copy solution where write operations are completed on both copies (local and remote sites) before they are considered to be complete. This type of remote mirroring is normally used for short distances to minimize the effect of I/O delays that are inherent to the distance to the remote site. Synchronous replication ensures that the data volumes on both the source and target storage systems are exact mirrors. Typically, the host application might notice that the writes take longer to process due to the distance delays (latency) in sending the update to the secondary and receiving the response.
- ▶ Asynchronous replication: Remote Mirroring can also be an asynchronous solution where consistent sets of data are copied to the remote location at specified intervals and host I/O operations are complete after writing to the primary. This is typically used for long distances between sites. In asynchronous replication, the host update is acknowledged immediately and replication of the updates is sent later to the remote system. In this case, the host avoids the write latency inherent to synchronous mirror designs. Although asynchronous replication has performance benefits, it is a non-zero RPO, which means some data loss is acceptable when a failover is initiated. Long distances require asynchronous replication in most cases.
- ▶ Consistency group: A consistency group consists of volumes that share atomic point in time (RPO). When a recovery is issued at the remote site, all volumes recover at the same point. In VMware vCenter SRM installations, consistency groups must be used only if actual dependencies across storage system volumes exist (such as data stores with multiple extents).

9.1.3 Minimum IBM Spectrum Accelerate family and VMware vCenter SRM solution prerequisites

In a typical SRM installation, the protected site provides business-critical data center services, and 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 thousands of miles away or in the same site. In the typical case, the recovery site is in a facility that is unlikely to be affected by any environmental, infrastructure, or other disturbances that affect the protected site.

To build an SRM solution featuring the robust storage system remote mirroring technology, the vSphere and storage system configurations that are deployed at both sites must meet the following minimum hardware and software requirements.

The vSphere has these requirements:

- ▶ Each site must include at least one vSphere data center with the following components:
 - The vCenter server and VMware vCenter SRM server must be configured at both sites.
 - ESX hosts must exist at both the protected site and the recovery site.
 - At least one VM must be on a replicated data store at the protected site.
- ▶ An VMware vCenter SRM license must be installed with enough per-VM licenses to cover the systems that are protected at each site.

Databases must be installed and configured at each site to support vCenter Server and VMware vCenter SRM or the embedded databases must be used. Supported databases and related interoperability information can be found at the following website:

http://partnerweb.vmware.com/comp_guide2/sim/interop_matrix.php?

- ▶ Server hardware resources that can support the same VMs and associated workloads as the protected site must be provisioned at the DR site.

- ▶ A list of VMware vCenter SRM product documentation can be found at the following site:
https://www.vmware.com/support/pubs/srm_pubs.html

The SAN and Networking requirements are the following:

- ▶ Transmission Control Protocol (TCP) connectivity must be available between the VMware vCenter SRM servers and the vCenter servers.
- ▶ The protected and recovery sites must be connected by a reliable FC and TCP/IP network with sufficient bandwidth that is available to meet RPOs and RTOs in the case of asynchronous remote mirroring, and both application workload requirements and RTOs in the case of synchronous remote mirroring.
- ▶ The recovery site must have access to the same public and private networks as the protected site, although not necessarily the same range of network addresses.
- ▶ At both sites, applicable networking and domain name server (DNS) resources must be set up, configured, and tested before installing VMware vCenter SRM.
- ▶ Network hardware resources that can support the same VMs and associated workloads as the protected site must be provisioned at the DR site.

IBM Spectrum Accelerate family has these requirements:

- ▶ The logical volumes backing the data stores must be on the XIV, the IBM FlashSystem A9000 or IBM FlashSystem A9000R, or on the IBM Spectrum Accelerate systems, or on a XIV Storage System and a IBM Spectrum Accelerate system, at both the protected site and the recovery site.
- ▶ The Spectrum Accelerate family SRA software must be downloaded from the VMware website and then installed and configured.
- ▶ Both storage systems must be physically and logically configured for remote mirroring. For more information about this topic, see *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401.
- ▶ Storage system and SAN hardware resources that can support the same VMs and associated workloads as the protected site must be provisioned at the DR site.
- ▶ Sufficient hard capacity must be available on the storage system at each site as dictated by the specific implementation and must include a relatively small amount of capacity that is dedicated to *placeholder data stores*, which are used by VMware vCenter SRM to store VM placeholder files, each roughly 1 KB.

The conceptual diagram that is presented in Figure 9-1 offers a comprehensive view of the topology of the VMware vCenter SRM environment harnessing IBM Spectrum Accelerate family Remote Mirroring with all essential components and their relationships.

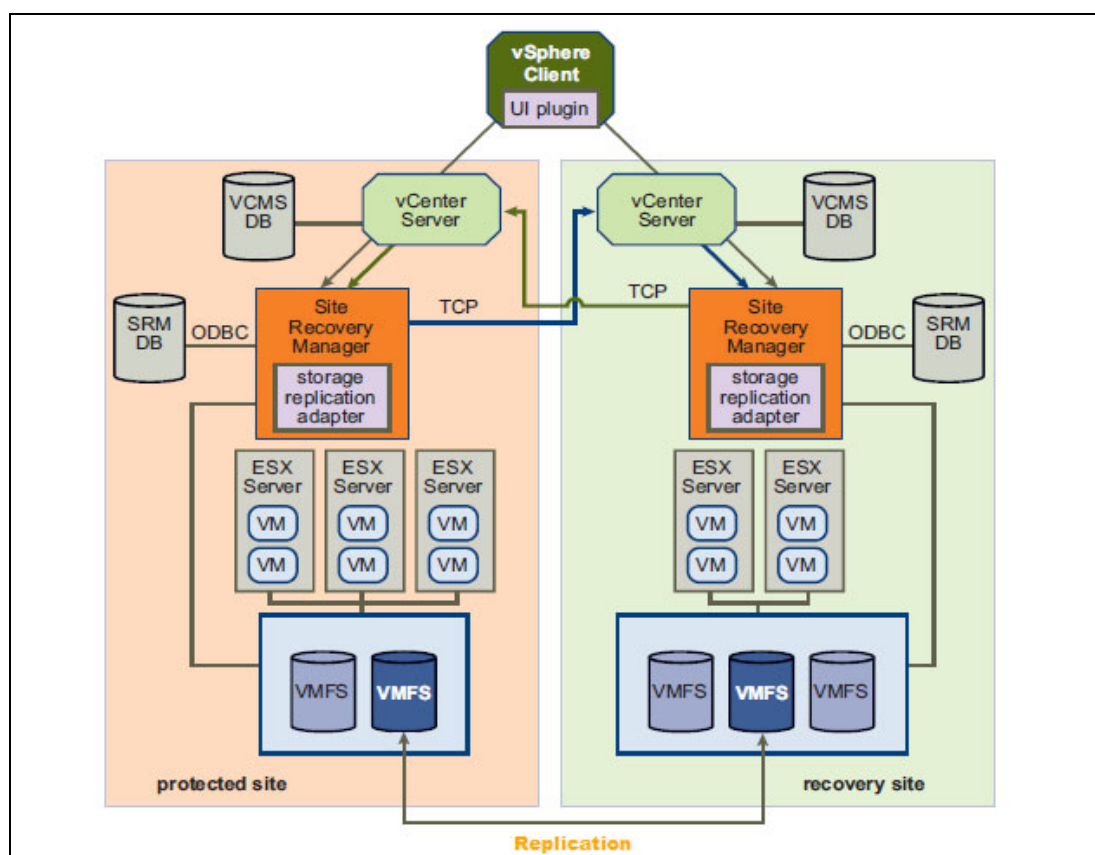


Figure 9-1 IBM Spectrum Accelerate family and VMware vCenter SRM environment with minimum required components

9.1.4 VMware vCenter SRM integration with the IBM Spectrum Accelerate Family Storage Replication Adapter

As shown in Figure 9-1, the Spectrum Accelerate family SRA functions as an interface between the storage system and VMware vCenter SRM by converting standardized commands that are generated by VMware vCenter SRM into IBM Spectrum Accelerate family-specific commands. For example, SRA enables VMware vCenter SRM commands that encompass execution of the vSphere-centric tasks and workflows, including querying replicated data stores and promoting replicated data stores, to proceed in concert with IBM Spectrum Accelerate family-specific functions, including remote mirroring and snapshot management:

- ▶ Discovering LUNS and their associations in the context of the remote mirroring relationships spanning storage systems at both the primary and DR sites.
- ▶ Running test failover, which is used to test the implementation of the planned DR workflows by invoking snapshots to create copies of the data stores without impacting operations at either the primary or DR environments.

- ▶ Automating the failover of a storage system at the primary VMware vCenter SRM site to a storage system at a recovery (secondary) VMware vCenter SRM site.

Immediately upon a failover, the ESXi servers at the secondary VMware vCenter SRM site start using the replicated data stores on the mirrored volumes of the secondary storage system.

- ▶ Invoking reprotect for either the entirety of the VMware environment or a subset.
- ▶ Running a failback following a previously completed failover and reprotect.

In summary, the Spectrum Accelerate family extends VMware vCenter SRM capabilities and allows it to seamlessly employ replication and mirroring as part of the VMware vCenter SRM comprehensive disaster recovery planning (DRP) solution.

At the time of writing, the current release of Spectrum Accelerate family SRA (version 3.0.0) supports the following versions of VMware vCenter SRM server:

- ▶ Version 6.1
- ▶ Version 6.5

For earlier versions of Site Recovery Manager, use XIV SRA 2.3.0.

9.1.5 VMware vCenter Site Recovery Manager operations

VMware vCenter SRM is tightly integrated with VMware vSphere. VSphere administrators use the product to initiate automated failover from a primary (protected) site to a secondary (recovery) site. Starting with VMware v5.0, VMware vCenter SRM also automates failback to the primary site.

Figure 9-2 illustrates a high-level view of the remote mirroring capability that is orchestrated by using key vCenter integration components to provide comprehensive vSphere data center protection.

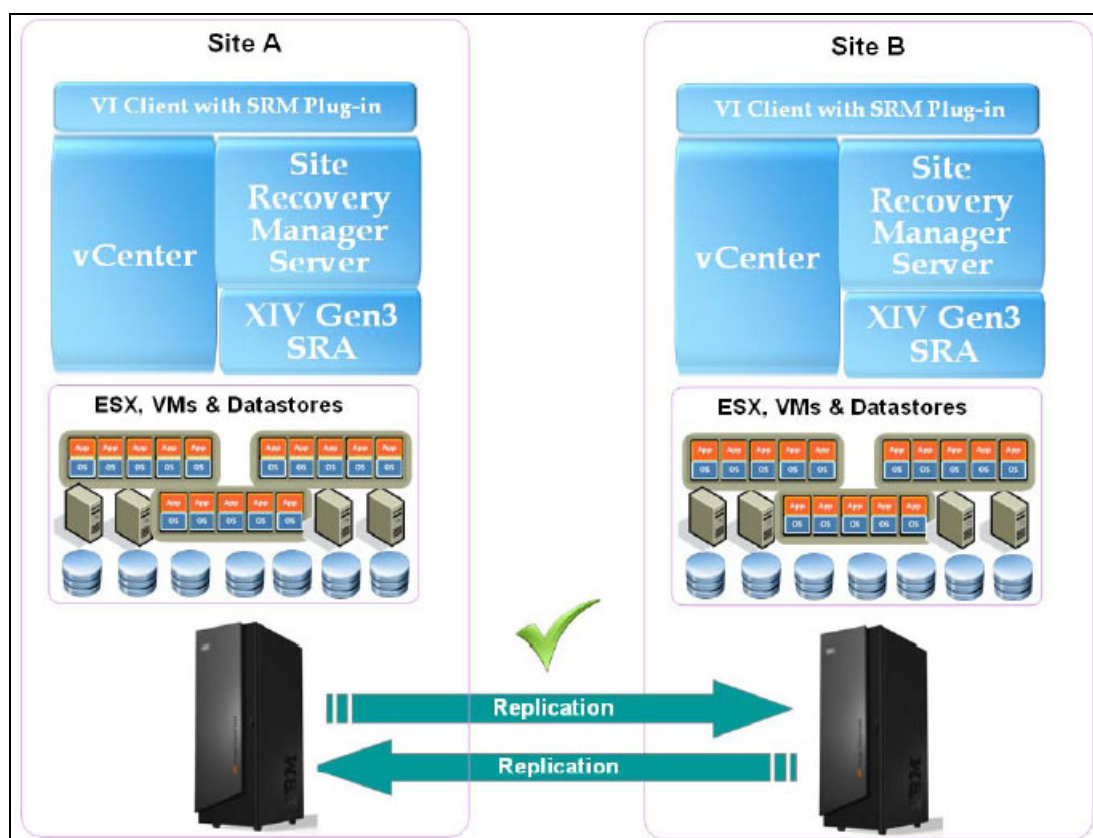


Figure 9-2 Remote mirroring and integration components for vSphere data center protection

Continuous protection

In normal production, the VMs run on ESX hosts and storage devices in the primary data center. Additional ESX servers and storage devices are on stand by in the backup data center. As a brief overview, the following common configuration steps are necessary to implement disaster recovery plans:

- ▶ Define protection groups of associated VMs that must be recovered together, for example Infrastructure (Active Directory or DNS), Mission Critical, and Business Critical.
- ▶ Define actions before a test or failover as the recovery site, such as cloning or suspending low-priority VMs to free recovery resources at the recovery site.
- ▶ Define the allocation of resources and any networking changes that are required by the VMs.
- ▶ Build call-outs that cause test or failover process to pause and present instructions to the administrator, or specify scripts in the recovery process.
- ▶ Identify finite values of time or specific numbers of heartbeats to wait for VMs to respond after their power-on process is complete.

VMware vCenter SRM does not automatically trigger the failover of a recovery plan; instead, human intervention is required to evaluate and declare a disaster. Following the initiation of storage remote mirroring failover processes, VMware vCenter SRM runs the earlier steps and automatically manages the mapping of compute resources, network resources, and recovery of VMFS, associated data stores, and ultimately VM and associated applications to the recovery site.

During failover

The remote mirroring function maintains a copy of the data with the necessary scope of consistency and currency to meet the predefined RPO on the storage system at the DR location.

If a failover process is triggered, all VMs shut down at the primary site if still possible or required. They are restarted on the ESX hosts at the backup data center, accessing the data on the backup storage system.

VMware vCenter SRM servers coordinate the operations of the replicated storage systems and vCenter servers at two sites with the following sequence of steps:

1. Shuts down the protected VMs if there is still connectivity between the sites and they are online.

When VMs at the protected site are shut down, VMs at the recovery site start. The VMs then access the data that was previously replicated from the protected site to assume responsibility for providing the same services.

2. Commands remote replication to synchronize any final data changes between sites.
3. Suspends data replication on the storage system and the storage system provides read/write replica devices that are mapped to vSphere at the recovery site.
4. Rescans the ESX servers at the recovery site to find devices and mounts the data stores.
5. Registers the replicated VMs.

