Redbooks
ibm.com/redbooks

Configuring VMware Virtual Volumes for Systems Powered by IBM Spectrum Virtualize

Adam Reid
Jacob Winch
Warren Hawkins
Jon Tate

Storage

IBM

Redbooks

International Technical Support Organization

Configuring VMware Virtual Volumes for Systems Powered by IBM Spectrum Virtualize

March 2016

SG24-8328-00

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

First Edition (March 2016)

This edition applies to the products and code levels described in Chapter 2, “Prerequisites” on page 5.

Contents

Notices vii
Trademarks viii

IBM Redbooks promotions ix

Preface xi
Authors xi
Now you can become a published author, too xiii
Comments welcome xiii
Stay connected to IBM Redbooks xiii

Chapter 1. Implementing VMware Virtual Volumes for systems powered by IBM Spectrum Virtualize 1
1.1 What are Virtual Volumes? 2
1.2 Overview of key components 2
1.2.1 vSphere 6.0 3
1.2.2 IBM Spectrum Control Base Edition 3
1.2.3 Storage system managed by IBM Spectrum Virtualize 3

Chapter 2. Prerequisites 5
2.1 Compatible storage systems 6
2.2 Prerequisites 6
2.2.1 Code levels 6
2.2.2 NTP Server Configuration 7
2.2.3 Server configuration for IBM Spectrum Control Base Edition 7
2.2.4 Licensing requirements 7
2.3 Example configuration 8

Chapter 3. Terminology 9
3.1 IBM Spectrum Virtualize 10
3.2 Storage system 10
3.3 Parent pool 10
3.4 Child pool 10
3.5 Virtual Volumes 10
3.6 VVols-enabled host 11
3.7 Protocol endpoint 12
3.8 Bind and unbind operations 13
3.9 Utility volume 14
3.10 VASA Provider 14
3.11 IBM Spectrum Control Base Edition 14
3.12 High availability group 14
3.13 IBM Spectrum Control Base Edition user 14
3.14 Storage service 15
3.15 Storage resource 15
3.16 Storage space 15
3.17 Storage Container 15
3.18 VVol Datastore 16
3.19 Storage Policy-Based Management 16

Chapter 4. Installation and configuration 17
4.1 Configuring the storage system 18
4.1.1 Enabling VVols 18
4.1.2 Enabling VMware ESXi hosts for VVols 21
4.2 Configuring IBM Spectrum Control Base Edition 22
4.2.1 Downloading the IBM Spectrum Control Base Edition installation files 22
4.2.2 Installing IBM Spectrum Control Base Edition 23
4.2.3 Changing the IBMSC administration account 23
4.2.4 Accessing the IBM Spectrum Control Base Edition web interface 24
4.2.5 An overview of the SCB web interface 24
4.2.6 Changing the default administrator password 26
4.2.7 Storage credentials 27
4.2.8 VASA credentials 27
4.2.9 General settings 28
4.2.10 Server certificate 29
4.2.11 Adding a storage system 30
4.2.12 Adding a storage space 31
4.2.13 Adding a storage service 32
4.2.14 Defining and attaching storage resources 34
4.3 Configuring VMware vSphere Web Client 35
4.3.1 Adding a storage provider 35
4.3.2 Adding a VVol datastore in VMware vSphere Web Client 37
4.3.3 Setup complete 38
4.4 Configuring additional SCB servers for high availability 39
4.4.1 Downloading the IBM Spectrum Control Base Edition installation files 39
4.4.2 Installing IBM Spectrum Control Base Edition 40
4.4.3 Changing the IBMSC administration account 40
4.4.4 Accessing the IBM Spectrum Control Base Edition web interface 41
4.4.5 Changing the default administrator password 41
4.4.6 Storage Credentials 42
4.4.7 VASA credentials 43
4.4.8 General settings 43
4.4.9 Server certificate 44
4.4.10 Adding a storage system 44
4.4.11 Configuring the VMware vSphere Web Client 46

Chapter 5. Managing and maintaining a Virtual Volumes environment 49
5.1 Example workflows 50
5.1.1 Allocating more storage to the vSphere administrator 50
5.1.2 Reclaiming storage from the vSphere administrator 52
5.2 Storage system tasks 53
5.2.1 Managing VMware ESXi hosts 53
5.2.2 Managing the utility volume 54
5.3 IBM Spectrum Control Base Edition tasks 55
5.3.1 Managing storage spaces 55
5.3.2 Managing storage services 56
5.3.3 Managing storage resources (child pools) 59
5.3.4 Adding more SCB instances 64
5.3.5 Managing Storage Systems 64
5.4 VMware vSphere Web Client tasks 66
5.4.1 Configuring multipathing 66
5.4.2 Managing queue depths 66
5.4.3 Managing storage providers 67

5.4.4 Managing VVol datastores 68
5.4.5 Working with Storage Policy-Based Management 71
5.4.6 Migrating virtual machines onto VVol datastores 71

Chapter 6. Considerations and best practices 73
6.1 Storage system considerations 74
6.1.1 Defining Parent Pools 74
6.1.2 Mirroring the utility volume 76
6.1.3 Performing an upgrade on a storage system with Virtual Volumes enabled. 76
6.2 IBM Spectrum Control Base Edition considerations 77
6.2.1 High availability 77
6.2.2 Defining storage spaces 78
6.2.3 Using the IBM Storage Enhancements for VMware vSphere Web Client in a Virtual Volumes environment 81

Chapter 7. Storage Policy-Based Management 83
7.1 Overview of Storage Policy-Based Management 84
7.2 Storage capabilities advertised by SCB 85
7.2.1 Encryption 86
7.2.2 Flash 86
7.2.3 Space efficiency (compression) 86
7.3 Defining policies in vSphere Web Client 87
7.4 Use case scenarios and examples 88
7.4.1 Scenario 1 88
7.4.2 Scenario 2 88
7.4.3 Scenario 3 89
7.4.4 Configuring storage services with multiple child pools 91

Chapter 8. Advanced concepts, operations, and functions in the Virtual Volumes environment 93
8.1 Types of VVols 94
8.2 Number of VVols 94
8.3 IBM FlashCopy functionality in a VVols environment 95
8.3.1 A basic introduction to IBM FlashCopy 95
8.4 Clone operations 96
8.5 Snapshot operations 98
8.5.1 Taking a snapshot. 98
8.5.2 Reverting to a snapshot 100
8.5.3 Deleting snapshots 102
8.6 Cold migration and storage vMotion in a Virtual Volumes environment 103
8.7 vMotion in a Virtual Volumes environment 104

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.

Easy Tier®
FlashCopy®
FlashSystem™
IBM®
IBM FlashSystem®
IBM Spectrum™
IBM Spectrum Control™
IBM Spectrum Virtualize™
Real-time Compression™
Redbooks®
Redbooks (logo)  ®
Storwize®
System Storage®

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.

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

Find and read thousands of IBM Redbooks publications

- Search, bookmark, save and organize favorites
- Get personalized notifications of new content
- Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App

iOS **Download Now** Android

Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!

ibm.com/Redbooks
About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

IBM® Spectrum Virtualize and VMware's Virtual Volumes (VVols) are paving the way toward a true IBM Software Defined Environment (SDE). IBM Spectrum™ Virtualize is at the core of software-defined storage. The addition of VVols enables a fundamentally more efficient operational model for storage in virtualized environments, centering it around the virtual machine (VM) rather than the physical infrastructure.

This IBM Redbooks® publication provides an overview of the VVols management framework and its implementation on storage systems managed by IBM Spectrum Virtualize™.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), San Jose Center, and IBM Hursley Park, UK.

Adam Reid has more than 15 years of computer engineering experience. Focused more recently on IBM Spectrum Virtualize, he's been deeply involved with VMware and the testing and configuration of virtualized environments pivotal to the future of software-defined storage. Adam has designed, tested, and validated systems to meet the demands of a wide range of mid-range and enterprise environments.

Jacob Winch is a Test Engineer in IBM Systems at the IBM Manchester Lab, UK. He works as part of the global team responsible for developing, testing, and supporting the SAN Volume Controller and IBM Storwize® Family Products. He has worked closely with VMware during the development of Virtual Volumes support for IBM storage systems.

Warren Hawkins has a background in Infrastructure Support in predominantly Microsoft Windows and VMware environments, and has gained 13 years experience working in 2nd and 3rd line support in both public and private sector organizations. Since joining IBM in 2013, Warren has played a crucial part in customer engagements. Using his field experience, he established himself as the Test Lead for the IBM Spectrum Virtualize product family, focusing on clustered host environments.

Jon Tate is a Project Manager for IBM System Storage® storage area network (SAN) solutions at the International Technical Support Organization (ITSO), San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2 and Level 3 support for IBM mainframe storage products. Jon has 30 years of experience in storage software and management, services, and support. He is an IBM Certified IT Specialist, an IBM SAN Certified Specialist, and is Project Management Professional (PMP) certified. He is also the UK Chairman of the Storage Networking Industry Association (SNIA).

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

Carlos Fuente
Leigh Reed
Richard Bordoli
Katja Gebuhr
Djihed Afifi

SAN Volume Controller/Storwize Test and Development Teams

Yi Ming Zhang

IBM Spectrum Control Base Edition Development Team

Ben Meadowcroft
Ken Werneburg

VMware Virtual Volumes Team

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 2 - 6 weeks in length, and you can participate either in person or as a remote resident working from your home base.

Learn more about the residency program, browse the residency index, and apply online:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us.

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

- Use the online **Contact us** review Redbooks form:

  ibm.com/redbooks

- Send your comments in an email:

  redbooks@us.ibm.com

- Mail your comments:

  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

1

Implementing VMware Virtual Volumes for systems powered by IBM Spectrum Virtualize

IBM Spectrum Virtualize and VMware's Virtual Volumes (VVols) are paving the way toward a true IBM Software Defined Environment. IBM Spectrum Virtualize is at the core of software-defined storage. The addition of VVols enables a fundamentally more efficient operational model for storage in virtualized environments, centering it around the virtual machine (VM) rather than the physical infrastructure.

This IBM Redbooks publication provides an overview of the VVols management framework and its implementation on storage systems managed by IBM Spectrum Virtualize.

This chapter provides information about the following topics:

- 1.1, "What are Virtual Volumes?" on page 2
- 1.2, "Overview of key components" on page 2

1.1 What are Virtual Volumes?

Before VVols, a virtual machine disk (VMDK) would be presented to a VM in the form of a file. This file represents a disk to the VM, which is then accessed by the guest operating system in the same way as a physical disk is accessed on a physical server. This VMDK is stored on a VMware file system (VMFS) formatted datastore.

The VMFS datastore is hosted by a single volume on a storage system, such as the IBM Storwize V7000. A single VMFS datastore, sometimes referred to as the *VMFS blender*, can have hundreds or even thousands of VMDKs.

VVols provides a one-to-one mapping between the VM's disks and the volumes (VVols) hosted by the storage system. This VVol is wholly owned by the VM. Making the VVol available at the storage level enables storage system-based operations at the granular VM level. For example, capabilities such as compression and encryption can be applied to an individual VM. Similarly, IBM FlashCopy® can be used at the VVol level, when performing snapshot and clone operations.

1.2 Overview of key components

Figure 1-1 provides a high-level overview of the three key components that enable the VVols management framework.

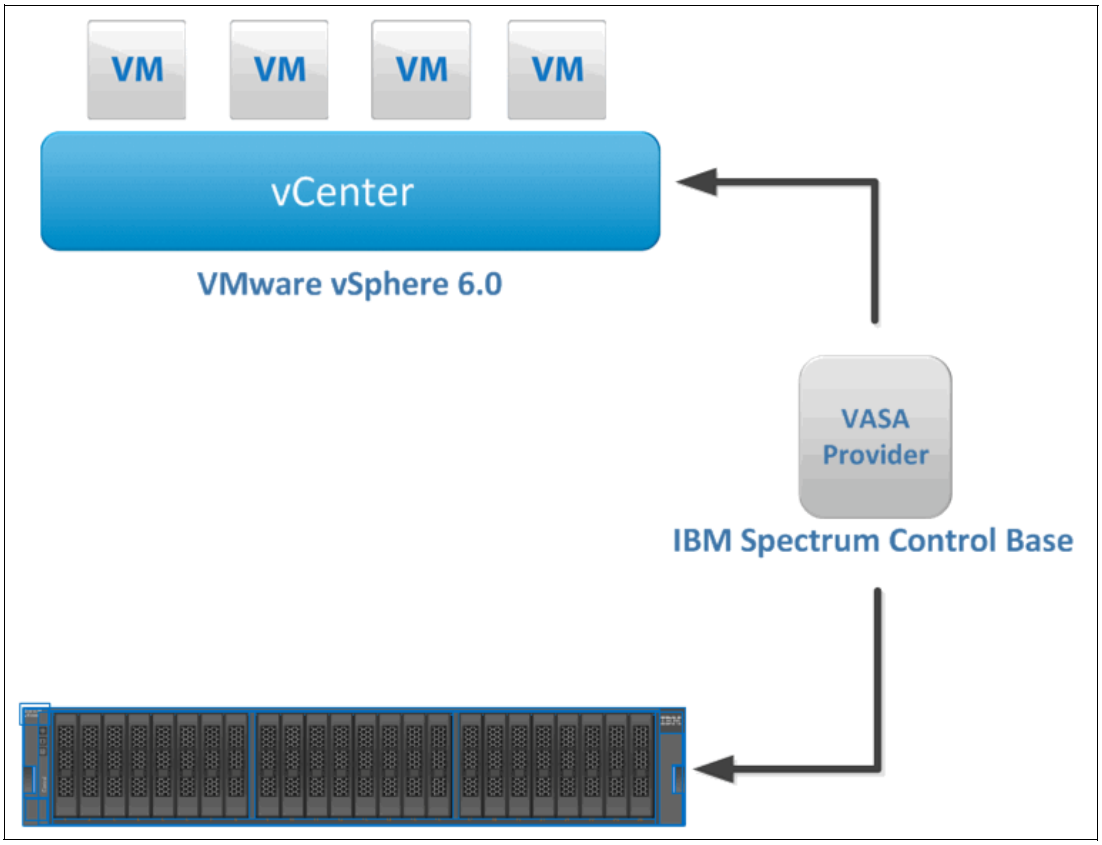

Figure 1-1 Virtual Volumes management framework

1.2.1 vSphere 6.0

VMware introduced the Virtual Volumes framework with vSphere 6.0. The integration of Virtual Volumes with IBM Spectrum Virtualize managed storage systems is dependent upon the vSphere application programming interfaces (APIs) for Storage Awareness (VASA). These APIs facilitate VM-related tasks initiated at the vSphere level to be communicated down to the storage system.

1.2.2 IBM Spectrum Control Base Edition

IBM support for VASA is provided by IBM Spectrum Control™ Base Edition (SCB). SCB is an out-of-band VASA Provider, which enables the communication between vSphere and the storage system along the control plane.

1.2.3 Storage system managed by IBM Spectrum Virtualize

IBM Spectrum Virtualize manages virtual machine disks (VVols) at the storage level. IBM Spectrum Virtualize-managed storage systems enable the flexible and dynamic provisioning of virtual machine storage that is required of a truly software-defined storage environment.

2

Prerequisites

This chapter contains information about the storage systems that are compatible with VMware Virtual Volumes (VVols), and lists the prerequisites that must be met before implementing Virtual Volumes in your environment. A reference guide to the software used when creating this document is also provided.

This chapter provides information about the following topics:

- 2.1, “Compatible storage systems” on page 6
- 2.2, “Prerequisites” on page 6
- 2.3, “Example configuration” on page 8

2.1 Compatible storage systems

At the time of writing, VVols can be implemented with the following storage systems:

- IBM SAN Volume Controller
- IBM Storwize Family
- IBM FlashSystem™ V9000

These storage systems support a single deployment of VVols across an entire storage infrastructure, by enabling users to virtualize third-party and existing storage under IBM Spectrum Virtualize, as demonstrated in Figure 2-1.

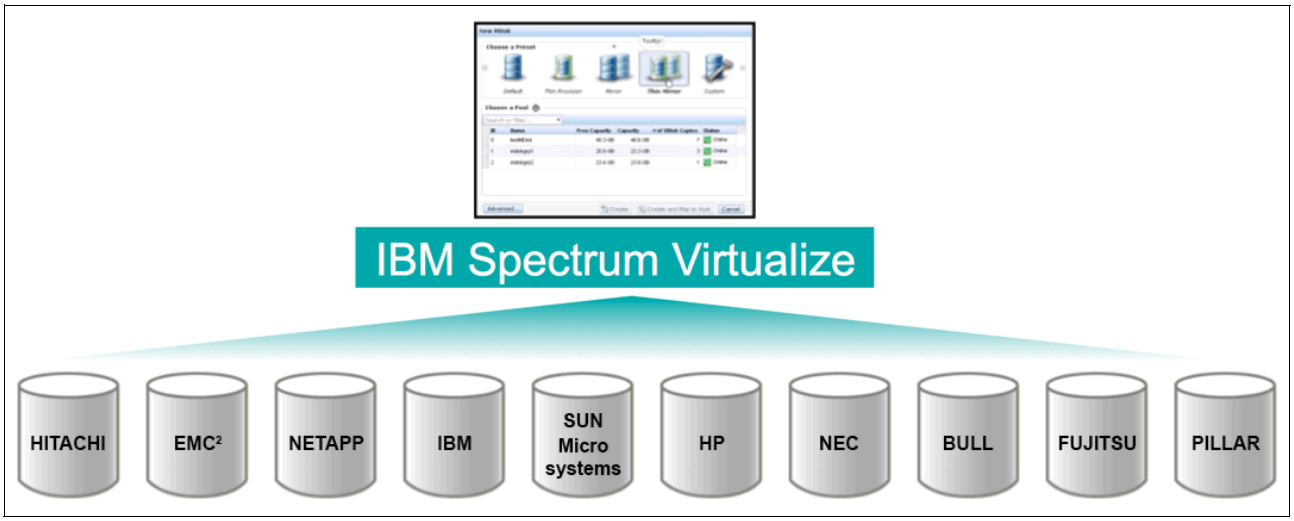

Figure 2-1 Virtualization of third-party and existing storage using IBM Spectrum Virtualize

2.2 Prerequisites

The following prerequisites must be met before attempting to implement Virtual Volumes in your environment.

2.2.1 Code levels

Three core software components are needed for a successful implementation of VVols. The minimum code level required for each component is provided in the following list:

- VMware vSphere 6.0 (or higher):
  - VMware ESXi version 6.0
  - VMware vCenter Server 6.0 or vCenter Server Appliance (vCSA) 6.0 (vSphere Web Client 6.0)
- IBM Spectrum Control Base Edition (SCB) version 2.2.1 (or higher)
- Compatible storage system running IBM Spectrum Virtualize version 7.6.0 (or higher)

Obtaining IBM Software

To download the current software for your storage system, browse for the relevant software on the following website:

https://ibm.biz/storage-system

Download the latest IBM Spectrum Control Base Edition installation package:

https://ibm.biz/ibm-scb

2.2.2 NTP Server Configuration

Ensure that Network Time Protocol (NTP) server is configured across the following components:

- Storage system
- IBM Spectrum Control Base Edition
- vSphere environment

NTP ensures that time settings are consistent across the entire VVols environment.

2.2.3 Server configuration for IBM Spectrum Control Base Edition

SCB can be deployed on a physical or virtual server. We suggest using a virtual machine, to take advantage of features such as VMware High Availability. For more information about this consideration, see 6.2.1, “High availability” on page 77.

At the time of writing, the server that hosts SCB must satisfy the following requirements:

- 64-bit dual-core processor (CPU)
- At least 4 gigabytes (GB) of random access memory (RAM)
- At least 16 GB of free disk space
- Operating system (OS): Red Hat Enterprise Linux (RHEL) 6.3-6.6 (x64)

For the current information about server requirements for IBM Spectrum Control Base Edition, see the release notes:

https://ibm.biz/ibm-scb-kc

The installation and configuration steps for IBM Spectrum Control Base Edition are covered in Chapter 4, “Installation and configuration” on page 17.

2.2.4 Licensing requirements

IBM does not require an additional license to enable Virtual Volumes on a storage system running IBM Spectrum Virtualize.

Remember: Users should be aware that VVols functionality uses IBM FlashCopy. Therefore, the appropriate FlashCopy license for your storage system should be applied before implementing VVols.

At the time of writing, the VVols feature is included as standard with VMware vSphere 6.0.

2.3 Example configuration

The following system configuration was used in the examples cited in this IBM Redbooks publication:

- VMware vSphere version 6.0u1:
  - VMware vCSA version 6.0u1
  - VMware ESXi version 6.0u1
- IBM Spectrum Control Base Edition version 2.2.1

  Deployed in an active/passive configuration on RHEL 6.5 (virtual machines)
- IBM Storwize V7000 Gen

  Powered by IBM Spectrum Virtualize, version 7.6.0

3

Terminology

This chapter provides an overview of the terminology that is used throughout this IBM Redbooks publication. We suggest familiarizing yourself with all of the terms introduced here, and then bookmarking this chapter for use as a quick reference guide as you proceed through the document.

The core concepts involved in the VMware Virtual Volumes (VVols) framework are introduced in the main body of the chapter. These sections provide the requisite information to read and understand this book. For a deeper technical explanation of a particular concept (for advanced or interested readers), see the shaded sections offering *Additional information*.

This chapter provides information about the following topics:

- 3.1, “IBM Spectrum Virtualize” on page 10
- 3.2, “Storage system” on page 10
- 3.3, “Parent pool” on page 10
- 3.4, “Child pool” on page 10
- 3.5, “Virtual Volumes” on page 10
- 3.6, “VVols-enabled host” on page 11
- 3.7, “Protocol endpoint” on page 12
- 3.8, “Bind and unbind operations” on page 13
- 3.9, “Utility volume” on page 14
- 3.10, “VASA Provider” on page 14
- 3.11, “IBM Spectrum Control Base Edition” on page 14
- 3.12, “High availability group” on page 14
- 3.13, “IBM Spectrum Control Base Edition user” on page 14
- 3.14, “Storage service” on page 15
- 3.15, “Storage resource” on page 15
- 3.16, “Storage space” on page 15
- 3.17, “Storage Container” on page 15
- 3.18, “VVol Datastore” on page 16
- 3.19, “Storage Policy-Based Management” on page 16

3.1 IBM Spectrum Virtualize

IBM Spectrum Virtualize is the name of the software used in combination with the following storage systems:

- IBM SAN Volume Controller
- IBM Storwize Family
- IBM FlashSystem V9000

3.2 Storage system

The term *storage system* is used throughout this document to refer to any VVol-compatible storage system running IBM Spectrum Virtualize code. For a list of compatible storage systems, see 2.1, “Compatible storage systems” on page 6.

3.3 Parent pool

A *parent pool* is a group of managed disks (MDisks), also known as Redundant Array of Independent Disks (RAID) arrays. This concept is often referred to simply as a *storage pool*. We explicitly state whether we are referring to a parent or child pool throughout this document, to avoid ambiguity.

Parent pools are created by the storage administrator using the storage system graphical user interface (GUI) or management command-line interface (CLI).

3.4 Child pool

A *child pool* is a logical quota of storage, which the storage administrator allocates from an existing parent pool. Unlike the parent pool, the capacity of a child pool is not directly tied to the underlying physical storage, and can consequently be resized without adding or removing drives or MDisks.

In the VVols environment, child pools are created and maintained using IBM Spectrum Control Base Edition (SCB).

Additional information: Creating child pools using SCB enables this server to “own” the child pools. For further information, see 3.13, “IBM Spectrum Control Base Edition user” on page 14.