With the introduction of *Enhanced Dependency Definition* in VMware vCenter SRM 5.X, the transfer of services from one site to the other is controlled by a recovery plan that specifies the order in which VMs are shut down and started, and the allocation of host and network resources that might be accessed.

6. Completes the power-up of replicated protected VMs in accordance with the recovery plan.

VMware vCenter SRM can automatically perform all these steps and fail over complete virtual environments with one click. This process saves time, eliminates user errors, and provides a detailed documentation of the disaster recovery plan.

After failover

Following a failover operation and the subsequent restoration of the original protected site, site protection can be reinstated by enacting the reversal of site mirroring roles in the VMware vCenter SRM 5.x by using the Reprotect option.

The Reprotect option helps to communicate with the Spectrum Accelerate family SRA to reverse the direction of replication. Protection groups now are replicated from the recovery site to the original primary site, as shown in Figure 9-3.

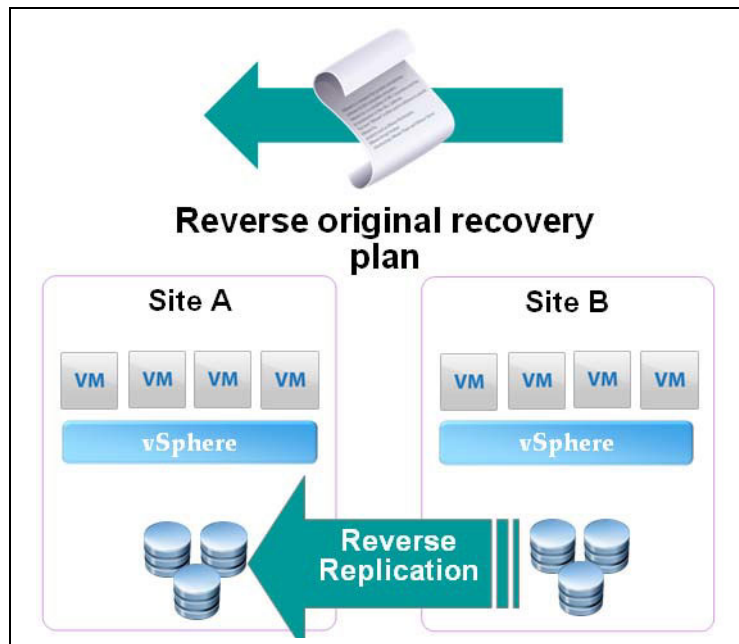


Figure 9-3 VMware vCenter SRM failback

During the reversal, IBM Spectrum Accelerate family system switches the roles of the volumes and replication is reversed. When this is done, only the changes between the volumes are replicated. An entire resynchronization of the volume is not necessary unless the original volume was lost in a disaster event.

After the data is synchronized to the original site, you can fail over and then select the **Reprotect** option again to restore the original operational configuration, as shown in Figure 9-4.

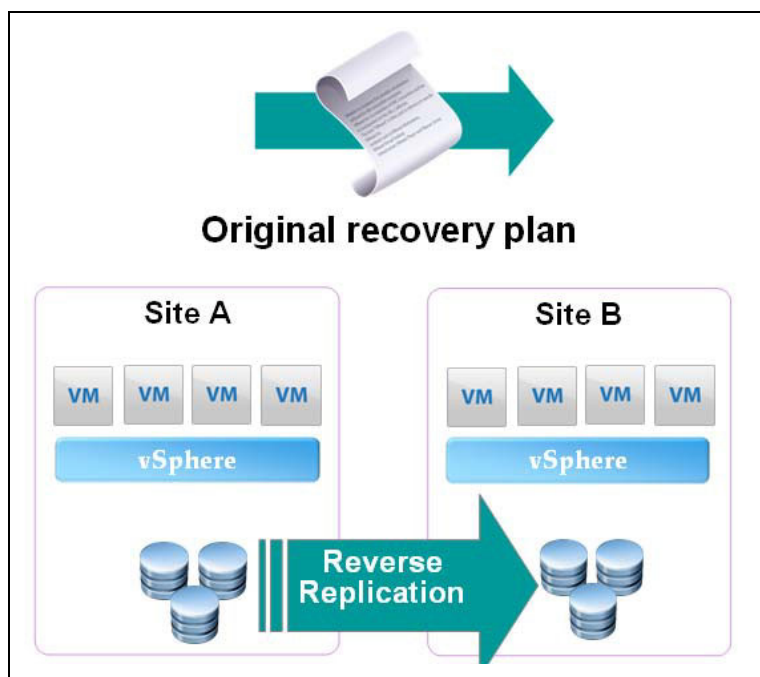


Figure 9-4 VMware vCenter SRM final reprotect

Testing failover

VMware vCenter SRM can also perform a test of the failover plan by creating an additional copy of the data at the backup site and starting the VMs from this copy without connecting them to any network. This feature allows vSphere administrators to test recovery plans without interfering with the production environment.

In addition, DR scripts can be tested by using storage system high performance, point-in-time, and incremental snapshots at the remote site that occur transparently to the ongoing data replication processes. These tests provide a deep understanding and functional assessment of the viability of achieving business-driven RPOs and RTOs, and provide an invaluable opportunity to pre-tune performance characteristics of hardware resources and address gaps in business continuance strategies before an actual disaster.

Figure 9-5 shows the fundamental sequence of steps that are invoked by VMware vCenter SRM testing capability, and demonstrates the operational synergy of VMware and the IBM Spectrum Accelerate family snapshot technology that is necessary for highly efficient and transparent disaster recovery plan testing.

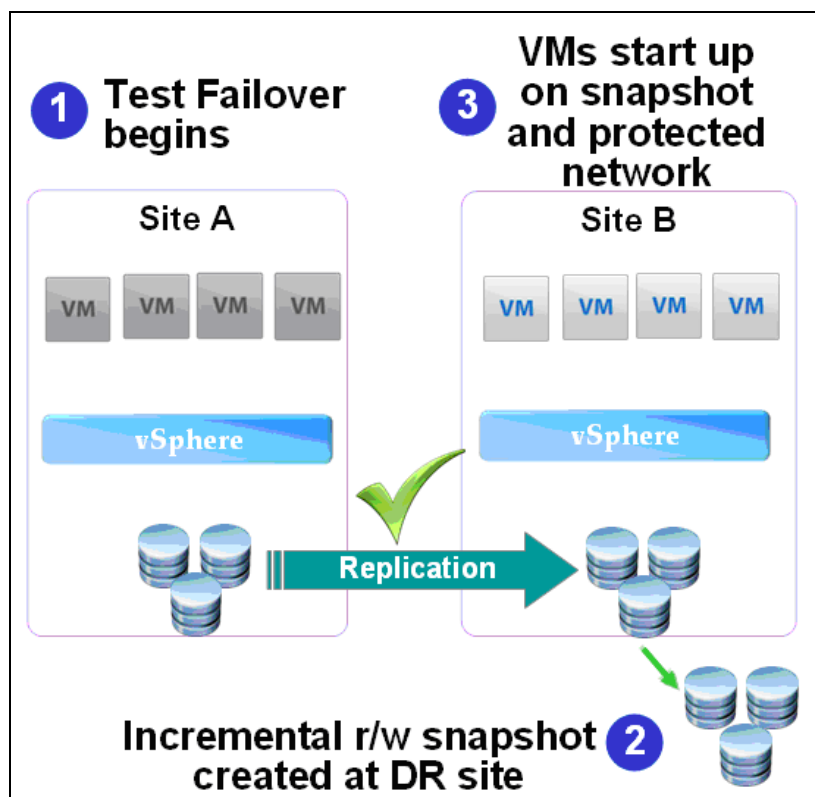


Figure 9-5 VMware vCenter SRM test failover plan leveraging snapshots

For the purposes of testing the DRP, VMware vCenter SRM performs the following tasks:

- ▶ Creates a test environment, including a network infrastructure.
- ▶ Rescans the ESX servers at the recovery site to find iSCSI and FC devices, and mounts the Spectrum Accelerate family snapshot.
- ▶ Registers the replicated VMs.
- ▶ Suspends nonessential VMs (if specified) at the recovery site to free resources.
- ▶ Completes the power-up of replicated, protected VMs in accordance with the recovery plan.
- ▶ Automatically deletes temporary files and resets storage configuration in preparation for a failover or for the next scheduled VMware vCenter SRM test.
- ▶ Provides a report of the test results.

High-performance snapshots on the storage system open additional usage cases during test failovers. These can include the following:

- ▶ Backup at the remote site
- ▶ Data mining
- ▶ Data analytics
- ▶ Test and development

VMware vCenter SRM Disaster Recovery plan configuration

This section offers a step-by-step guide that illustrates the flexible and comprehensive VMware vCenter SRM DR features as they are implemented in the VMware vCenter SRM utility within vSphere Web Client. To this end, the functional attributes of creating an active/passive disaster recovery plan and subsequent invocation of the VMware vCenter SRM testing, recovery, and reprotect capabilities are presented.

For an overview of the physical and logical configuration steps that are necessary to deploy Spectrum Accelerate family systems in a dual-site DR topology, see *VMware vCenter Site Recovery Manager Version 5.x guidelines for IBM XIV Gen3 Storage System*, found at:

<https://ibm.biz/Bds83Z>

This document includes guidance that addresses the configuration of connectivity, performance tuning settings, and provisioning mirrored logical volumes and associated consistency groups.

VMware vCenter SRM setup preferred practices

In addition to performing installation and setup with the proper order of operations, it is also important to follow the preferred practices when initially creating the VMware vCenter SRM configuration.

The VMware vCenter SRM 5.x guidelines for an IBM XIV Gen3 Storage System are the following:

- ▶ Specify a non-replicated data store for swap files.
This avoids unwanted bandwidth consumption and improves the recovery time as vCenter does not have to remove all the swap files from the VMs during recovery.
- ▶ Install VMware tools on all the VMs that are participating in a protection group.
Many recovery operations depend on proper VMware tools installation.
- ▶ Configure the VM dependencies across priority groups instead of setting it per VM.
This action ensures that VMs are started in parallel. The XIV Gen3 Storage System is optimized for parallel workloads so that this greatly improves performance.

For a comprehensive guide to VMware vCenter SRM 5.0 implementation preferred practices, see *VMware vCenter Site Recovery Manager 5.0 Performance and Best Practices*, found at:

<http://www.vmware.com/files/pdf/techpaper/srm5-perf.pdf>

9.2 Installing Spectrum Accelerate Family SRA for VMware vCenter SRM

This section provides the steps for installing the Spectrum Accelerate family SRA for VMware vCenter SRM server Version 6.5 under Microsoft Windows Server 2008 R2 Enterprise. Download and install Spectrum Accelerate family SRA for VMware on each VMware vCenter SRM server in your business continuity and DR solution by completing the following steps:

1. Find the Spectrum Accelerate family SRA installation file. For VMware vCenter SRM 6.5.0, the VMware vCenter SRM and SRA installation files are found at the following website:
https://my.vmware.com/web/vmware/details?downloadGroup=SRM_SRA65&productId=526
2. Run the installation file on a Protected Site VMware vCenter SRM server.

3. The IBM Spectrum Accelerate family SRA installation wizard opens, as shown in Figure 9-6. Click **Next**.

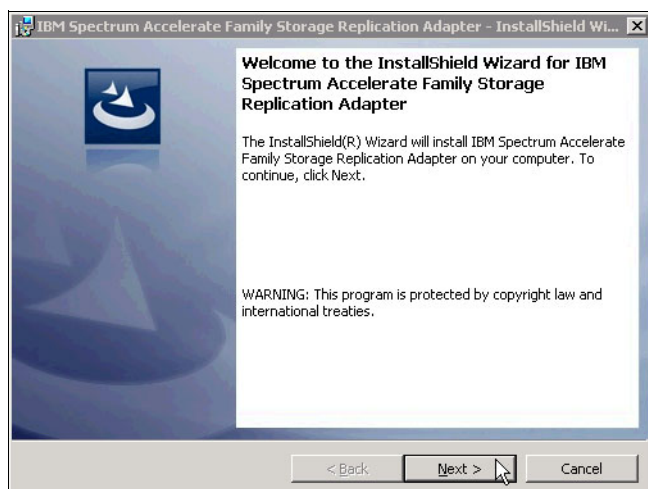


Figure 9-6 Welcome to the SRA installation wizard window

4. Follow the wizard guidelines to complete the installation.
5. Run the installation file on the Recovery Site VMware vCenter SRM server and repeat steps 3 and 4.

Tip: For step-by-step instructions that contain detailed installation guidance, error message definitions and resolutions, and basic configuration practices for the IBM Spectrum Accelerate family SRA for VMware Version 3.0.0, see IBM Knowledge Center: https://www.ibm.com/support/knowledgecenter/SSEQ8L_3.0.0/PDFs/IBM_SAF_SRA_3.0.0_UG.pdf?view=kc

9.2.1 Configuring the storage system for VMware vCenter SRM

Make sure that all VMs that you plan to protect are on IBM Spectrum Accelerate family volumes. If there are any virtual systems that are not on IBM Spectrum Accelerate family system, move them by completing the following steps:

1. Create volumes on the storage system.
2. Add the data store to the ESXi server.
3. Migrate or clone that VM to move it to the IBM Spectrum Accelerate family volumes.

For more information about connecting ESXi hosts to the storage system, see Chapter 3, “Attaching VMware ESXi” on page 25.

Create a storage pool on the storage system at the recovery site. The new storage pool contains the replicas of the ESXi host data stores that are associated with VMs that you plan to protect.

Remember: Configure a snapshot size of at least 20 percent of the total size of the recovery volumes in the pool. For testing failover operations that can last several days, increase the snapshot size to half the size of the recovery volumes in the pool. For longer-term or I/O intensive tests, the snapshot size might have to be the same size as the recovery volumes in the pool.

For information about XIV Storage System mirroring, see *IBM XIV Storage System Business Continuity Functions*, SG24-7759. For more information about IBM FlashSystem A9000 and IBM FlashSystem A9000R mirroring, see *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401.

At least one VM for the protected site must be stored on the replicated volume before you can start configuring the VMware vCenter SRM server and SRA. In addition, avoid replicating swap and paging files.

Start the mirroring for the volumes or consistency group that hold the data store.

The remainder of this chapter assumes that these prerequisite steps were completed, resulting in the state of the protected volume appearing similar to the volume *ITSO_Mirror* that is shown in Figure 9-7. To get to this view select **REMOTE VIEWS** → **Mirrored/HyperSwap Volumes (Availability)** and filter for your volumes in HyperScale Manager UI.