3.5 Virtual Volumes

A *Virtual Volume* (VVol) refers to an individual volume, which corresponds to a virtual machine (VM) disk. For example, when deploying a new virtual machine with a 40-gibibyte (GiB) data disk (for example, a Windows VM’s `C:\` drive), a corresponding 40 GiB volume is created on the storage system.

Figure 3-1 depicts the status of a VVol within the context of the storage system.

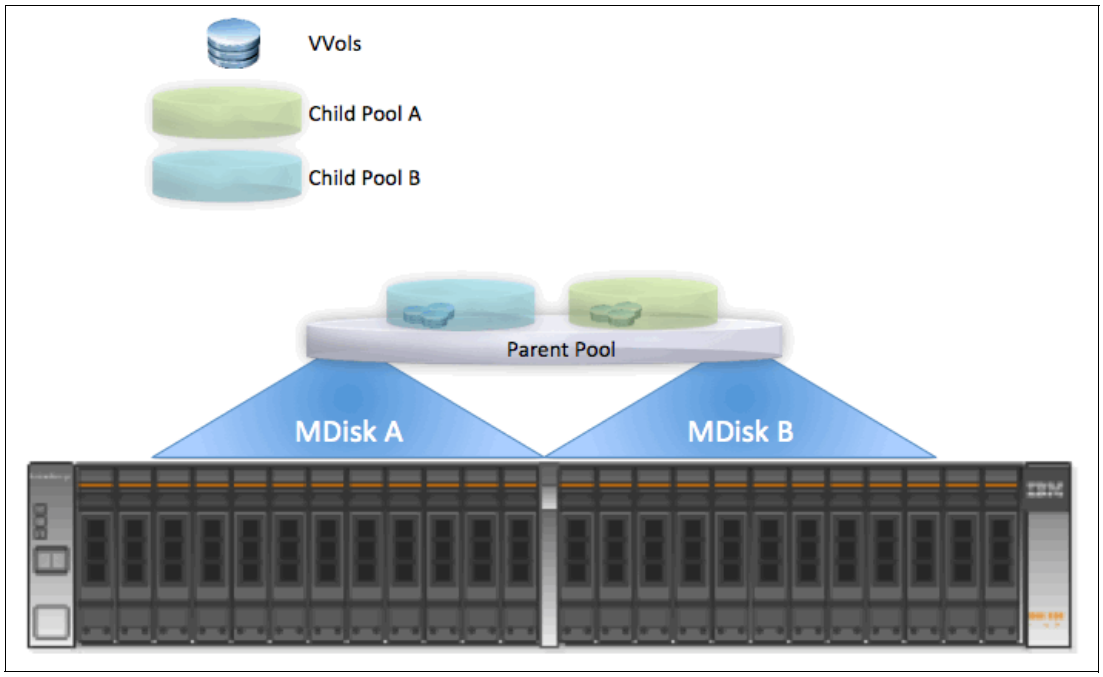

Figure 3-1 The relationship between parent pools, child pools, and VVols

The VVols framework is strikingly different from the widely used VMware file system (VMFS) model, where the storage corresponding to a virtual machine's 40 GiB data disk is embedded within the VMFS, and cannot be managed as an individual object by the storage system or the storage administrator.

In this document, we use the term *VVols* to differentiate these objects from regular volumes, because they are created and managed from within the VMware vSphere environment, rather than by directly interacting with the storage system.

VVols are automatically created when carrying out vSphere tasks, such as provisioning a VM. These tasks are initiated by the vSphere administrator.

Additional information: VVols are "owned" by the SCB server. For further information, see 3.13, "IBM Spectrum Control Base Edition user" on page 14.

3.6 VVols-enabled host

We use the term *VVols-enabled host* to describe a host that meets the following requirements:

- VMware ESXi (ESXi) version 6.0 or later installed.
- Registered with the storage system with host type set to `VVol`. Note that in the CLI, the `VVol` host type is referred to as `adminlun`.

Additional information: Setting the host type to `VVol` advertises protocol endpoints to a host. This step is essential for establishing a data path between a host and a VVol.

Consequently, changing the host type to `VVol` does *not* affect existing host mappings. It simply enables VVols functionality in addition to existing connectivity.

3.7 Protocol endpoint

ESXi hosts use *protocol endpoints* (PEs) to establish a data path to a VVol. Protocol endpoints act as a query interface, and can connect to potentially hundreds or thousands of VVols.

This design entails that, unlike conventional volumes, VVols do not need to be manually mapped to hosts. Instead, SCB binds a VVol to a specific protocol endpoint by creating a unique form of host mapping, known as a *subsidiary volume hostmap*. After the binding operation is complete, ESXi hosts can access VVols for input/output (I/O) operations using a PE.

Figure 3-2 depicts the role of protocol endpoints in the Virtual Volumes I/O stack.

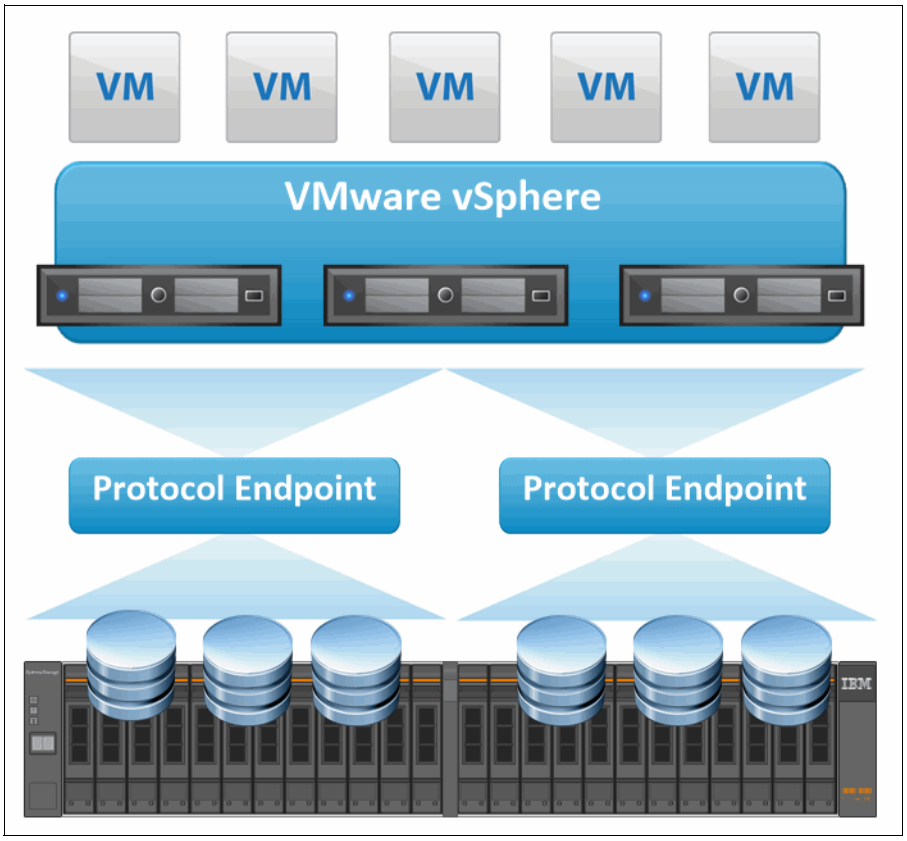

Figure 3-2 The data path in a Virtual Volumes environment

Protocol endpoints are generated internally by the storage system, and are automatically discovered by ESXi hosts when a VVol datastore is mounted. Consequently, they do not require any manual setup from the storage administrator or the vSphere administrator.

On the vSphere Web Client, protocol endpoints become visible after a VVol Datastore is mounted by an ESXi host. You can view this information by performing the following steps:

1. Navigate to Hosts and Clusters.
2. Select the relevant ESXi Host.
3. Select **Manage** → **Storage** → **Protocol Endpoints**, as shown in Figure 3-3.

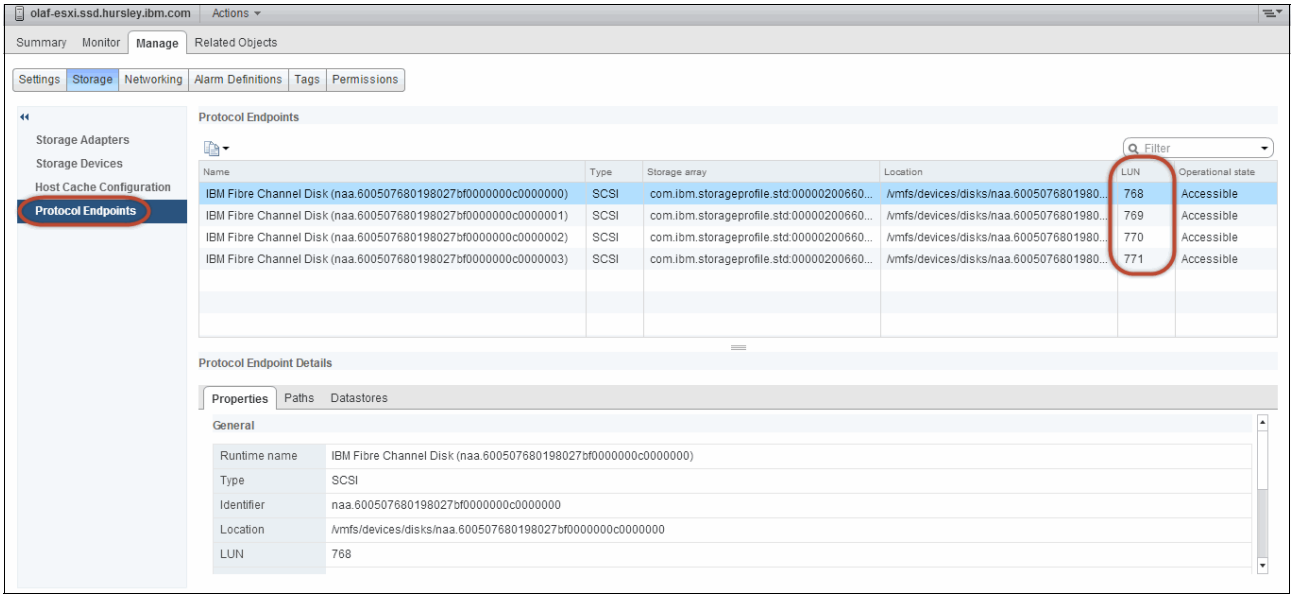

Figure 3-3 An example of the Protocol Endpoints tab in the vSphere Web Client

Additional information: On storage systems powered by IBM Spectrum Virtualize, a pair of protocol endpoints is created for every I/O group in the cluster. There is typically one protocol endpoint per node in the system.

Protocol endpoints are part of the storage fabric, and are visible to a VVols-enabled host down every path to the storage system. PEs are presented to ESXi hosts as Small Computer System Interface (SCSI) IDs 768 - 775, as shown in Figure 3-3.

3.8 Bind and unbind operations

ESXi hosts issue bind operations with vSphere application programming interfaces (APIs) for Storage Awareness (VASA), using IBM Spectrum Control Base Edition, to connect VVols to a protocol endpoint. When a host does not require a binding to a VVol (for example, when a VM is powered-off), an unbind command is issued. Bind and unbind operations are initiated by ESXi hosts, and do not require manual steps from storage or vSphere administrators.

Additional information: The bind and unbind operations sent by an ESXi host translate into the creation and deletion of subsidiary volume host maps on the storage system.

3.9 Utility volume

A *utility volume* is created on the storage system as part of the VVols configuration. This utility volume contains a database that is created and maintained by the IBM Spectrum Control Base Edition instance. The metadata stored in the database is critical to IBM Spectrum Control Base Edition operations and the VVols environment. If possible, it is advised to store a mirrored copy of the utility volume for redundancy.

Storage administrators create the utility volume when they enable VVols on their storage system.

3.10 VASA Provider

VASA Provider (VP) is the generic name given to an out-of-band software component that provides communications between vSphere components (ESXi hosts and vCenter) and the storage system. Each storage vendor who supports VVols implements their own VASA Provider. IBM Spectrum Control Base Edition is the VASA Provider for IBM.

3.11 IBM Spectrum Control Base Edition

IBM Spectrum Control Base Edition is the VASA Provider for storage systems running IBM Spectrum Virtualize code. SCB is external to the storage system, and is installed on a physical or virtual server running Red Hat Enterprise Linux (RHEL) V6.3 and later. We often refer to this server as the *SCB server*.

SCB is typically managed by the storage administrator.

3.12 High availability group

A *high availability group* (HA group) contains two or more IBM Spectrum Control Base Edition instances, which manage the same storage systems and database. Configuring an HA group ensures minimal downtime of virtual machine management if the active SCB instance experiences an outage. For further explanation of SCB HA, see 6.2.1, “High availability” on page 77.

3.13 IBM Spectrum Control Base Edition user

The SCB server requires a user account on the storage system, to satisfy VASA requests at the storage level. The IBM Spectrum Control Base Edition User has the *VASA Provider* security role, which allows for the ownership of VVols-specific objects. VASA Provider ownership prevents other users defined on the system from inadvertently disrupting your VVols environment.

3.14 Storage service

Storage services are created using SCB. Storage administrators use services to define the set of capabilities (for example, encryption) that are offered to the vSphere administrator, for Storage Policy-Based Management (SPBM).

3.15 Storage resource

In the VVols environment, *storage resources* are allocated from the storage system in the form of a child pool.

3.16 Storage space

Storage spaces are created using SCB. Storage administrators combine information about storage resources (allocated as child pools) and storage capabilities (defined as storage services) into a logical construct known as a *storage space*. This is presented to vSphere as a *Storage Container*.

3.17 Storage Container

Storage Containers are logical entities that are presented to vCenter for use as a VVol Datastore. A Storage Container (VMware terminology) directly correlates to a storage space (IBM terminology).

3.18 VVol Datastore

A *VVol Datastore* is the container for VVols. It is directly mapped to a storage space. VVol Datastores are typically created by the vSphere administrator.

Figure 3-4 shows the relationship between VVol datastores and storage spaces. In our example, the storage space encompasses a child pool and an encrypted storage service.

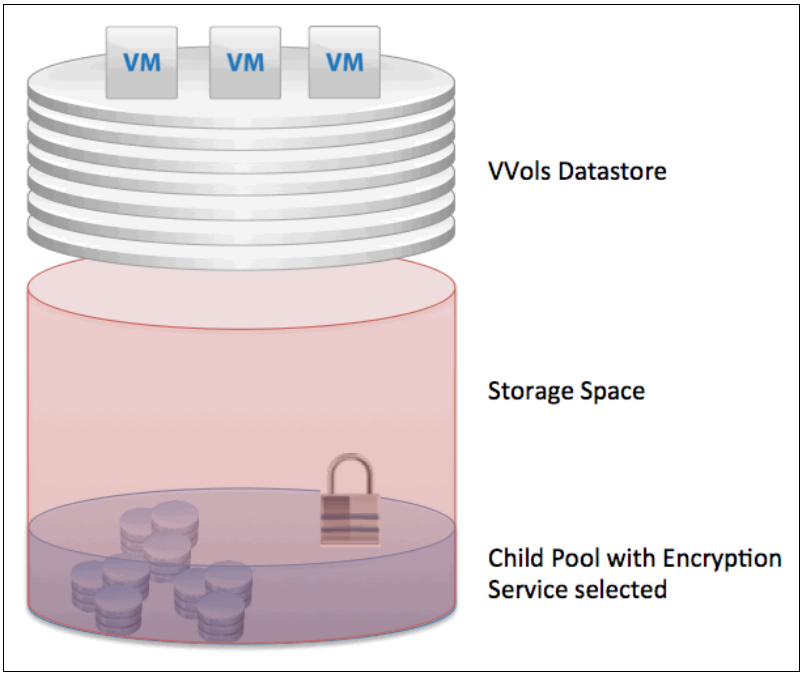

Figure 3-4 The one-to-one mapping between a VVol Datastore and a storage space

Additional information: Unlike a conventional VMFS Datastore, a VVol Datastore does not require a file system.

3.19 Storage Policy-Based Management

Storage Policy-Based Management (SPBM) refers to the ability to create policies in vSphere that help to select the appropriate storage resource (VVol Datastore) when provisioning and managing virtual machines. This introduces opportunities for automation, and optimizes the provisioning of virtual machine storage at scale.

The capabilities offered when defining storage services are essential to SPBM, and should be discussed between vSphere and storage teams when configuring IBM Spectrum Control Base Edition.

Additional information: For more information about SPBM, see Chapter 7, “Storage Policy-Based Management” on page 83.

4

Installation and configuration

Chapter 1, “Implementing VMware Virtual Volumes for systems powered by IBM Spectrum Virtualize” on page 1 described the three key components involved in implementing the VMware Virtual Volumes (VVols) management framework. You might find it helpful to refer to Figure 1-1 on page 2 from time to time while you familiarize yourself with the new terminology and how it fits your current environment.

For the sake of simplicity, we approach the installation and configuration of our VVols environment by singling out each of these components. This chapter starts with the storage system, then moves onto IBM Spectrum Control Base Edition (SCB), and finally touches on adding a VVols Datastore to your vSphere environment.

This chapter provides information about the following topics:

- 4.1, “Configuring the storage system” on page 18
- 4.2, “Configuring IBM Spectrum Control Base Edition” on page 22
- 4.3, “Configuring VMware vSphere Web Client” on page 35
- 4.4, “Configuring additional SCB servers for high availability” on page 39

4.1 Configuring the storage system

During this chapter, you perform several steps required to enable VVols on your storage system. See the prerequisites in Chapter 2, “Prerequisites” on page 5 to ensure that your environment meets the minimum requirements.

Requirement: Parent pools must already be configured on the storage system before attempting to enable VVols.

4.1.1 Enabling VVols

To enable VVols, complete the following steps:

1. In the management GUI for your storage system, go to **Settings** → **System** and select **VVols**, as depicted in Figure 4-1.

   Figure 4-1 shows you where to find the control to enable VVols on your storage system.

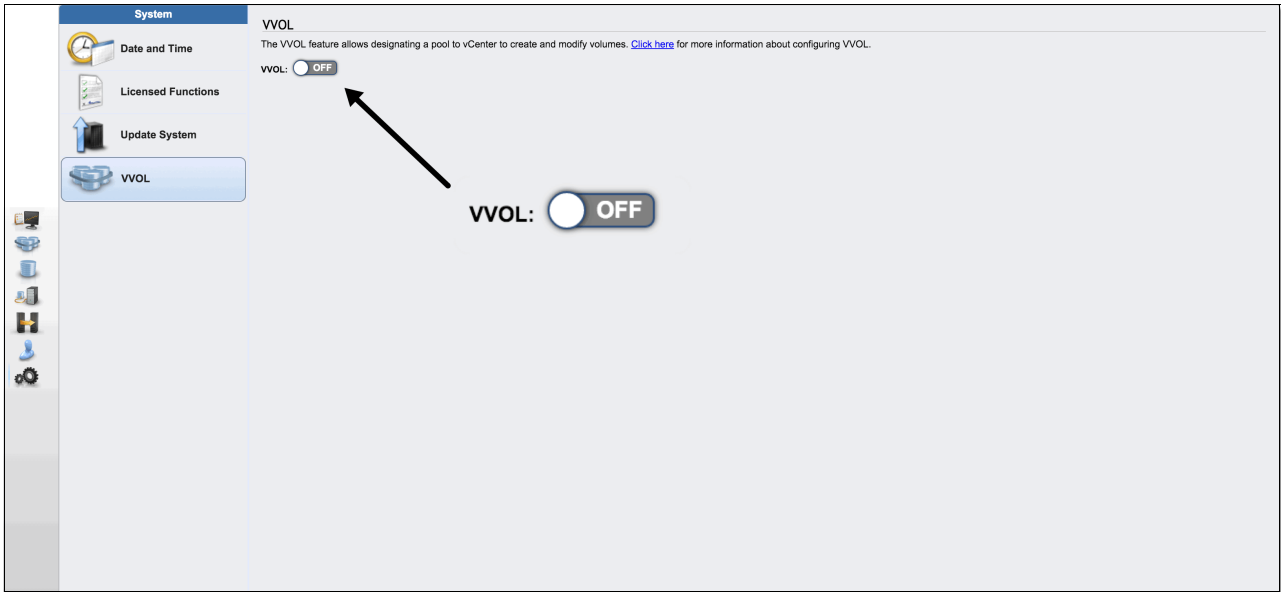

Figure 4-1 Storage system toggle the VVols management GUI to On

Tip: If you are unable to toggle VVols to **ON**, ensure that you have at lease one parent pool configured on your storage system.

2. Start the Virtual Volumes configuration by clicking the toggle switch. You are then presented with a short wizard that walks you through the steps required to complete the setup of VVols on your storage system.

Figure 4-2 shows the Enable VVols dialog. Here we select the parent pool on which to locate our utility volume, and create our user for SCB.

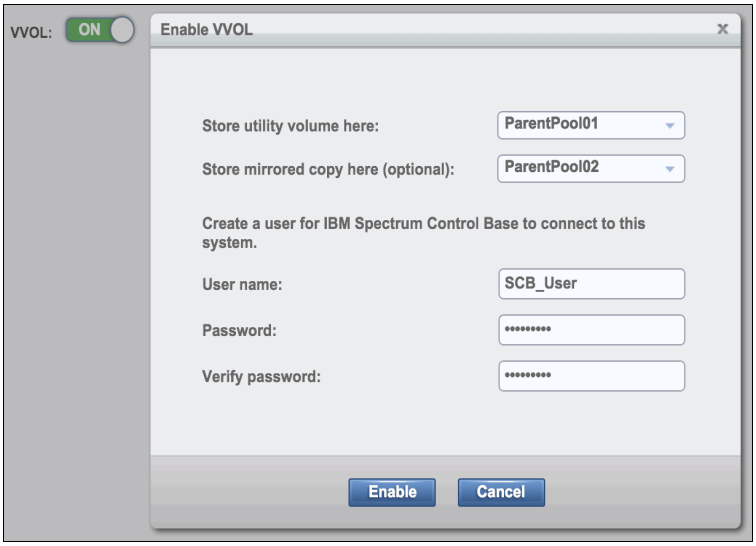

Figure 4-2 Storage system Enable VVols dialog box

3. Select an available parent pool on your storage system for your utility volume. As an optional task, select an additional parent pool where you want your utility volume to be mirrored, as illustrated in Figure 4-3.

Figure 4-3 Storage system parent pool selection box

Note: The utility volume contains a database that's critical to your VVols environment. We suggest mirroring your utility volume to a second parent pool that is in a separate failure domain for redundancy. For example, use a parent pool that is made from managed disks (MDisks) that are presented from different storage systems or a different input/output (I/O) group.

4. Create a user for IBM Spectrum Control Base Edition to communicate with the storage system. In this step, you are creating a user account that is automatically added to the security role `VASAProvider`, as shown in Figure 4-4.

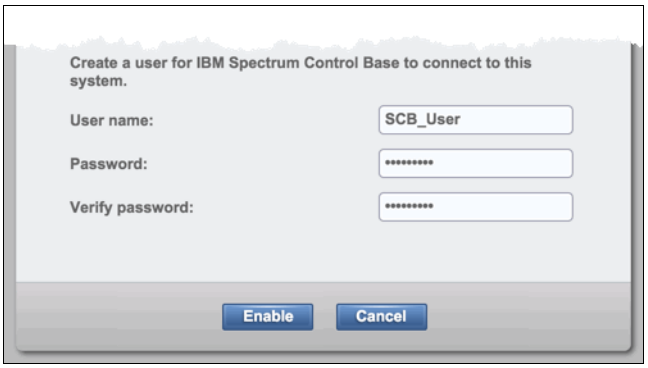

Figure 4-4 Storage system SCB user creation dialog

Tip: You will later require the SCB user information when setting up IBM Spectrum Control Base Edition.