Volume	System	Availability Role	Host
ITSO_Mirror	A9000	Primary	WS_ESX_7
ITSO_Mirror	A9000R	Secondary	0

Figure 9-7 Mirrored volumes

9.3 VMware vCenter Site Recovery Manager implementation and usage

Implementing VMware vCenter SRM effectively requires the following sequence of steps in addition to the optional but preferred step of testing the recovery plans:

1. Connect the VMware vCenter SRM instances at both sites by specifying the IP addresses and authentication credentials of the vCenter Servers.
2. Set up “Inventory Mappings” between sites:
 - a. *Resource mappings* specify the resources that are recovered for protected VMs.
 - b. *Folder mappings* correlate the folders between the sites.
 - c. *Network mappings* link the vSphere data center networks of the two sites.
3. Assign *placeholder data stores* at each site.
 These data stores serve as repositories for small VM files that function as placeholders, which can be registered and activated when VMware vCenter SRM restores operations at the recovery site.
4. Specify alarms and permissions as required.
5. Add and configure *array managers*.
6. Create *protection groups*.
7. Create a *recovery plan*.
8. Test and clean up *recovery plans*.
9. Start *recovery*.
10. Start *reprotect*.

The remainder of this section focuses primarily on the final five steps because they are specific to the implementation of VMware vCenter SRM DR planning and testing with the IBM Spectrum Accelerate family. The following procedures use VMware vCenter SRM 6.5.1 and vSphere Web Client.

9.3.1 Connecting the sites

To configure the VMware vCenter SRM server for the protected and recovery sites, complete the following steps:

1. Run the vSphere Client and connect to the vCenter server.
2. In the vSphere Web Client Home window, click **Site Recovery**, as shown in Figure 9-8.



Figure 9-8 Inventory: Site Recovery

3. The Site Recovery Getting Started tab opens. Select **Sites**, as shown in Figure 9-9.

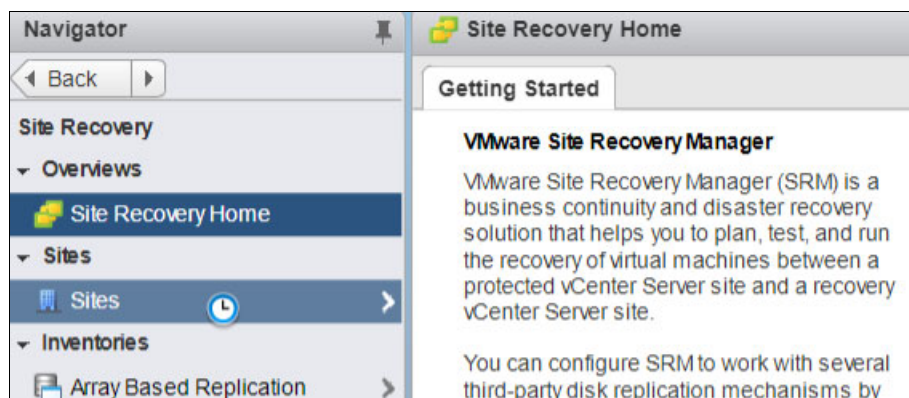


Figure 9-9 Getting Started with VMware vCenter Site Recovery Manager: Select Sites

4. Right-click the local site and select **Pair Site**, as shown in Figure 9-10.




Figure 9-10 VMware vCenter SRM Sites: Pair Site

5. In the Select Site window that opens, enter the Remote Site Information by specifying the IP address and communication port of the Platform Services Controller at the recovery site, as shown in Figure 9-11. Proceed by clicking **Next**.

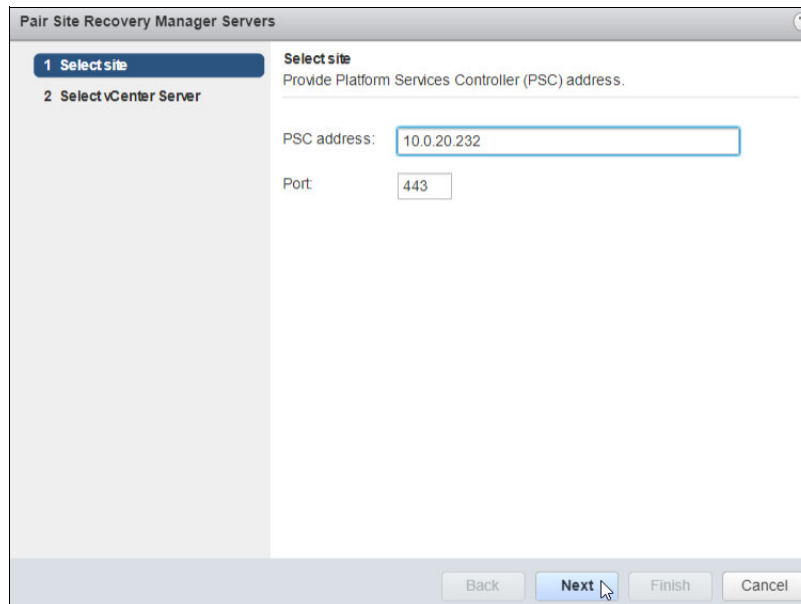


Figure 9-11 VMware vCenter SRM Pair Site: Select Site

6. In the Select vCenter Server window that opens, select the vCenter Server, complete the user name and password, and click **Finish**, as shown in Figure 9-12.

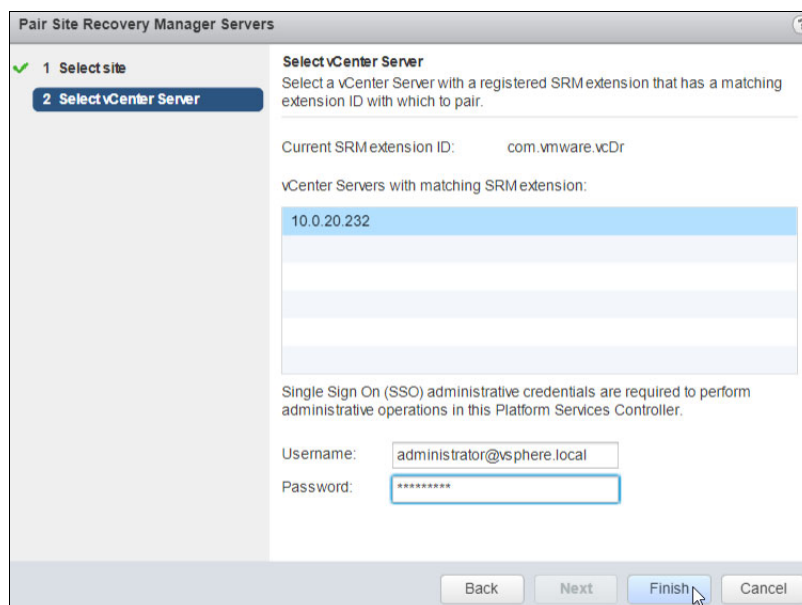


Figure 9-12 VMware vCenter SRM Pair Site: Select vCenter Server

7. Two Security Alerts open, one for each site, as shown in Figure 9-13. Ignore the Security Alerts and click **Yes**.

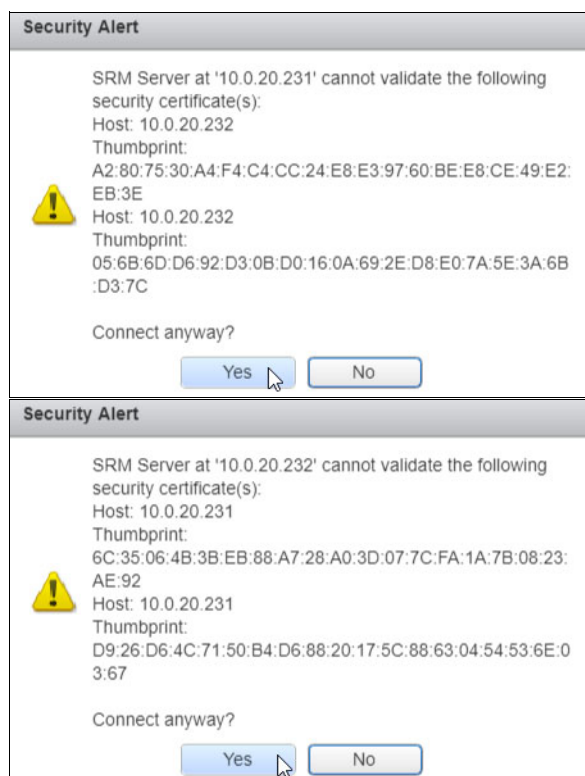


Figure 9-13 VMware vCenter SRM Server Certificate Security Alerts

8. Click the remote site. A window opens. Fill in the user name and password for the remote vCenter server. Click **Login** to proceed, as shown in Figure 9-14.




Figure 9-14 VMware vCenter SRM Login to the remote site

9.3.2 Setting up inventory mappings

Inventory mappings specify locations and networks for the vCenter server to use when placeholder VMs are initially created at the recovery site, effectively defining how VMware vCenter SRM maps resources at the protected site to resources at the recovery site. This phase of the VMware vCenter SRM implementation details the process of making inventory mappings for both the protected and recovery sites.

Complete the following steps:

1. Right-click the protected site and select **New Network Mapping**, as shown in Figure 9-15.

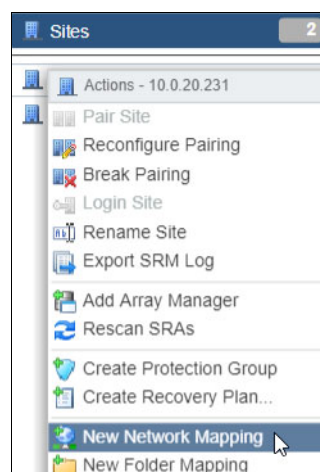


Figure 9-15 VMware vCenter SRM Configuration: New Network Mapping

2. Select **Automatically prepare mappings for networks with matching names** and click **Next**, as shown in Figure 9-16.

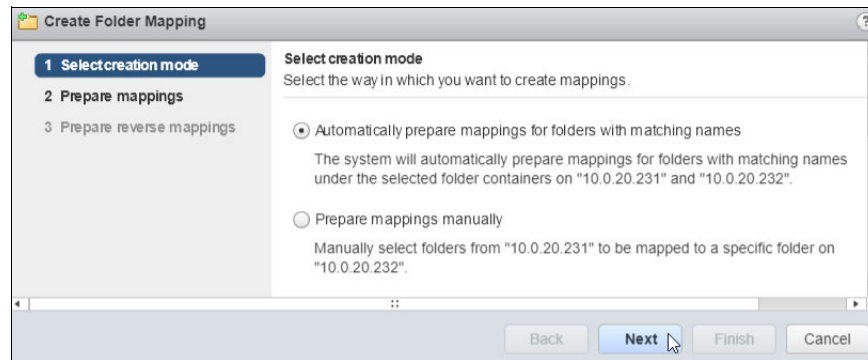


Figure 9-16 VMware vCenter SRM Network Mapping: Select creation mode

3. Select the protected site data center and recovery site data center, click **Add mappings**, and click **Next**, as shown in Figure 9-17.




Figure 9-17 VMware vCenter SRM Network Mapping: Prepare mappings

4. Select test network and click **Next**, as depicted in Figure 9-18.

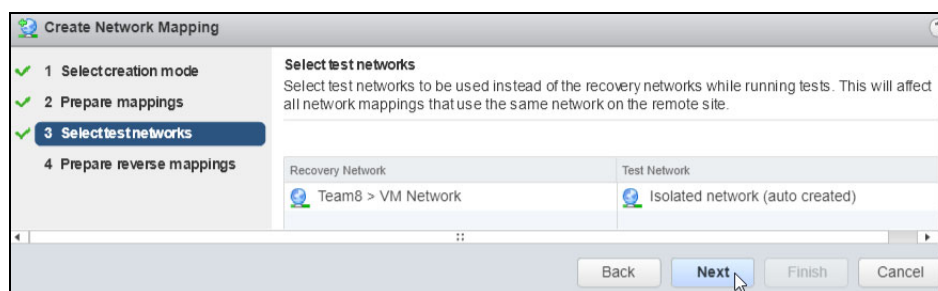


Figure 9-18 VMware vCenter SRM Network Mapping: Test Network

5. Check the reverse mapping and click **Finish**, as shown in Figure 9-19.

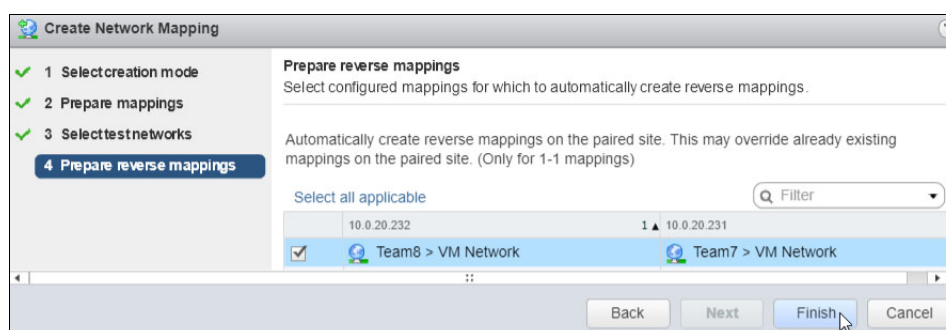


Figure 9-19 VMware vCenter SRM Network Mapping: Prepare reverse mappings

6. Right-click the protected site and select **New Folder Mapping**, as shown in Figure 9-20.

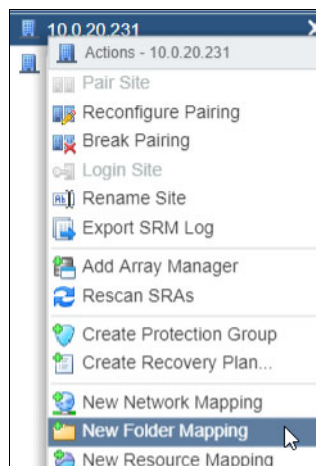


Figure 9-20 VMware vCenter SRM Configuration: New Folder Mapping

7. Select **Automatically prepare mappings for folders with matching names** and click **Next**, as shown in Figure 9-21.

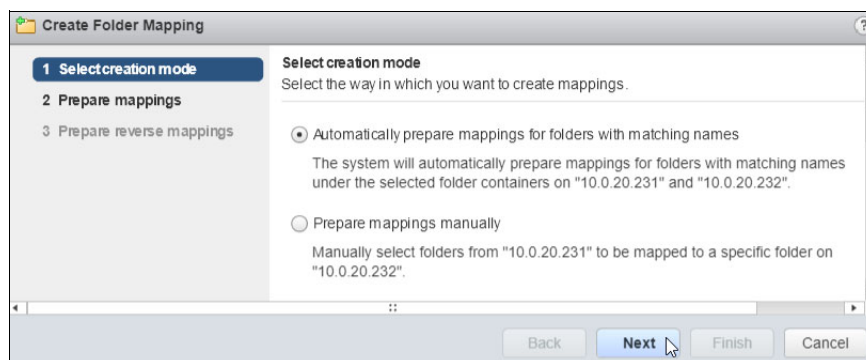


Figure 9-21 VMware vCenter SRM Folder Mapping: Select creation mode

8. Select the protected site data center and recovery site data center, click **Add mappings**, and click **Next**, as shown in Figure 9-22.