5. Click **Enable** to continue to the next step where you will receive confirmation that VVols has been enabled, as shown in Figure 4-5.

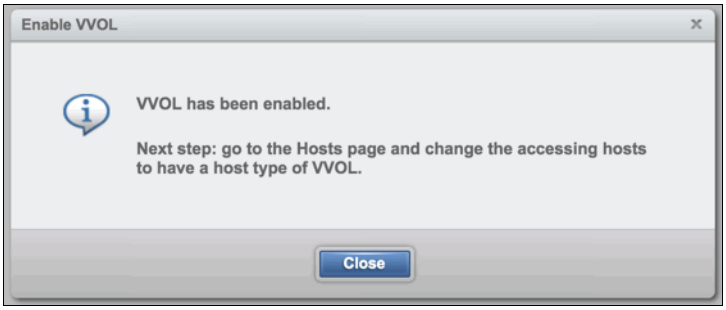

Figure 4-5 VVols enabled message

Important: The vSphere APIs for Storage Awareness (VASA) Provider role is strictly reserved for use by the IBM Spectrum Control Base Edition server. Users should *not* log in to the graphical user interface (GUI) or command-line interface (CLI) as a user with `VASAProvider` privileges. Users must not perform *any* manual tasks as the `VASAProvider` user, unless specifically advised to by IBM Support.

4.1.2 Enabling VMware ESXi hosts for VVols

Existing ESXi hosts on your storage system need to be enabled for VVols before you can use Virtual Volumes in the vSphere Web Client. If you are adding new ESXi hosts to your storage system, the same VVols enablement steps apply. See 5.2.1, “Managing VMware ESXi hosts” on page 53 for more information about adding hosts.

To enable ESXi hosts, complete the following steps:

1. Enable VVols on your existing hosts by selecting the ESXi host object from the Hosts view and then select **Action** → **Properties** → **Edit**. Here you can change the host type to `VVols`, as depicted in Figure 4-6. You need to repeat these steps for each ESXi host you want to enable for VVols.

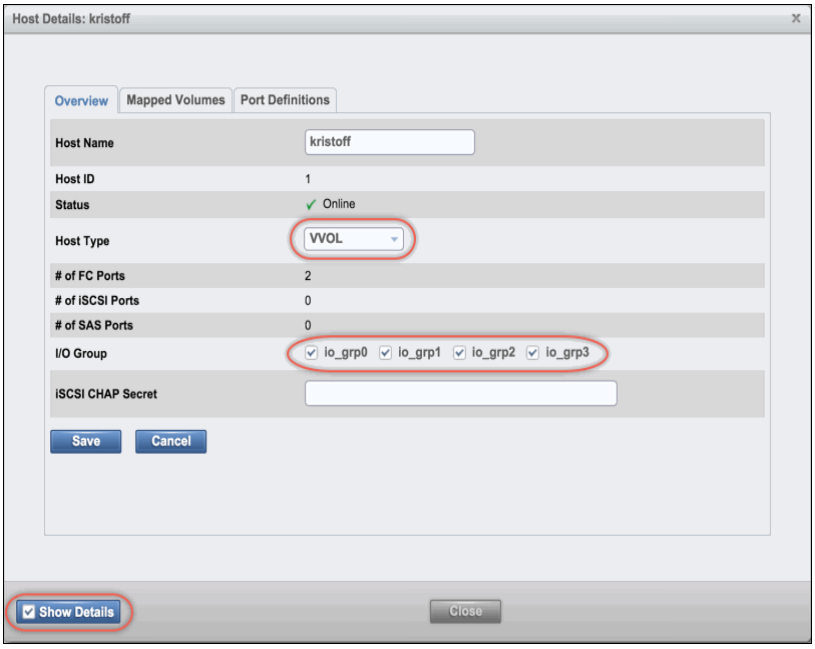

Figure 4-6 Storage system enable host type VVols

Tip: Existing hostmap configurations are unaffected by enabling your ESXi hosts for VVols. However, you might need to enable host access to all I/O groups if not already configured.

2. To add an ESXi host to the system (optional), select the **Hosts** view and select **Add Host** to start the Add Host dialog, as pictured in Figure 4-7 on page 22. Perform the following steps to complete the configuration:
   a. Select the connectivity type for the host.
   b. Enter a name for the host.
   c. Input the port details.
   d. Click **Advanced** to show more options.
   e. Open the Host Type menu. Select **VVol**.
   f. Select **Add** to complete the task.

Figure 4-7 shows the Add Host dialog.

Figure 4-7 Storage system Add Host dialog

The steps required for setting up your storage system have now been completed. Move onto the next section to configure the IBM Spectrum Control Base Edition instance.

4.2 Configuring IBM Spectrum Control Base Edition

IBM Spectrum Control Base Edition plays a pivotal role in the communication between your vSphere environment and the storage system. This section guides you through the installation and configuration options for installing IBM Spectrum Control Base Edition.

Note: It is assumed that you already have a virtual machine (VM) provisioned for the installation of IBM Spectrum Control Base Edition. For more information about the configuration of this VM, see 2.2.3, “Server configuration for IBM Spectrum Control Base Edition” on page 7.

4.2.1 Downloading the IBM Spectrum Control Base Edition installation files

Obtain the current IBM Spectrum Control Base Edition installation files from IBM Fix Central:

https://ibm.biz/ibm-scb

Extract them to a temporary folder, for example, `/scb_install`, on your IBM Spectrum Control Base Edition virtual machine, hereafter referred to as *SCB*. Example 4-1 shows the typical installation files for SCB after they have been extracted.

Note: Downloading the installation package from a trusted, Secure Sockets Layer (SSL)-protected resource, such as Fix Central, ensures authenticity and integrity. If you are in any doubt as to the authenticity of the package, see the IBM Spectrum Control Base Edition User Guide for instructions about using the Signing Key to validate the package. You can obtain the user guide from the previous link to Fix Central.

Example 4-1 SCB installation files

```
ibm_spectrum_control-2.2.1-xxxx-x86_64.bin
nginx-1.6.2-1.el6.ngx.x86_64.rpm
postgresql92-9.2.10-1PGDG.rhel6.x86_64.rpm
```

```
postgresql92-contrib-9.2.10-1PGDG.rhel6.x86_64.rpm
postgresql92-libs-9.2.10-1PGDG.rhel6.x86_64.rpm
postgresql92-server-9.2.10-1PGDG.rhel6.x86_64.rpm
```

4.2.2 Installing IBM Spectrum Control Base Edition

Tip: Now would be a great time to ask your vSphere administrator to clone your SCB VM to a template. You can later use the template to deploy further instances of SCB server as part of a high availability (HA) group.

To install SCB, complete the following steps:

1. Change to the folder on your SCB server where you previously extracted the installation files and start the installation of SCB.
2. On the command line, type `rpm -iv *.rpm`. Sample output is shown in Example 4-2.

*Example 4-2 Example output of command rpm -iv *.rpm*

```
warning: nginx-1.6.2-1.el6.ngx.x86_64.rpm: Header V4 RSA/SHA1 Signature, key ID
7bd9bf62: NOKEY
warning: postgresql92-9.2.10-1PGDG.rhel6.x86_64.rpm: Header V4 DSA/SHA1
Signature, key ID 442df0f8: NOKEY
warning: uuid-1.6.1-10.el6.x86_64.rpm: Header V3 RSA/SHA256 Signature, key ID
fd431d51: NOKEY
Preparing packages for installation...
postgresql92-libs-9.2.10-1PGDG.rhel6
postgresql92-9.2.10-1PGDG.rhel6
uuid-1.6.1-10.el6
postgresql92-contrib-9.2.10-1PGDG.rhel6
postgresql92-server-9.2.10-1PGDG.rhel6
nginx-1.6.2-1.el6.ngx
```

3. Next, we turn our attention to installing the `.bin` file. At the command line, type `./ibm_spectrum_control-2.2.1-xxxx-x86_64.bin` (note that you may have a later file version).
4. When prompted, enter `1` to accept the agreement. After a minute or two, you receive a `Installation completed successfully` message, and you can move on to the next step.

4.2.3 Changing the IBMSC administration account

As part of the installation process, a user account named `IBMSC` is created for administration purposes. To avoid unauthorized access to the SCB server, it is strongly advised to change the default password. To modify the IBMSC user account, type `passwd ibmsc` on the command line and enter a new password.

This is the last task we need to run at the command line. All further steps are run from the SCB server web interface.

Suggestion: Configuring a second SCB instance is required for redundancy (advised). This is described in more detail in 4.4, “Configuring additional SCB servers for high availability” on page 39.

4.2.4 Accessing the IBM Spectrum Control Base Edition web interface

To access the SCB web interface, complete the following steps:

1. Open a web browser and connect to your SCB server's web interface using the following URL: `https://<SCB Server Name>:8443`. You might need to *trust* the certificate before continuing to the logon page, as shown in Figure 4-8.

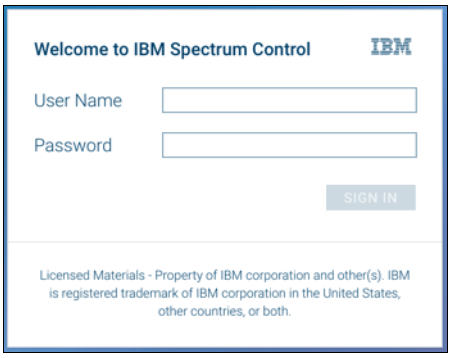

Figure 4-8 IBM Spectrum Control Base Edition web interface logon

2. Log in to the SCB web interface with the following default credentials:
   - **User Name**: `admin`
   - **Password**: `admin1!`

4.2.5 An overview of the SCB web interface

The first time you log in to the SCB web interface you are presented with the Welcome window. Click past this and you are presented with the SCB home window, as shown in Figure 4-9. Take some time to familiarize yourself with the necessary components and views to continue the configuration of SCB.

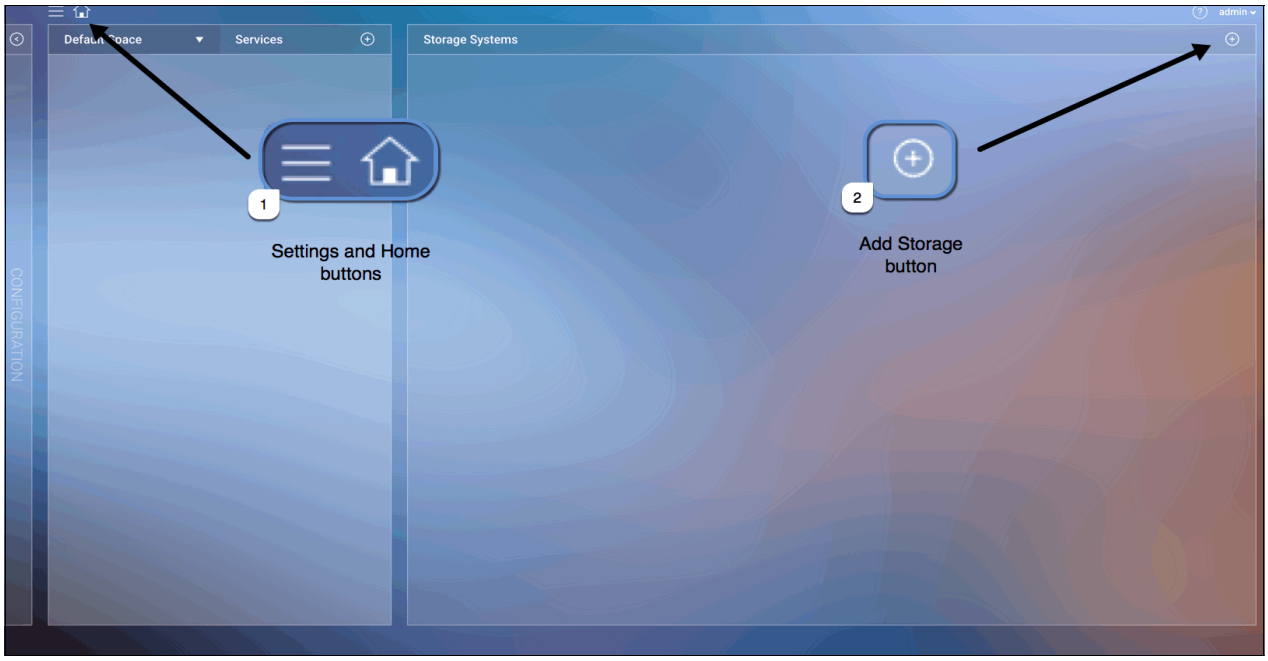

Figure 4-9 IBM Spectrum Control Base Edition Home window

The SCB web interface consists of four management panes. In the default Home view, you are presented with the Default Space and Services pane on the left and the Storage Systems pane on the right, as pictured in Figure 4-9 on page 24. For clarity, we have indicated the locations of the **Settings** and **Add Storage** buttons.

Note: The remaining two panes for Applications and Monitoring are not required for configuring Virtual Volumes, and are outside the intended scope of this book.

Table 4-1 summarizes the functionality of the SCB GUI elements.

Table 4-1 IBM Spectrum Control Base Edition GUI Elements

GUI Element	Description
	This is the **Settings** button. Use this button to access the Settings menu.
	This is the **Home** button. Click this button to display the Home window (Spaces, Services, and Storage Systems view).
	This is the **Add** button. Click this button to add new objects, such as storage systems and storage services.
	This is the **Edit** button, which is displayed when selecting or highlighting an object. Click this button to configure the object (system, pool, server, and so on) or remove it.
	This is the **Remove** button. Click this button to remove a storage element or delete a user from the User list.
	This is the **Help** button. Click this button to display the Welcome page, which is also available after login.
	These are the right and left navigation pointers. They indicate if additional panes are available to the right or to the left of the current pane.
	This is the **Resource Attach** button. It is available when an unattached storage resource is selected in a Storage Resources table. Click this button to attach the storage resource to a previously selected storage service.
	This is the **Resource Detach** button. It is available when an attached storage resource is selected in a Storage Resources table. Click this button to detach the storage resource from a previously selected storage service.

During this chapter, we configure the following items, in order, using the SCB web interface:

1. Changing the default administrator password
2. Storage credentials
3. VASA credentials
4. General settings
5. Server certificate
6. Adding a storage system
7. Adding a storage space
8. Adding a storage service
9. Defining and attaching storage resources

4.2.6 Changing the default administrator password

One of the first steps you want to take is to change the default password for the administrator account:

1. Click the **Settings** button and select **Users** in the menu. You are presented with the Users section box. Highlight the **admin** user account by clicking it. Two small buttons, **Edit User** and **Delete User**, appear to the right, as shown in Figure 4-10.

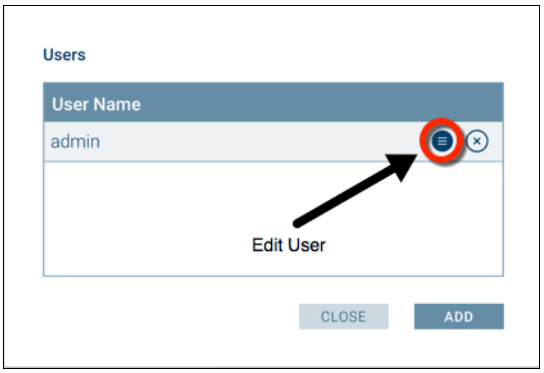

Figure 4-10 SCB Users selection box

2. Click the **Edit Users** button to modify the **admin** user properties, as shown in Figure 4-11.

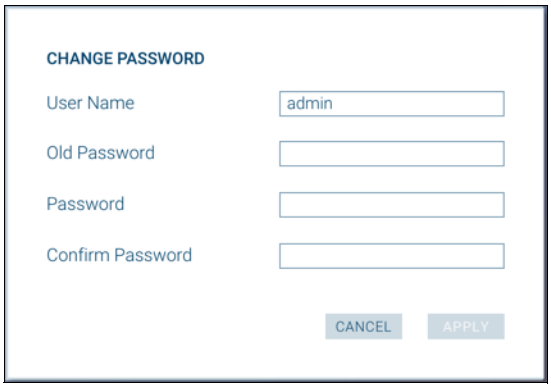

Figure 4-11 SCB User dialog

3. Enter the existing password of `admin1!` and then enter a new password for your admin user. Click **APPLY** → **Close**.

4.2.7 Storage credentials

Earlier in this chapter (Figure 4-4 on page 20), we created a user on our storage system for use by SCB. By setting these credentials in SCB, we enable communication to the storage system in order to perform various tasks, one of which is adding the storage system into SCB:

1. Click the **Settings** button and select **Storage credentials**, as shown in Figure 4-12.

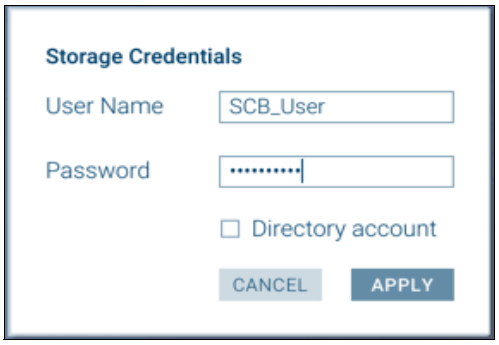

Figure 4-12 SCB Storage Credentials dialog

2. Enter the User Name and Password of the SCB user account created on the storage system.
3. Click **APPLY**.

> **Note:** The storage credentials must exist on all of the storage systems that you intend to connect to your IBM Spectrum Control Base Edition server.

4.2.8 VASA credentials

The VASA credentials are defined in SCB and used by the VMware vSphere Web Client to connect to the SCB server. To configure VASA, complete the following steps:

1. Click **Settings** → **VASA Credentials**, as shown in Figure 4-13.

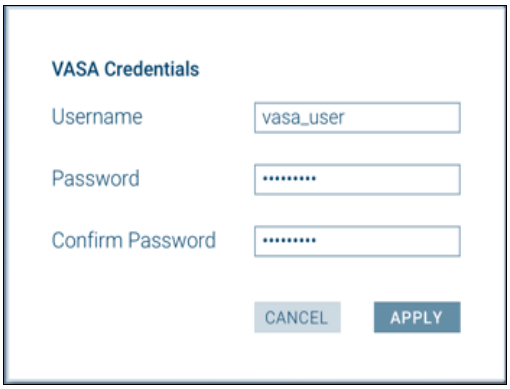

Figure 4-13 SCB VASA Credentials dialog

2. Enter the Username and Password, then click **APPLY**.

> **Tip:** The VASA user credentials are required later when adding the SCB storage provider in the vSphere Web Client.

4.2.9 General settings

This step is of particular importance to the configuration of high availability for our VVols environment. To enter the general settings, complete the following steps:

1. Click the **Settings** button and select **General Settings** on the Settings menu, as shown in Figure 4-14.

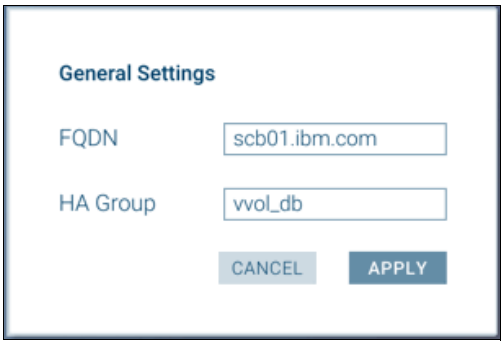

Figure 4-14 SCB General Settings

2. Enter the fully qualified domain name (FQDN) of your IBM Spectrum Control Base Edition server.
3. Enter the name for the HA Group. In this example, we have called it `vvol_db`.
4. Click **APPLY**.

Remember: High availability (HA) of SCB relies on the HA Group. See 6.2.1, “High availability” on page 77 for further details about HA groups and the use of multiple SCB servers for redundancy.

4.2.10 Server certificate

During the installation of IBM Spectrum Control Base Edition, an SSL certificate is generated with generic information. To provide secure communication between your SCB server and the vSphere Web Client, you need to generate a new certificate with the FQDN of your SCB server. Figure 4-15 shows the Server Certificate dialog for the test SCB server.

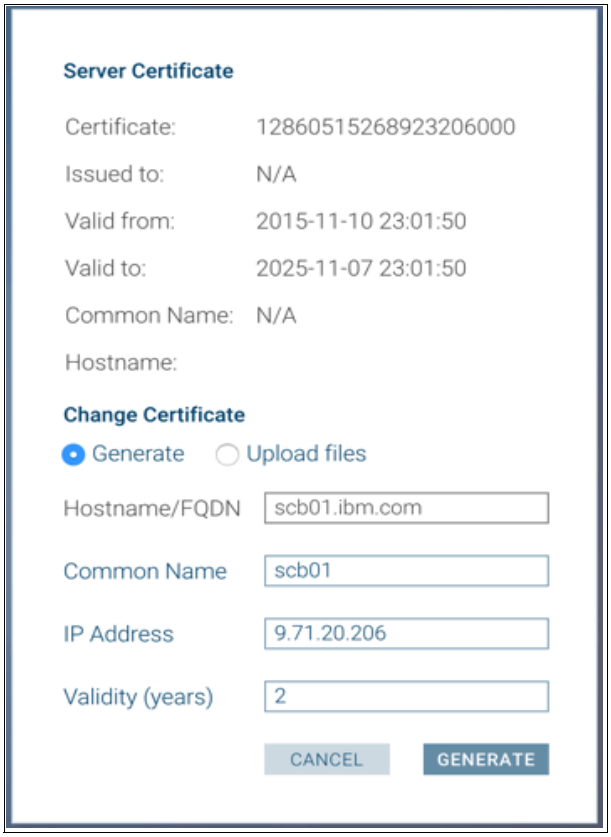

Figure 4-15 SCB Server Certificate dialog

1. Click the **Settings** button and select **Server certificate** on the Settings menu.
2. Enter the FQDN of your SCB server.
3. Enter a Common Name for your SCB server.
4. Enter the IP Address of your SCB server.
5. Choose a certificate Validity period.
6. Click **GENERATE**.
7. You receive an alert asking you to refresh your browser, as shown in Figure 4-16.

Figure 4-16 SCB reload web page

Tip: If you've already been issued with a trusted certificate that you want to use, you can replace the self-signed certificate with an existing trusted certificate. For more information, on this process, see the IBM Knowledge Center for IBM Spectrum Control Base Edition:

`https://ibm.biz/ibm-scb-kc`

4.2.11 Adding a storage system

Adding a storage system to IBM Spectrum Control Base Edition relies on having correctly entered the storage credentials. If you have not yet configured your storage credentials, see 4.2.7, “Storage credentials” on page 27. To add a storage system, complete the following steps:

1. Click the **Add** button on the Storage Systems pane, as shown in Figure 4-17. See 4.2.5, “An overview of the SCB web interface” on page 24 if you are unsure as to the location of the **Add** button.

Figure 4-17 SCB Add New Array dialog

2. Enter the Management IP address or Hostname of the storage system and select the storage system Type. In our example, we are using an IBM Storwize V7000 storage system with a Hostname of `cheech-cl`.
3. Click **ADD**. If the credentials are correct, the storage system is added to the Storage Systems pane, as depicted in Figure 4-18 on page 31.

Note: If you receive an `Unable to connect to the storage array due to a credentials error`, verify the user name and password entered earlier in step 4.2.7, “Storage credentials” on page 27. Remember that *both* the user name and password are case-sensitive.