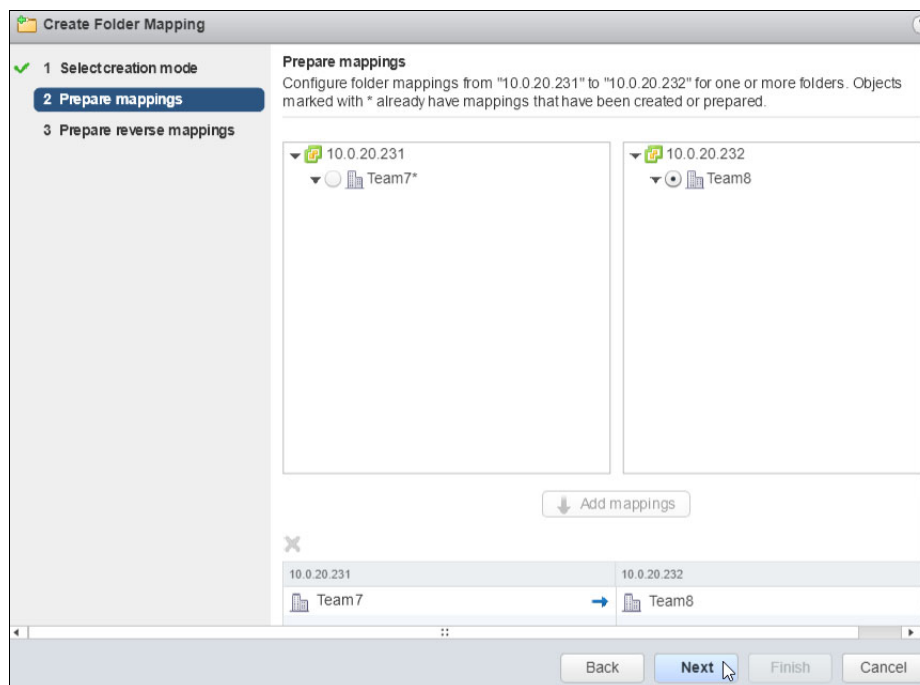


Figure 9-22 VMware vCenter SRM Folder Mapping: Prepare mappings

9. Check the reverse mapping and click **Finish**, as shown in Figure 9-23.

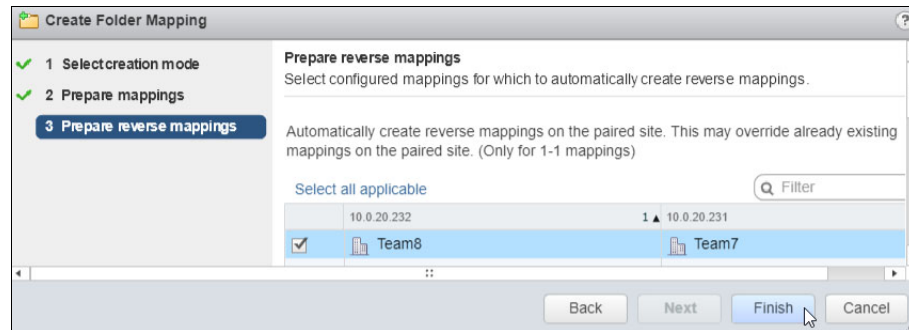


Figure 9-23 VMware vCenter SRM Folder Mapping: Prepare reverse mappings

10. Right-click the protected site and select **New Resource Mapping**, as shown in Figure 9-24.

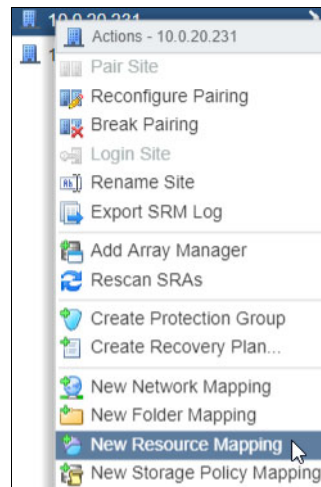


Figure 9-24 VMware vCenter SRM Configuration: New Resource Mapping

11. Select the protected site ESXi and recovery site ESXi server, click **Add mappings**, and click **Next**, as shown in Figure 9-25.

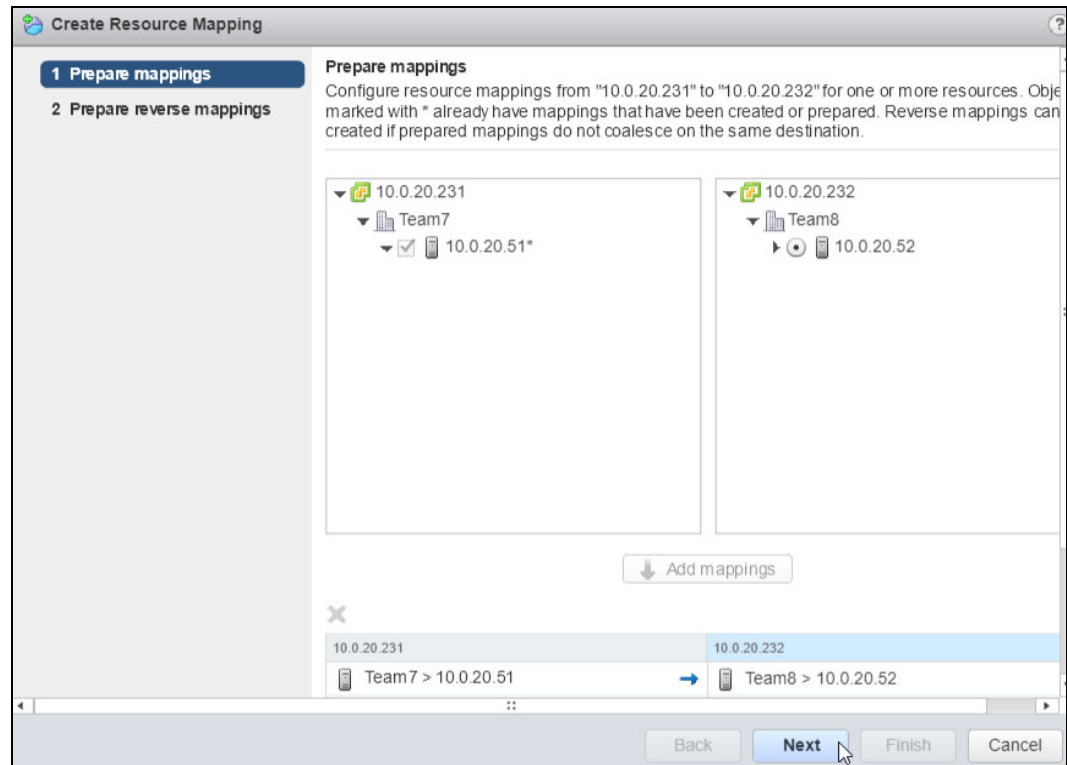


Figure 9-25 VMware vCenter SRM Resource Mapping: Prepare mappings

12. Check the reverse mapping and click **Finish**, as shown in Figure 9-26.

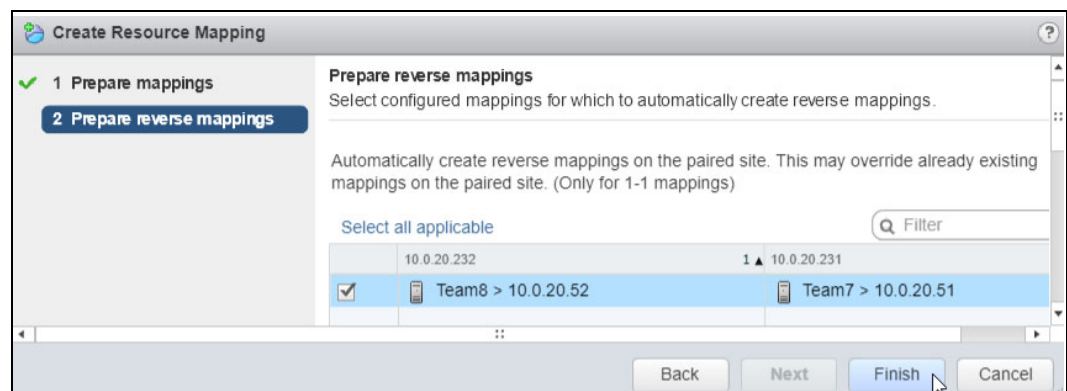


Figure 9-26 VMware vCenter SRM Resource Mapping: Prepare reverse mappings

9.3.3 Configuring placeholder data stores

Placeholder data stores must be assigned at both sites to function as repositories for a subset of metadata files that are associated with the protected VMs that are used to register them in vCenter inventory following an actual failover, failback, or planned migration.

Placeholder data stores can be relatively trivial in size, so for most applications, backing them with a storage system logical volume with the minimum allocation size is probably sufficient.

Placeholder data stores can be placed on local ESXi server data stores as well. These logical volumes and placeholder data stores must be created and configured in the vSphere client before completing the following steps that are necessary to define them in the VMware vCenter SRM configuration:

1. Right-click the protected site and select **Configure Placeholder Datastore**, as shown in Figure 9-27.

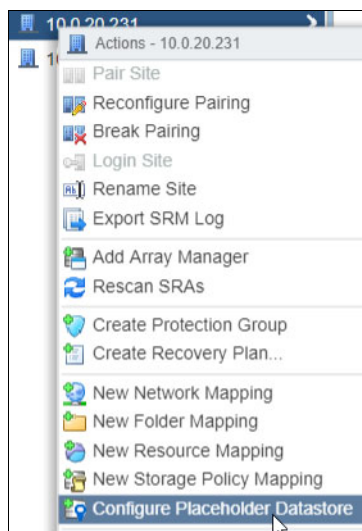


Figure 9-27 VMware vCenter SRM Configuration: Configure Placeholder Datastore

2. Select the local data store and click **OK**, as shown in Figure 9-28.

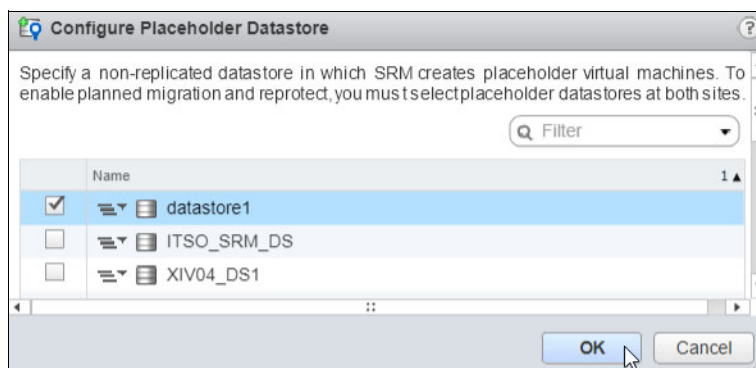


Figure 9-28 VMware vCenter SRM: Configure Placeholder Datastore

3. Repeat the previous two steps, but this time for the recovery site.

9.3.4 Adding and configuring array managers

In order for Spectrum Accelerate family SRA to start the remote mirroring processes that underpins the data center site protection workflows in VMware vCenter SRM, you must designate the storage subsystems at each site that are managed by SRA. To configure the SRA within the vSphere Web Client, make sure that the SRA is registered with the VMware vCenter SRM servers on both sites. Furthermore, make sure that none of the names of both storage systems have changed after the mirroring targets have been defined on FlashSystem A9000 and A9000R. Otherwise both arrays will not pair in SRM. The name of the target system can be changed by using HyperScale Manager or by using the XCLI command **target_rename**.

Complete the following steps:

1. Click the protected site, and if the Spectrum Accelerate family SRA does not have the OK status, click the Monitor tab, click **SRAs**, and click **Rescan**, as shown in Figure 9-29.

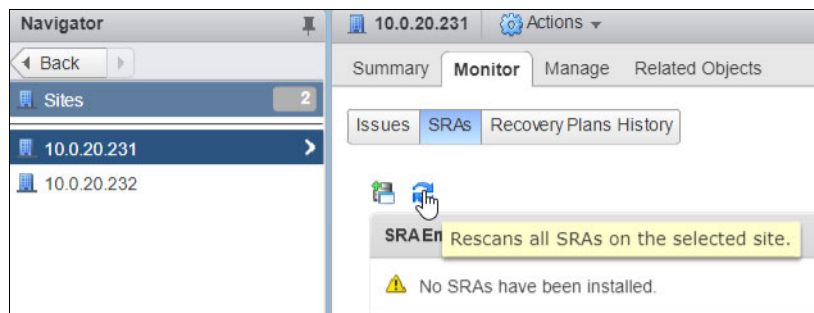


Figure 9-29 VMware vCenter SRM: Rescan SRAs on the protected site

2. Click the recovery site, and if the Spectrum Accelerate family SRA does not have the OK status, click the Monitor tab, click **SRAs**, and click **Rescan**, as shown in Figure 9-30.




Figure 9-30 VMware vCenter SRM: Rescan SRAs on the recovery site

- The Spectrum Accelerate family SRA is now registered, as shown in Figure 9-31. The IBM Spectrum Accelerate family SRA supports XIV, IBM FlashSystem A9000 and IBM FlashSystem A9000R, and IBM Spectrum Accelerate systems.



Issues SRAs Recovery Plans History	
 	
IBM Spectrum Accelerate Family SRA	
SRA:	IBM Spectrum Accelerate Family SRA
Status:	OK
Version:	3.0.0
Vendor:	IBM Corp.
Install Location:	C:\Program Files\VMware\VMware vCenter Site Recovery Manager\storage\sra\IBM SAF
Vendor URL:	http://www.ibm.com/support/knowledgecenter/STJTAG/hsg/hsg_sra_xiv_kcwelcome.htm
Supported Array Models:	IBM Corp., IBM XIV IBM Corp., IBM Spectrum Accelerate IBM Corp., IBM FlashSystem A9000/A9000R
Supported Software:	IBM XIV Remote Mirroring10.2 IBM XIV Remote Mirroring11 IBM FlashSystem A9000/A9000R Remote Mirroring12
Stretched Storage:	Supported

Figure 9-31 VMware vCenter SRM: IBM Spectrum Accelerate family SRA

- Right-click the protected site and select **Add Array Manager**, as shown in Figure 9-32.