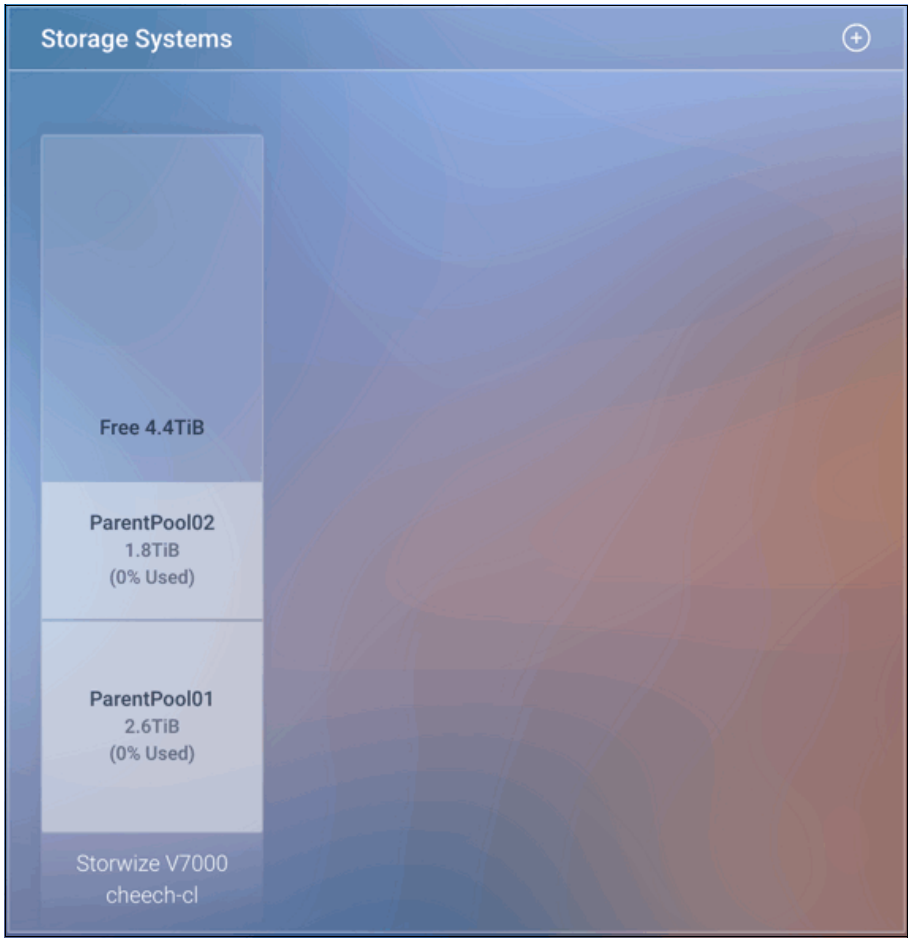

Figure 4-18 IBM Spectrum Control Base Edition: Storage Systems pane with a storage system

4.2.12 Adding a storage space

Creating a storage space in IBM Spectrum Control Base Edition enables you to present storage resources to the vSphere Web Client. For more information about storage spaces, see 6.2.2, “Defining storage spaces” on page 78.

To add a storage space, complete the following steps:

1. In the Spaces and Services pane, click the **Spaces** drop-down and select **Add New Space**, as shown in Figure 4-19.

> **Note:** The drop-down list for storage space selection shows the name of the currently selected storage space, and the default is **Default Space**.

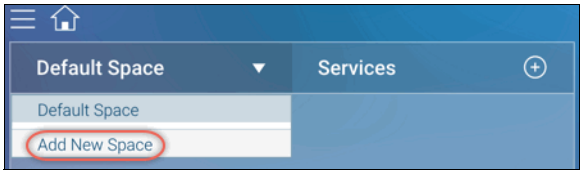

Figure 4-19 SCB adding a storage space

2. Enter a name and description for the new storage space and click **APPLY**. In our example, we chose to name it `SC01` (storage container). See Figure 4-20.

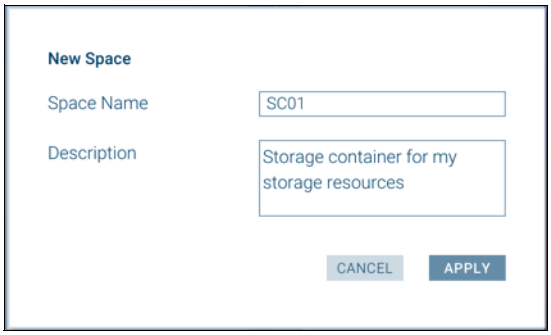

Figure 4-20 SCB configure New Space

3. You can confirm that you have successfully added your new storage space by again clicking the drop-down **Spaces** list and selecting your new storage space, as shown in Figure 4-21.

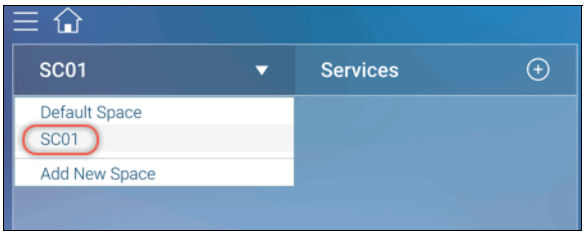

Figure 4-21 SCB spaces list

4.2.13 Adding a storage service

Now that you have created a storage space, the next step is to define the service or services that the space will be offering up to the vSphere Web Client. Essentially, a storage service advertises the capabilities of our storage resources, such as encryption, to the vSphere Web Client through the storage space. More information about the configuration and use of storage services is provided in 6.2.2, “Defining storage spaces” on page 78.

To add a storage service, complete the following steps:

1. Click the drop-down **Spaces** list and select your new storage space, as shown in Figure 4-21.
2. Click the **Add** button on the Storage and Services pane, as indicated in Figure 4-22.

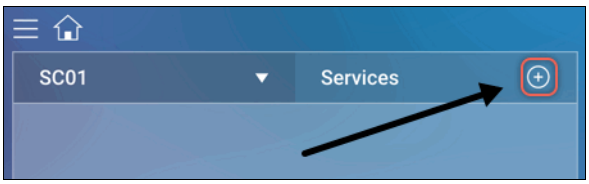

Figure 4-22 SCB Add service button

3. Enter a name and description for your new service. In the following example, we called the storage service `SC01_Service`. *Ensure that you have selected the box for VVol Service* (Figure 4-23).

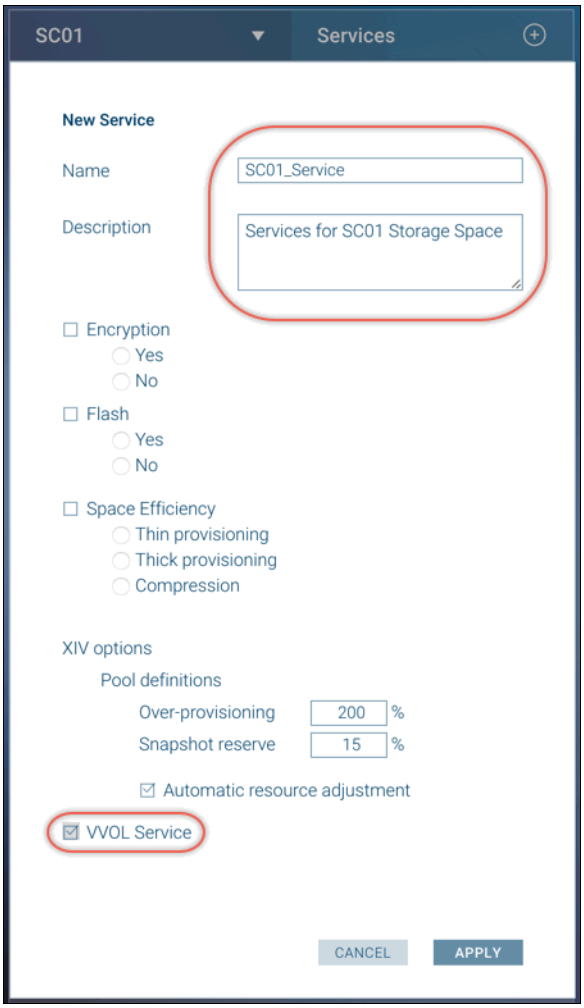

Figure 4-23 SCB New storage service configuration

4. Click **APPLY**. You should now have a *storage space* with a *storage service* configured much like that shown in Figure 4-24.

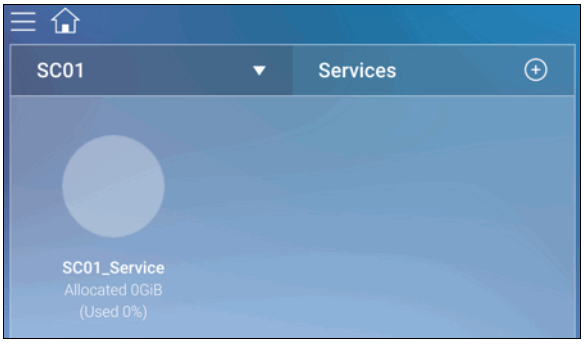

Figure 4-24 SCB Storage and Services view

Note: In our example, we create a simple VVols-enabled storage service. With *Storage Policy-Based Management* (SPBM), we look at configuring the storage services to provide a more granular and automated VVols environment. See Chapter 7, “Storage Policy-Based Management” on page 83 for more information.

4.2.14 Defining and attaching storage resources

Now that we have a *storage space* (`SC01`) with a *storage service* (`SC01_Service`) defined, you can add storage resources (child pools) to your storage space:

1. Click the drop-down **Spaces** list and select the *storage space* that you want to add storage resources to. In this example, we are using `SC01`. The VVols-enabled *storage service* that you previously configured for the *storage space* is displayed in the pane directly below the **Spaces** drop-down.
2. Click the *storage service*. The storage service becomes highlighted to indicate selection. See step 1 in Figure 4-25.
3. On the Storage Systems pane, hover over the storage system to display the **Edit** button, as depicted in step 2 in Figure 4-25.
4. Click the **Edit** → **Modify**, as shown in step 3 of Figure 4-25. The Array Settings dialog box is displayed.

   Figure 4-25 shows the steps involved in adding a storage resource to a storage service.

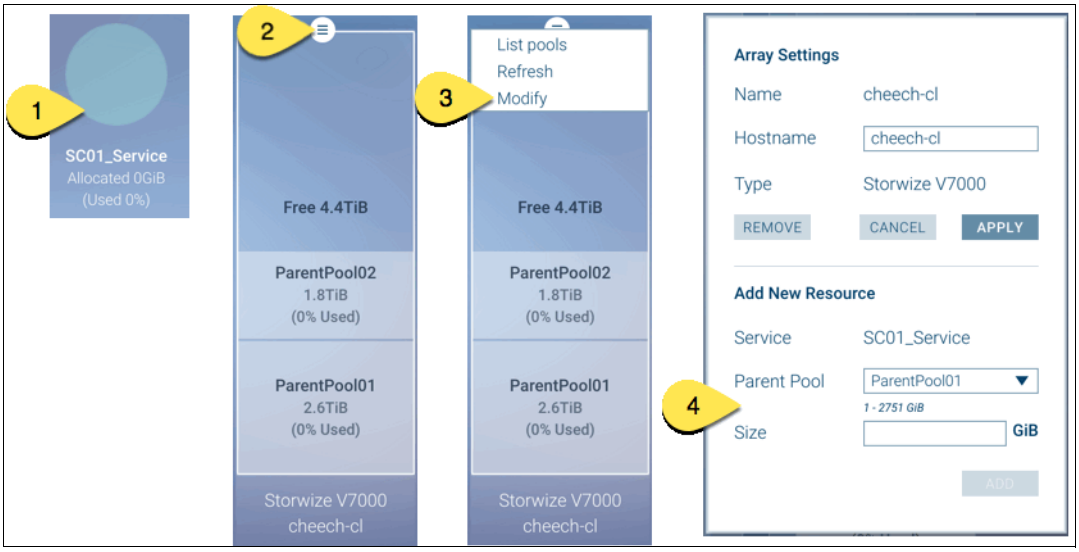

Figure 4-25 SCB adding a storage resource

5. In the Add New Resource section of the Array Settings dialog, select a parent pool and specify a size in gibibytes (GiB) for the new storage resource (child pool).
6. Click **Add** to complete the Add New Resource task. The storage resource is created on the storage system (as a child pool) and attached to the service.