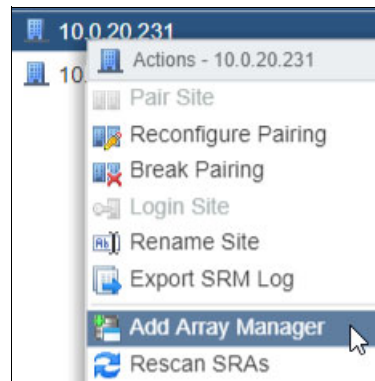


Figure 9-32 VMware vCenter SRM: Add Array Manager

5. Select **Add a pair of array managers** and click **Next**, as shown in Figure 9-33.

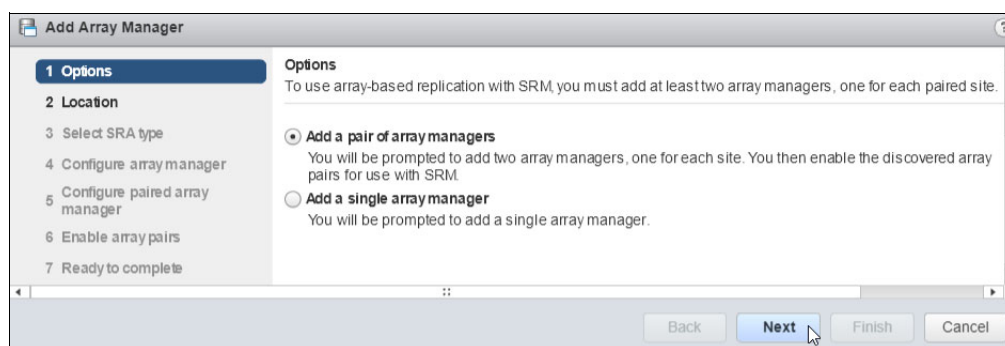


Figure 9-33 VMware vCenter SRM Add Array Manager: Options

6. Select a pair of sites and click **Next**, as shown in Figure 9-34.

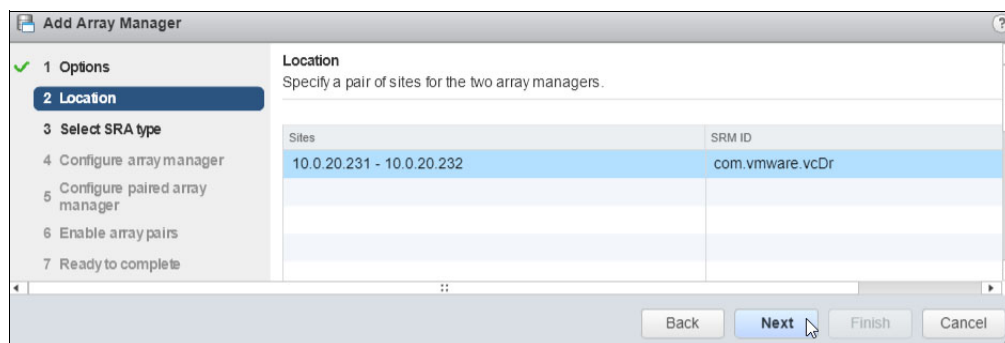


Figure 9-34 VMware vCenter SRM Add Array Manager: Location

7. Select **IBM Spectrum Accelerate Family SRA** as the SRA type and click **Next**, as shown in Figure 9-35.

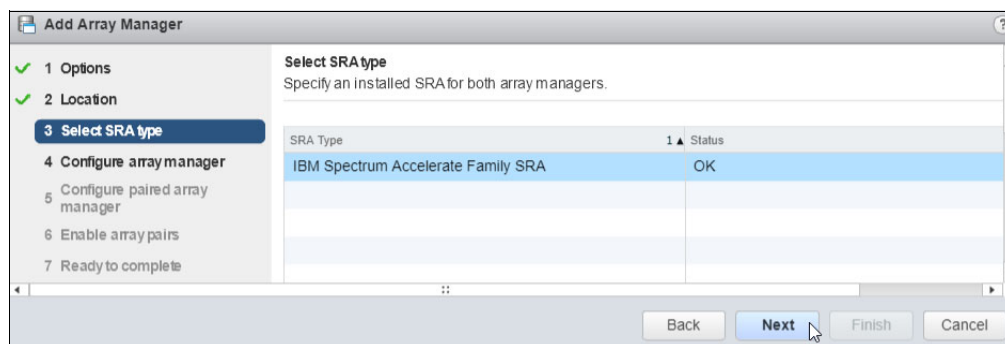


Figure 9-35 VMware vCenter SRM Add Array Manager: SRA type

8. Fill in Display Name, IP addresses, Username, and Password of the storage system on the protected site and click **Next**, as shown in Figure 9-36.

The screenshot shows the 'Add Array Manager' wizard in VMware vCenter SRM. The left sidebar lists the steps: 1 Options, 2 Location, 3 Select SRA type, 4 Configure array manager (selected), 5 Configure paired array manager, 6 Enable array pairs, and 7 Ready to complete. The main panel is titled 'Configure array manager' and contains the following fields:

- Display Name:** A9000
- Storage System:**
 - First management IP address / hostname:** 10.0.20.108
 - Second management IP address / hostname:** 10.0.20.109
 - Third management IP address / hostname:** 10.0.20.110
 - Username:** admin
 - Password:** (masked with asterisks)

At the bottom right, there are four buttons: Back, Next (highlighted with a mouse cursor), Finish, and Cancel.

Figure 9-36 VMware vCenter SRM Add Array Manager: Add array manager protected site

9. Fill in Display Name, IP addresses, Username, and Password of the storage system on the recovery site and click **Next**, as shown in Figure 9-37.

The screenshot shows the 'Add Array Manager' wizard in VMware vCenter SRM. The left sidebar lists the steps: 1 Options, 2 Location, 3 Select SRA type, 4 Configure array manager, 5 Configure paired array manager (selected), 6 Enable array pairs, and 7 Ready to complete. The main panel is titled 'Configure paired array manager' and contains the following fields:

- Display Name:** A9000R
- Storage System:**
 - First management IP address / hostname:** 10.0.20.100
 - Second management IP address / hostname:** 10.0.20.101
 - Third management IP address / hostname:** 10.0.20.105
 - Username:** admin
 - Password:** (masked with asterisks)

At the bottom right, there are four buttons: Back, Next (highlighted with a mouse cursor), Finish, and Cancel.

Figure 9-37 VMware vCenter SRM Add Array Manager: Add array manager recovery site

10. Check the array pair to enable and click **Next**, as shown in Figure 9-38.

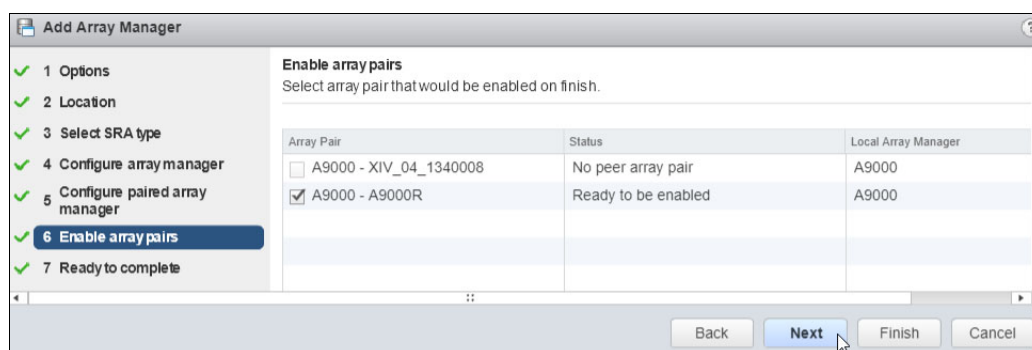


Figure 9-38 VMware vCenter SRM Add Array Manager: Enable array pairs

11. Review the settings and click **Finish**, as shown in Figure 9-39.

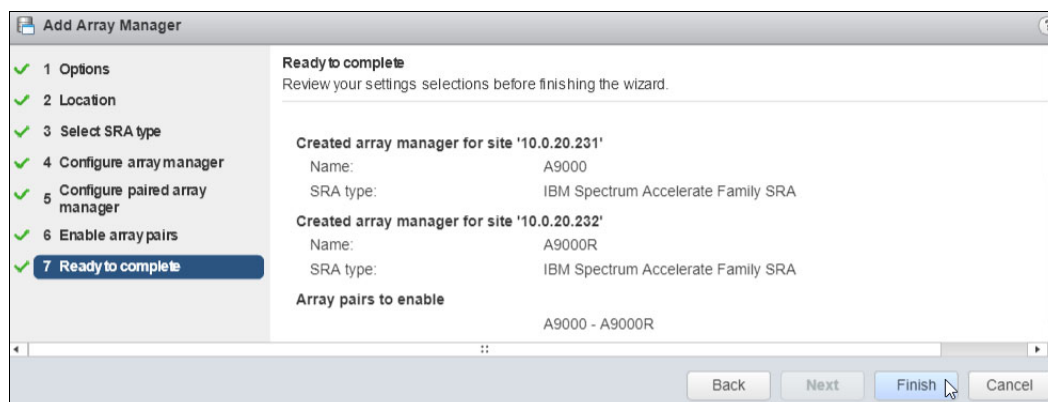


Figure 9-39 VMware vCenter SRM Add Array Manager: Ready to complete

12. Click **Site Recovery** → **Array Based Replication**, click the **Manage** tab, and click the **Array Pairs** tab. The array pair with its replication direction and local data store is displayed, as shown in Figure 9-40.

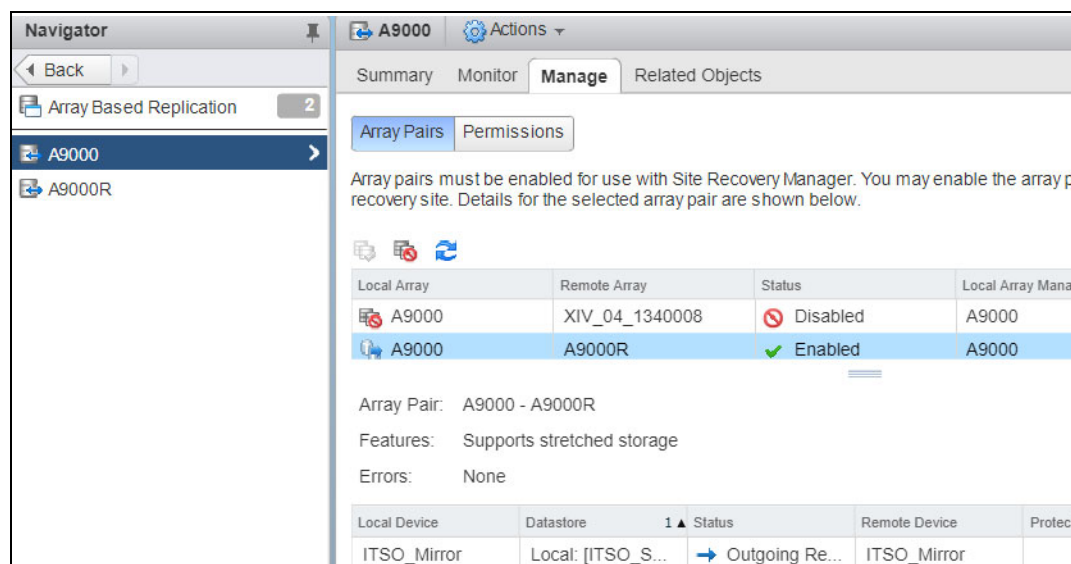


Figure 9-40 VMware vCenter SRM Array Based Replication

9.3.5 Creating protection groups

Protection groups consist of pointers to the replicated vSphere data stores containing collections of VMs that are failed over from the protected site to the recovery site during VMware vCenter SRM DR, planned migration, or testing operations. In a way, protection groups are the VMware equivalent of storage consistency groups in that they are logical groupings that are defined at the data store level instead of the logical volume level.

The process of creating protection groups within VMware vCenter SRM consists of the following steps, and is a prerequisite to running VMware vCenter SRM workflows:

1. Click **Site Recovery** → **Sites**, right-click the protected site, and click **Create Protection Group**, as shown in Figure 9-41.

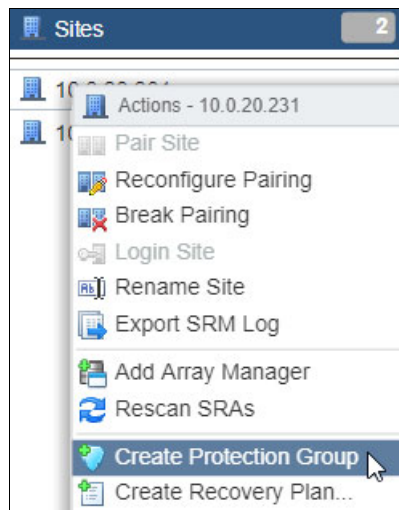


Figure 9-41 VMware vCenter SRM Create Protection Group

2. Fill in a name for the protection group, select the location, and click **Next**, as shown in Figure 9-42.

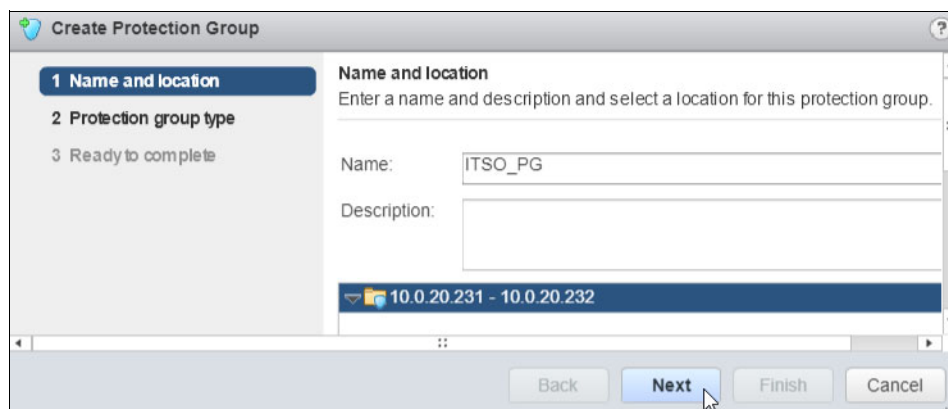


Figure 9-42 VMware vCenter SRM Create Protection Group: Name and location

3. Select the protection group type and click **Next**, as depicted in Figure 9-43.

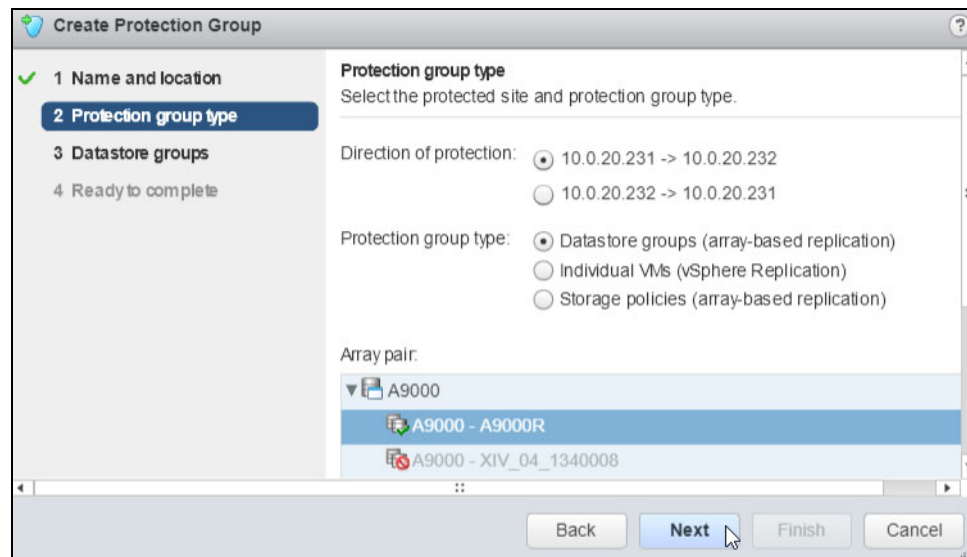


Figure 9-43 VMware vCenter SRM Create Protection Group: Type

4. Select the data store groups to protect, and the protected VMs are shown. Click **Next**, as shown in Figure 9-44.

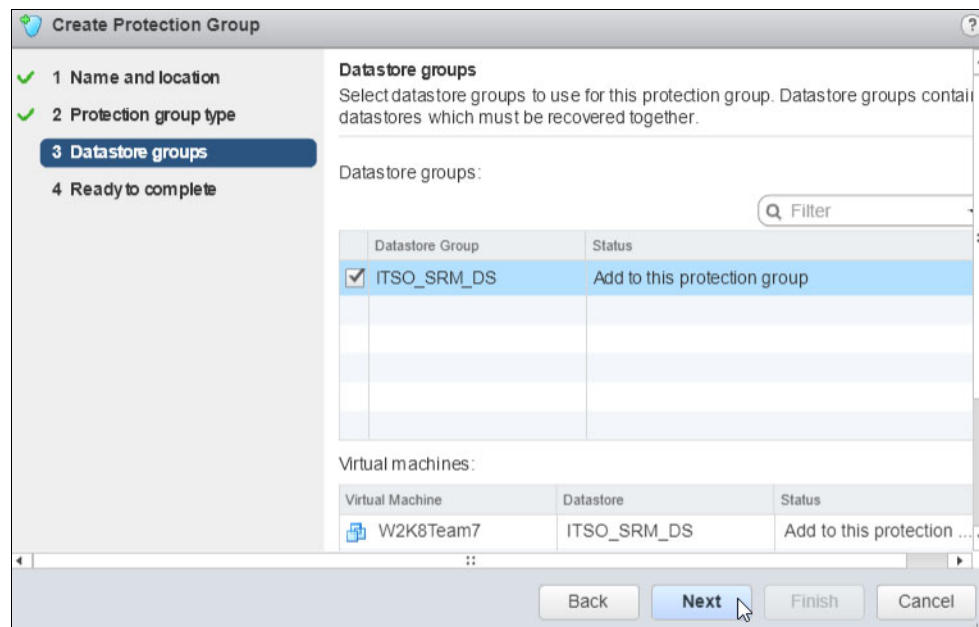


Figure 9-44 VMware vCenter SRM Create Protection Group: Datastore groups

5. Review the settings and click **Finish**, as shown in Figure 9-45.

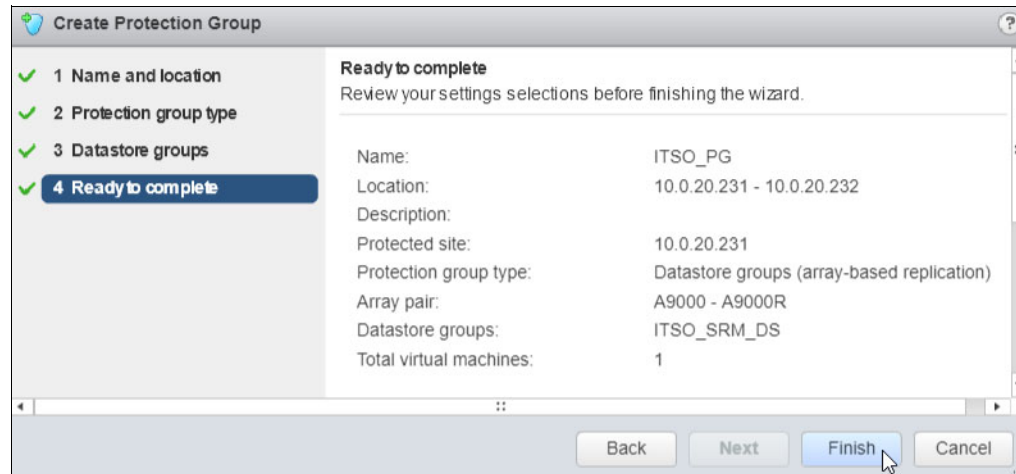


Figure 9-45 VMware vCenter SRM Create Protection Group: Ready to complete

6. Click **Site Recovery** → **Protection Groups**. The status of the protection group is shown in Figure 9-46.

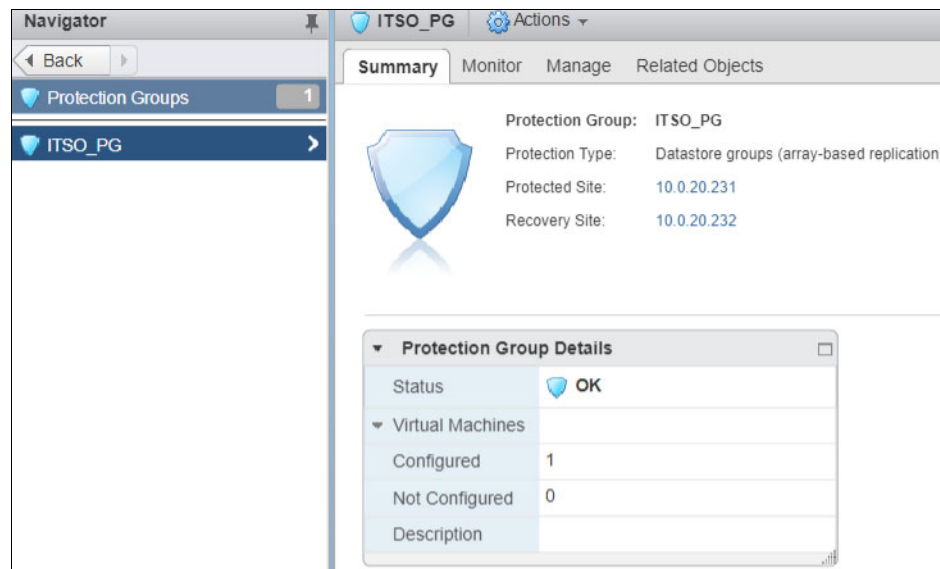


Figure 9-46 VMware vCenter SRM Protection Group

9.3.6 Creating recovery plans

Recovery plans govern how VMs in one or more protection groups are restored at the recovery site. The steps that follow demonstrate how to customize the plan to meet specific needs:

1. Click **Site Recovery** → **Sites**, right-click the protected site, and click **Create Recovery Plan**, as shown in Figure 9-47.




Figure 9-47 VMware vCenter SRM Create Recovery Plan

2. Fill in a name for the recovery plan, select the location, and click **Next**, as shown in Figure 9-48.

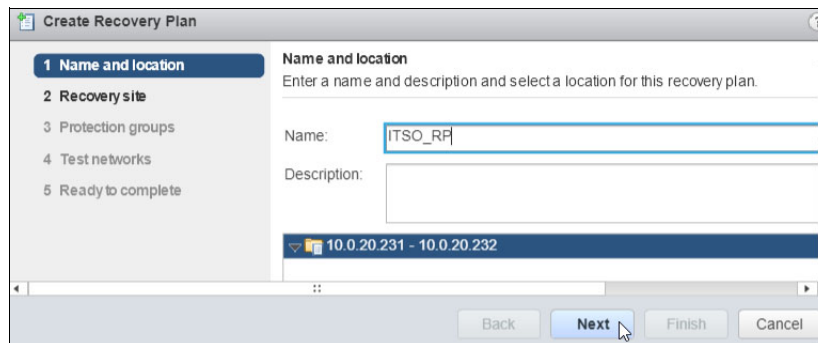


Figure 9-48 VMware vCenter SRM Create Recovery Plan: Name and location

3. Select the recovery site and click **Next**, as shown in Figure 9-49.



Figure 9-49 VMware vCenter SRM Create Recovery Plan: Recovery site

4. Select the protection groups for the recovery plan and click **Next**, as shown in Figure 9-50.

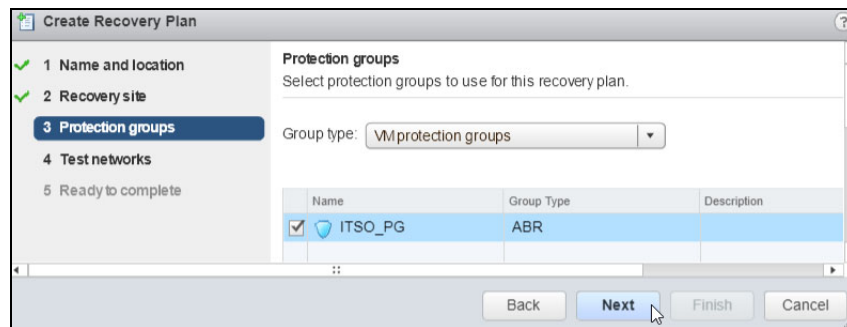


Figure 9-50 VMware vCenter SRM Create Recovery Plan: Protection Groups

5. Select which network to use when testing the recovery plan and click **Next**, as shown in Figure 9-51.

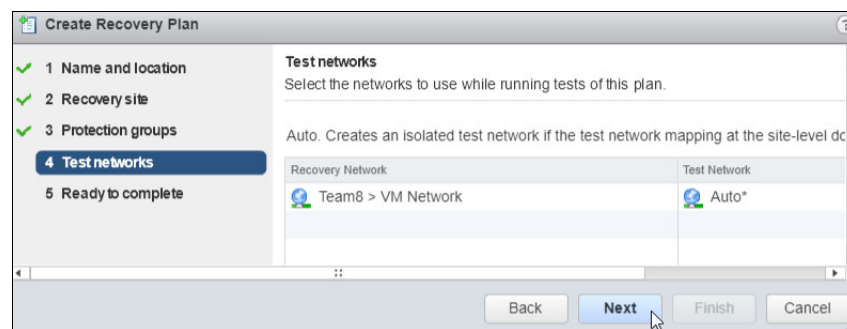


Figure 9-51 VMware vCenter SRM Create Recovery Plan: Test networks

6. Review the settings and click **Finish**, as shown in Figure 9-52.

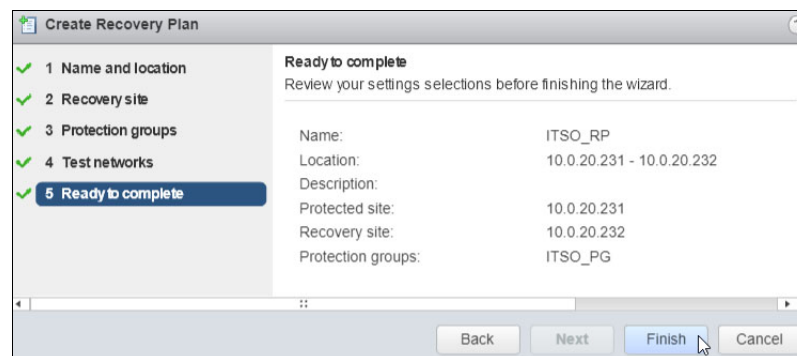


Figure 9-52 VMware vCenter SRM Create Recovery Plan: Ready to complete

7. Click **Site Recovery** → **Recovery Plans**. The status of the recovery plan is displayed, as shown in Figure 9-53.

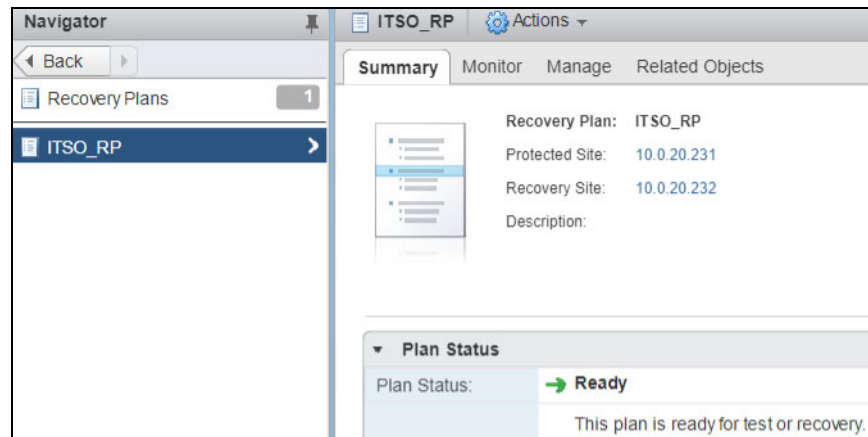


Figure 9-53 VMware vCenter SRM Recovery Plan

8. By default, every recovery plan includes a sequence of steps, also known as a workflow, which consists of pre-set values and directives to control how VMs in a protection group are recovered at the recovery site. These workflows can be uniquely defined for each of the activity categories, which consists of Test Steps, Cleanup Steps, Recovery Steps, and Reprotect Steps.

Click the Monitor tab and then the Recovery Steps tab, and a default workflow sequence representing the Test Steps is shown in Figure 9-54. For more information about developing recovery plans by specifying customized recovery steps, consult the VMware vCenter SRM documentation, found at:

https://www.vmware.com/support/pubs/srm_pubs.html

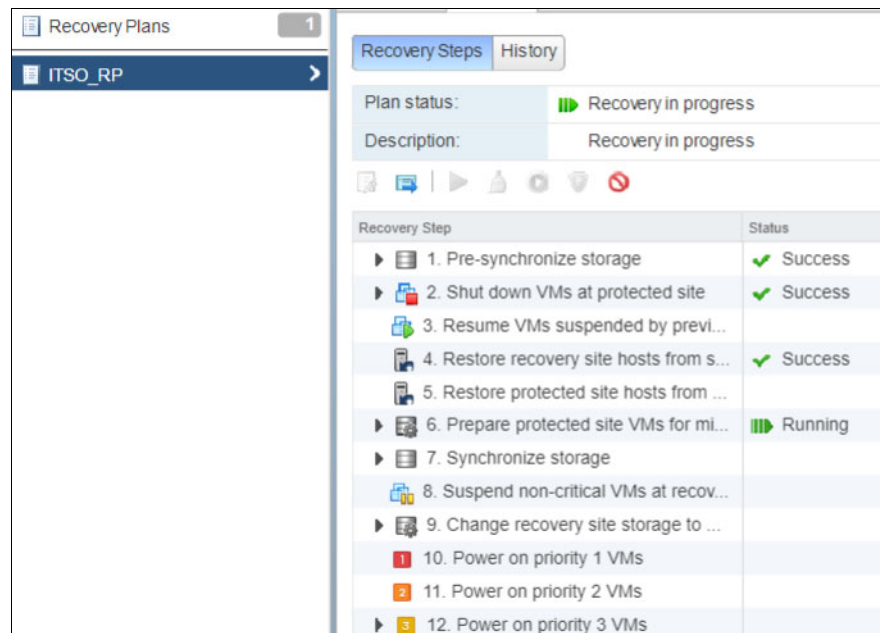


Figure 9-54 VMware vCenter SRM Recovery Plan: Review Recovery Steps

Note: Consider these basic guidelines when defining recovery plan workflows:

- ▶ Configure the VM dependencies across priority groups instead of setting VM dependencies individually per VM, which ensures that VMs are started in parallel. The storage systems are optimized for parallel workloads, so this greatly improves performance.
- ▶ Define an action before a test or failover at the recovery site, such as cloning down or suspending low-priority VMs to free recovery resources at the recovery site.
- ▶ Define the allocation of resources and any networking changes that are required by the VMs.
- ▶ Build call-outs that cause test or failover process to pause and present instructions to the administrator, or specify scripts in the recovery process.

9.3.7 Testing recovery plans

The VMware vCenter SRM recovery plan testing capability represents an invaluable asset in the development of robust recovery strategies because it empowers administrators to eliminate uncertainty and dramatically reduce the risk of failed recovery by identifying and addressing any issues pro-actively without interrupting the operation of the production environment. The two primary phases of this process include the test itself and the subsequent removal of temporary elements that are created during the test and restoring the recovery plan to its initial state by using the cleanup process. Both of these procedures are illustrated in the next series of steps:

1. Click **Test recovery plan**, as shown in Figure 9-55.

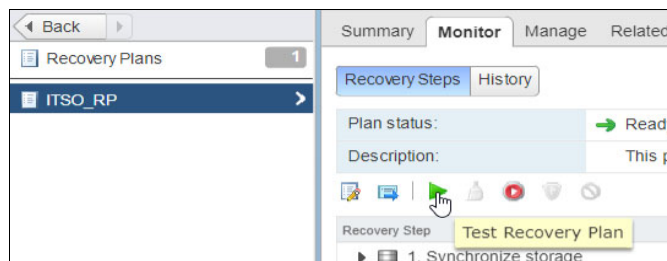


Figure 9-55 VMware vCenter SRM Test

2. Select **Replicate recent changes to recovery site** and click **Next** (see Figure 9-56).



Figure 9-56 VMware vCenter SRM Test: Confirmation Options

- Review the settings and click **Finish**, as shown in Figure 9-57.



Figure 9-57 VMware vCenter SRM Test: Ready to complete

- The test starts and the current recovery steps are displayed, as shown in Figure 9-58.

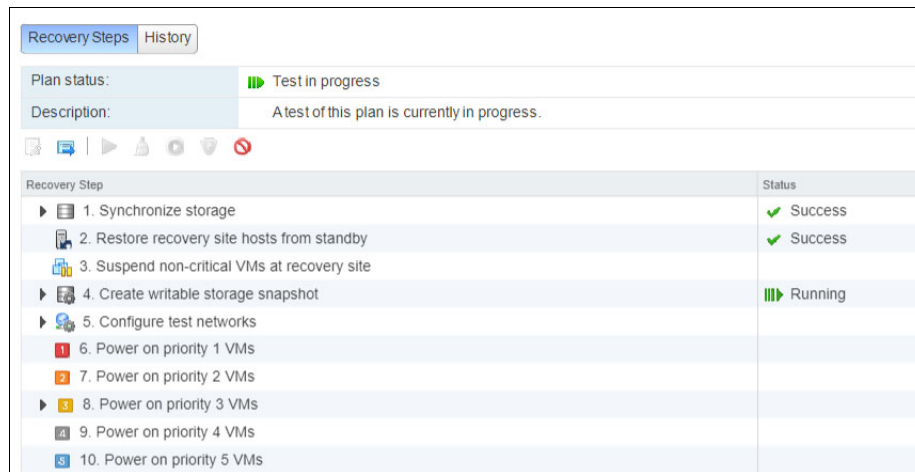


Figure 9-58 VMware vCenter SRM Test: Recovery Steps

- A snapshot is created on the recovery site storage system. Go to Hyper-Scale Manager (GUI) and click **Pools and Volumes View** → **Snapshots**, as illustrated in Figure 9-59.

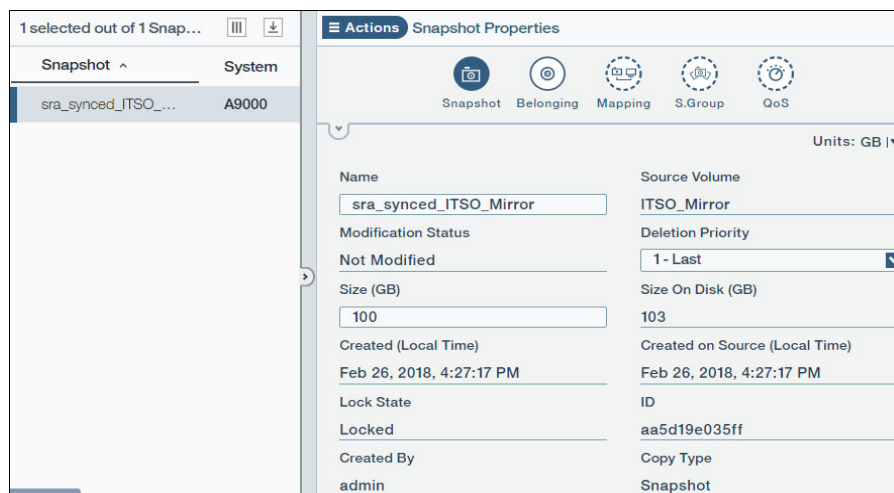


Figure 9-59 VMware vCenter SRM Test: Snapshot on the storage system recovery site

- Go to the vSphere Web Client on the recovery site vCenter and select **Hosts and Clusters**, as shown in Figure 9-60. The VM is running in a safe network environment.

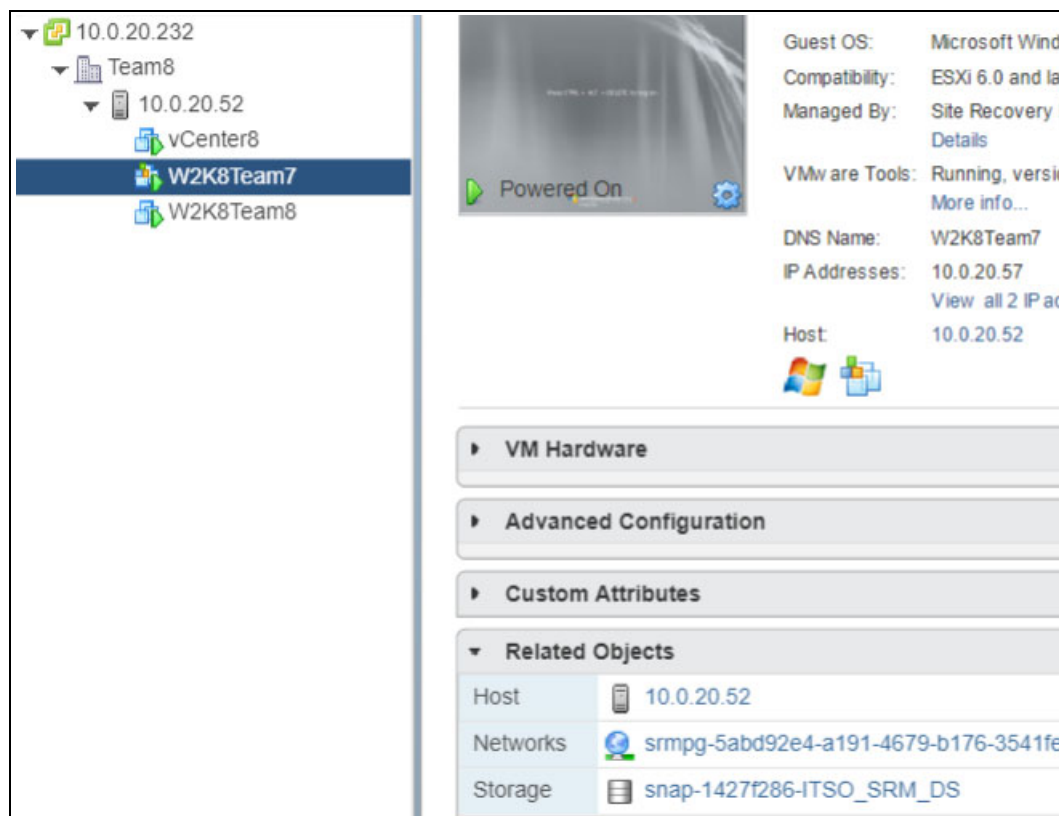


Figure 9-60 VMware vCenter SRM Test: VM on the recovery site

9.3.8 Cleanup

The cleanup process is necessary to restore the recovery plan to its initial state after a test is performed. Complete the following steps:

- Following the test completion, administrators receive a notification of the test completion status. Click the **Cleanup recovery plan** button, as shown in Figure 9-61.

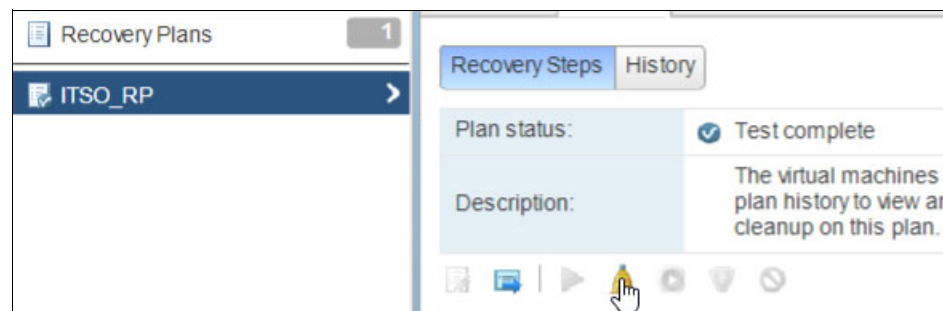


Figure 9-61 VMware vCenter SRM Cleanup

- Click **Next** to remove the test environment, as shown in Figure 9-62.

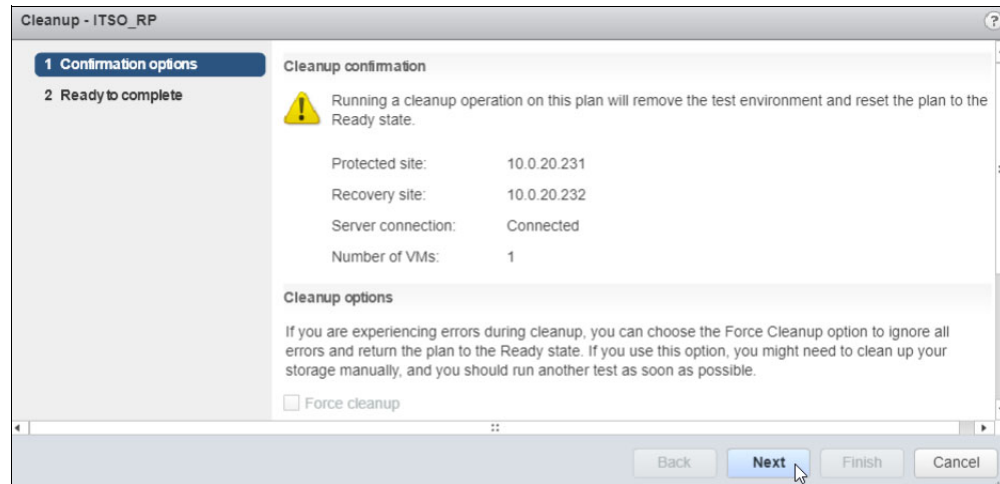


Figure 9-62 VMware vCenter SRM Cleanup: Confirmation Options

- Review the settings and click **Finish**, as shown in Figure 9-63.

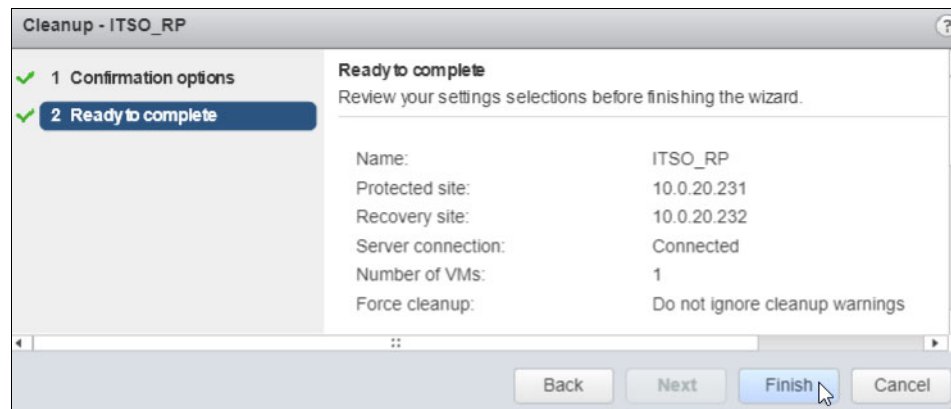


Figure 9-63 VMware vCenter SRM Cleanup: Ready to complete

The cleanup starts and the current recovery steps are displayed, as shown in Figure 9-64.

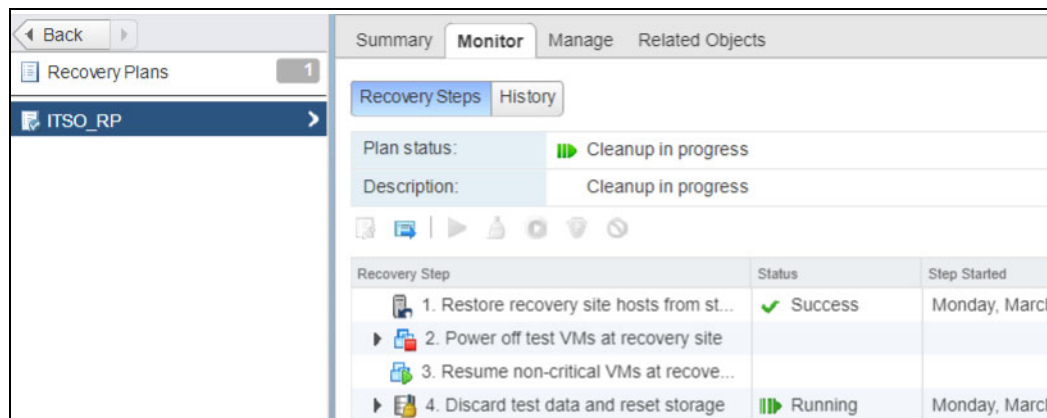


Figure 9-64 VMware vCenter SRM Cleanup: Recovery Steps

Note: If the cleanup process fails to successfully complete, revisit the Cleanup wizard to specify the **Force Cleanup** option and return the recovery plan to the ready state. The caveat of this method is the potential requirement for administrator intervention to restore the storage to its initial state.

9.3.9 Recovery

Although the workflows that are outlined in the VMware vCenter SRM recovery plans are fully automated, vSphere administrators are responsible for initiating a site recovery by completing the following steps:

1. Click the **Run recovery plan** icon to initiate recovery, as shown in Figure 9-65.

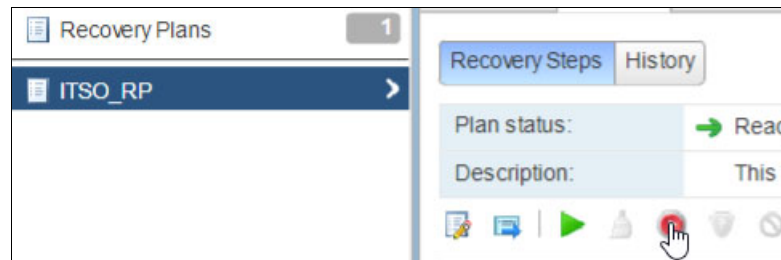


Figure 9-65 VMware vCenter SRM Recovery

2. Select **Planned Migration**, select the check box that explains that you understand the process, and click **Next**, as shown in Figure 9-66.

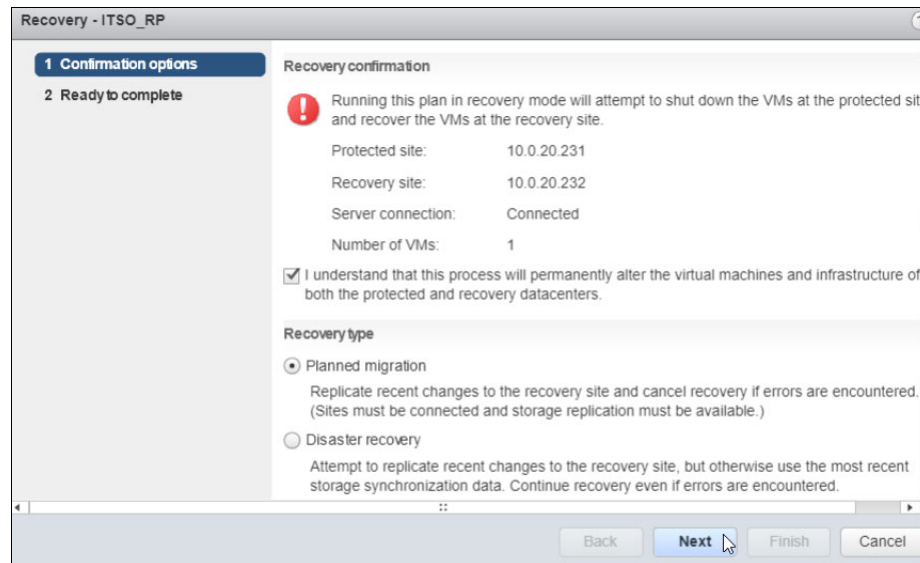


Figure 9-66 VMware vCenter SRM Recovery: Confirmation Options

- Review the settings and click **Finish**, as shown in Figure 9-67.

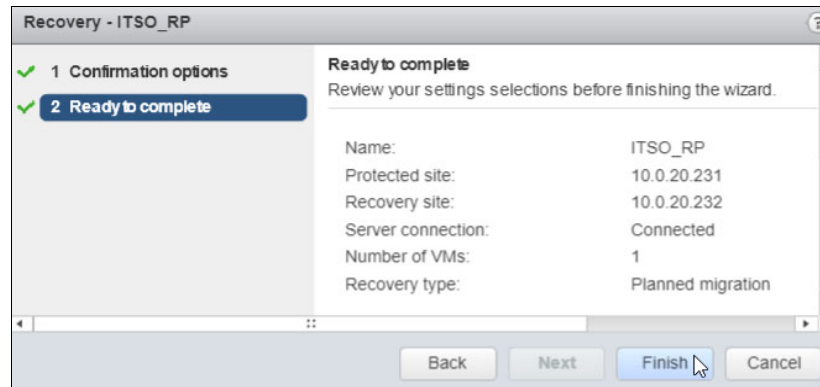


Figure 9-67 VMware vCenter SRM Recovery: Ready to complete

- The recovery starts and the current recovery steps are displayed, as shown in Figure 9-68.

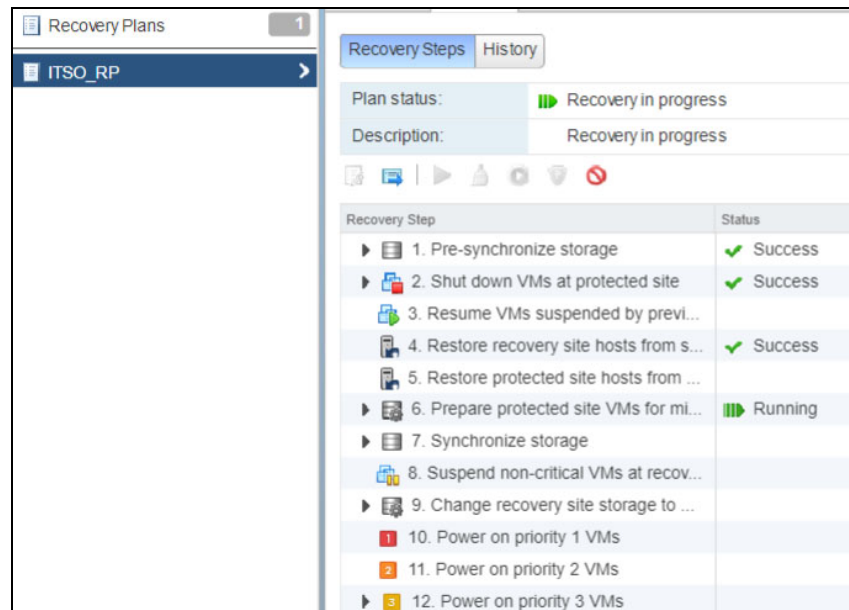


Figure 9-68 VMware vCenter SRM Recovery: Recovery Steps

- The VM shuts down on the protected site, the mirror replicates changes, and the Availability Role of the volume on the recovery site storage system changes to “Primary (not as designated)”, as shown in Figure 9-69. The volume is mapped to the ESXi host on the recovery site. The ESXi host rescans the storage, mounts the data store, and starts the VM.

1 selected out of 2 Volumes			
	Volume	System	Availability Role
	ITSO_Mirror	A9000	Primary
	ITSO_Mirror	A9000R	Primary (not as designat...)

Figure 9-69 VMware vCenter SRM Recovery: Volumes on the protected and recovery sites

- Go to the vSphere Web Client on the recovery site vCenter and select **Hosts and Clusters**, as shown in Figure 9-70. The VM is running on the recovery site.




Figure 9-70 VMware vCenter SRM Recovery: VM on the recovery site

9.3.10 The reprotect process

The reprotect process can be started only after the successful completion of a recovery process, and it effectively automates the process of reversing the role of the protected site to facilitate the continuance of site-level protection for the vSphere production environment that is running at the recovery site following a planned or unplanned migration. Activating the reprotect process can be completed by completing the following steps:

- Click the **Reprotect recovery plan** icon to initiate the reprotect, as shown in Figure 9-71.

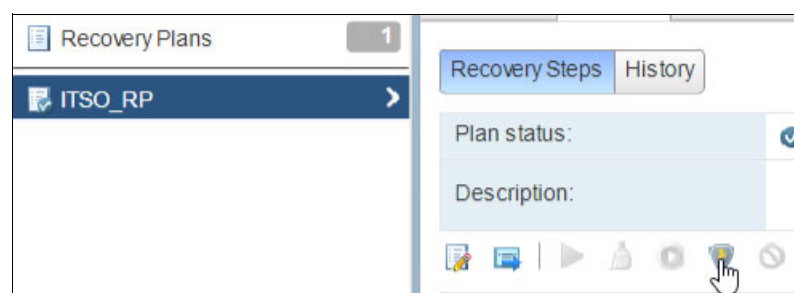


Figure 9-71 VMware vCenter SRM Reprotect

2. Select the **I understand that this operation cannot be undone** check box. The Force Cleanup check box is initially unavailable. Click **Next**, as shown in Figure 9-72.

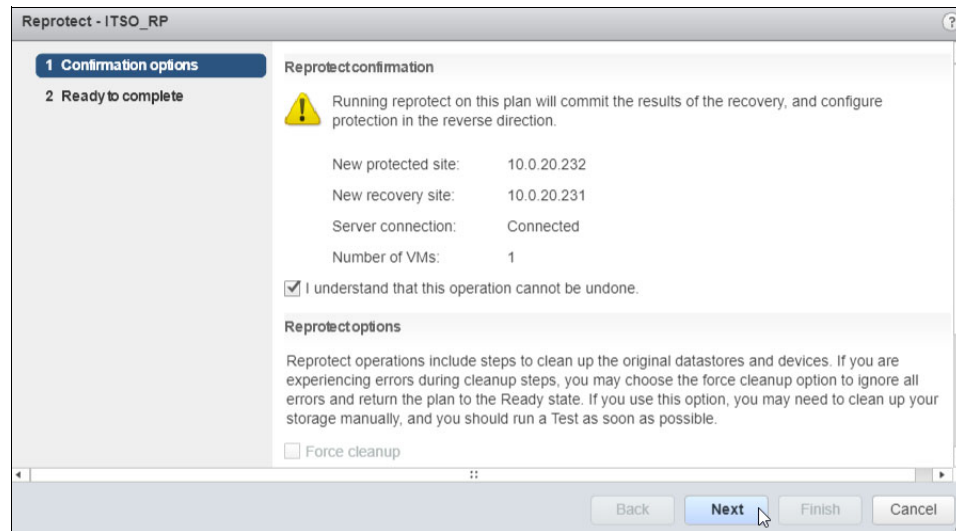


Figure 9-72 VMware vCenter SRM Reprotect: Confirmation Options

3. Review the settings and click **Finish**, as shown in Figure 9-73.

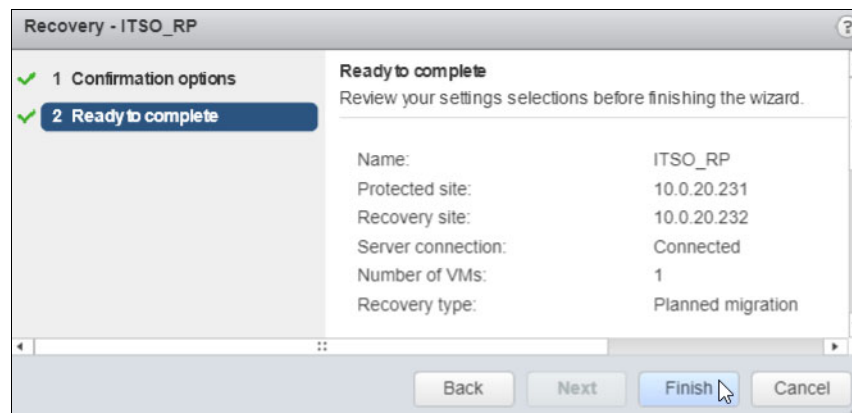


Figure 9-73 VMware vCenter SRM Reprotect: Ready to complete

4. The reprotect starts and the current recovery steps are displayed, as shown in Figure 9-74.

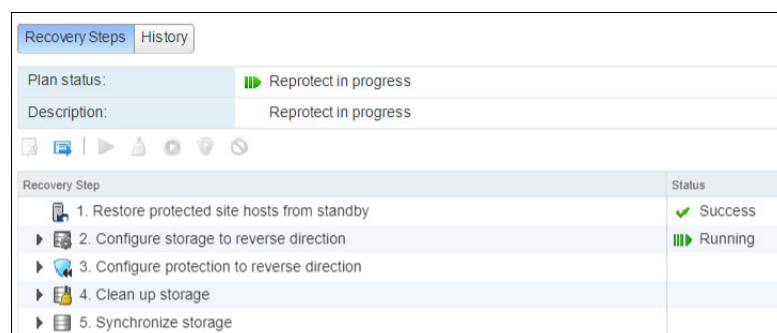


Figure 9-74 VMware vCenter SRM Reprotect: Recovery Steps

5. The Availability Role of the volume on the protected site storage system changes to Secondary, as shown in Figure 9-75.



The screenshot shows a table titled "1 selected out of 2 Volumes". The table has three columns: "Volume", "System", and "Availability Role". The first row is highlighted and shows "ITSO_Mirror" under Volume, "A9000R" under System, and "Primary" under Availability Role. The second row shows "ITSO_Mirror" under Volume, "A9000" under System, and "Secondary" under Availability Role. There is a small icon to the left of the first row.

1 selected out of 2 Volumes		
Volume	System	Availability Role
ITSO_Mirror	A9000R	Primary
ITSO_Mirror	A9000	Secondary

Figure 9-75 VMware vCenter SRM Reprotect: Volumes on the protected and recovery site

9.3.11 Failing back to the protected site

To fail back to the protected site, run another recovery, as described in 9.3.9, “Recovery” on page 166, and then run another reprotect, as described in 9.3.10, “The reprotect process” on page 168.

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Some publications that are referenced in this list might be available in softcopy only.

- ▶ *IBM FlashSystem A9000 and A9000R, IBM XIV, and IBM Spectrum Accelerate with IBM SAN Volume Controller Best Practices*, REDP-5408
- ▶ *IBM Hyper-Scale Manager for IBM Spectrum Accelerate Family: IBM XIV, IBM FlashSystem A9000 and A9000R, and IBM Spectrum Accelerate*, SG24-8376
- ▶ *IBM FlashSystem A9000 and A9000R Business Continuity Solutions*, REDP-5401
- ▶ *IBM FlashSystem A9000 and IBM FlashSystem A9000R Architecture and Implementation*, SG24-8345
- ▶ *IBM FlashSystem A9000, IBM FlashSystem A9000R, and IBM XIV Storage System Host Attachment and Interoperability*, SG24-8368
- ▶ *IBM XIV Storage System Architecture and Implementation*, SG24-7659
- ▶ *IBM XIV Storage System Business Continuity Functions*, SG24-7759
- ▶ *IBM XIV Storage System: Host Attachment and Interoperability*, SG24-7904

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, draft and additional materials, at the following website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services



REDP-5425-01

ISBN 0738456802

Printed in U.S.A.

Get connected