Summary: In our example, we have created a 2500 GiB storage resource (child pool) on storage space `SC01` with the associated storage service `SC01_Service`. This resource has been taken from `ParentPool01`. This, in turn, creates a child pool on our storage system with an automatically generated name, as shown in Figure 4-26 on page 35.

Figure 4-26 shows the storage resource added to storage space `SC01` with an auto generated name `cp_SC01_Service_xxxx`. On the right, you can see the same storage resource shown as the child pool created on the storage system.

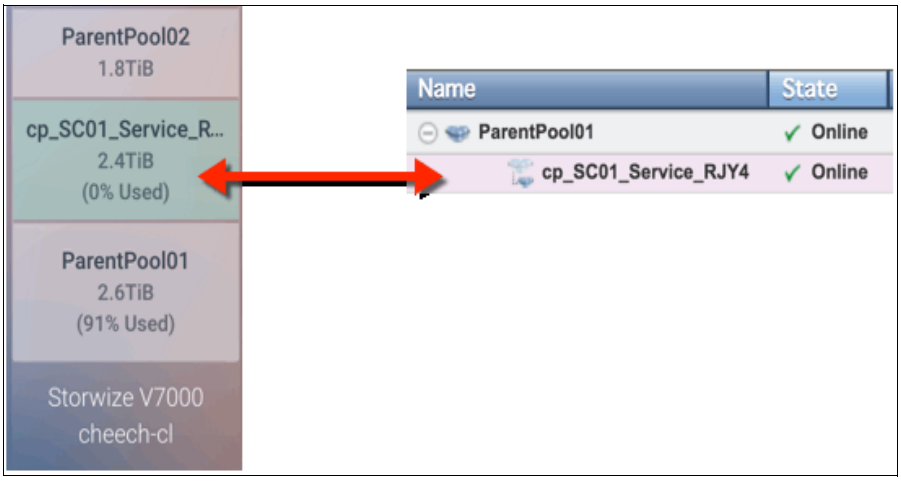

Figure 4-26 SCB storage resource and child pool on the storage system

4.3 Configuring VMware vSphere Web Client

The configuration of Virtual Volumes is now complete for our storage system and IBM Spectrum Control Base Edition. As such, the tasks required of the storage administrator have been completed. Finally, we move our attention to configuring the vSphere environment. Before we can use VVols, the following tasks should be performed by the vSphere administrator.

4.3.1 Adding a storage provider

Adding a storage provider to the vSphere Web Client provides the last, but pivotal step in configuring the VVols environment. The storage provider is the IBM Spectrum Control Base Edition instance that we configured earlier in the chapter.

To add the SCB instance as a storage provider in the vSphere Web Client, complete the following steps:

1. In your vSphere Web Client Home window, click **vCenter Inventory Lists**.
2. From the Inventory list on the left, click **vCenter Servers** under the Resources twistie.
3. Select your vCenter Server from the pane on the left.
4. Click the **Manage** tab in the right pane.
5. Click **Storage Providers** in the right pane.
6. Click the green Plus sign (+) button.

You should now have arrived at the same dialog as that shown in Figure 4-27.

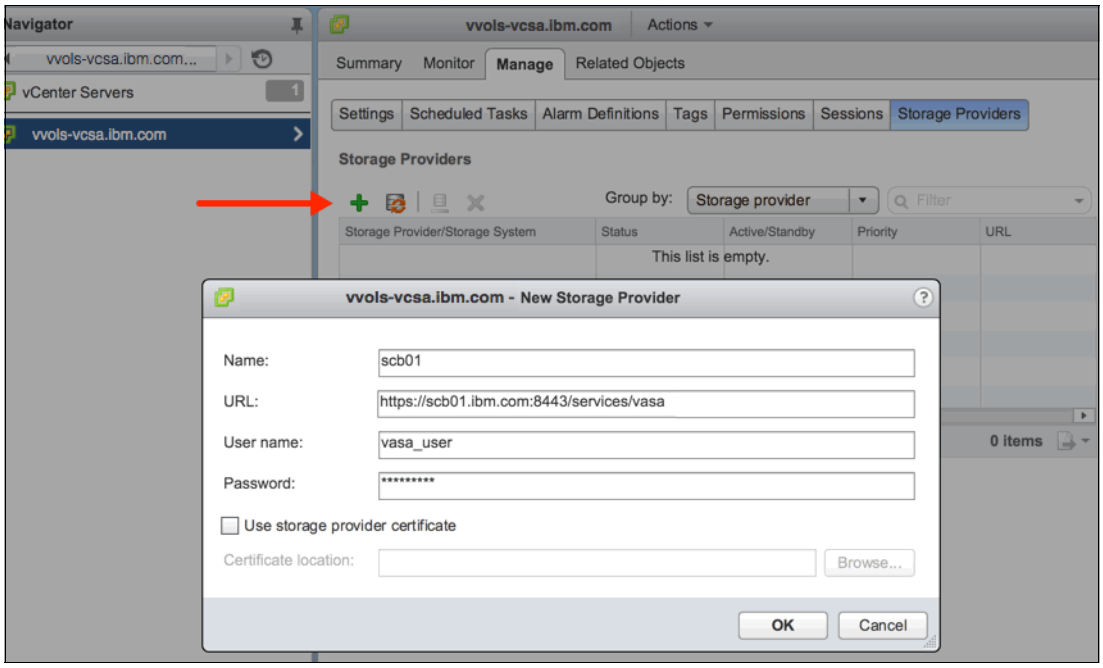

Figure 4-27 The vSphere Web Client adding a New Storage Provider

7. In the New Storage Provider dialog, complete the following fields:
   a. Enter a friendly **Name**.
   b. Enter the **URL** of your SCB instance with the required port and path details, for example, `https://<host.domain>.com:8443/services/vasa`.
   c. Enter the **User name** and **Password** we created earlier in 4.2.8, “VASA credentials” on page 27.
8. Click **OK** and, when prompted, select **Yes** to accept the security alert. This process can take a few minutes to complete.

   The SCB instance should now have established a secure SSL connection to the vCenter Web Client as an *Active* Storage Provider. See Figure 4-28 for reference.

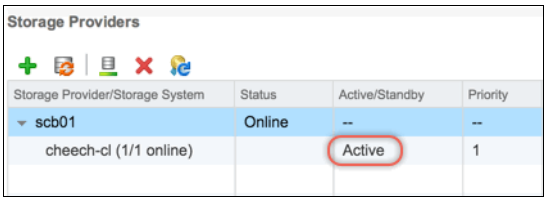

Figure 4-28 The vSphere Web Client Storage Providers pane

4.3.2 Adding a VVol datastore in VMware vSphere Web Client

To use the *storage space* that you defined earlier in your IBM Spectrum Control Base Edition instance, you need to add a VVol datastore to the vSphere Web Client.

To add a datastore, complete the following steps:

1. In your vSphere Web Client Home window, click **vCenter Inventory Lists** and select **Datastores**.
2. Click the **Create a New Datastore** icon in the pane to the right, as shown in Figure 4-29. This starts the New Datastore wizard.

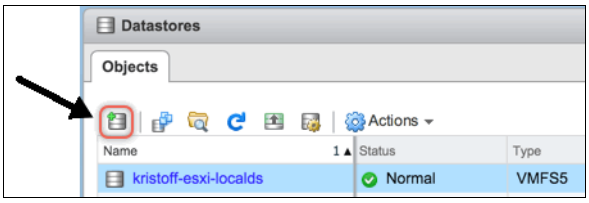

Figure 4-29 In vSphere Web Client, creating a new datastore

3. In the New Datastore wizard, select the placement location for the datastore (the cluster or hosts that you want to share the VVols datastore with).
4. Select **VVol** as the datastore type.
5. Select a backing storage container. Any storage spaces defined in IBM Spectrum Control Base Edition should be available to select as a storage container. The storage space names, as configured in SCB, are displayed in the **Backing Storage Container** field, as shown in Figure 4-30.

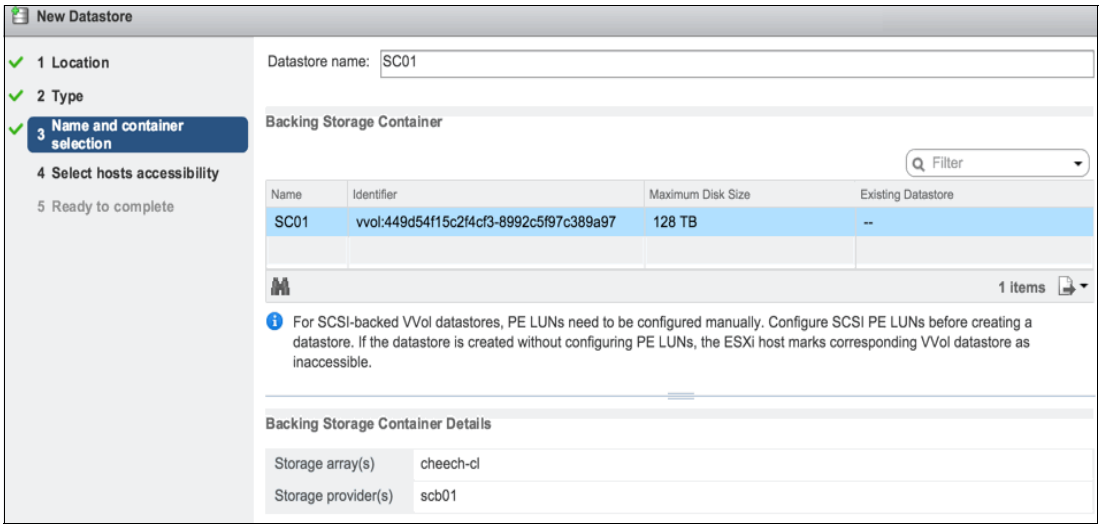

Figure 4-30 The vSphere Web Client New Datastore dialog

6. Select the hosts that require access to the VVol datastore.

7. Review the configuration options and click Finish. Figure 4-31 shows the settings used in our example configuration.

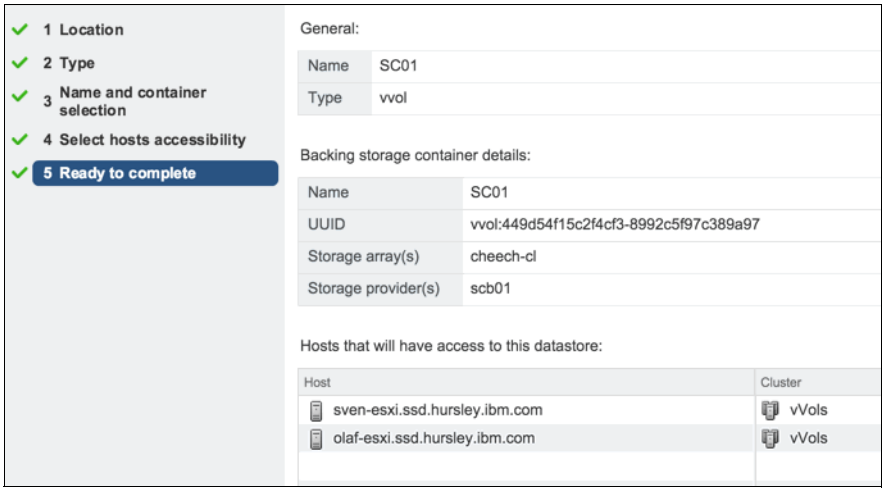

Figure 4-31 The vSphere Web Client new VVol datastore settings

Figure 4-32 shows our newly created VVol datastore.

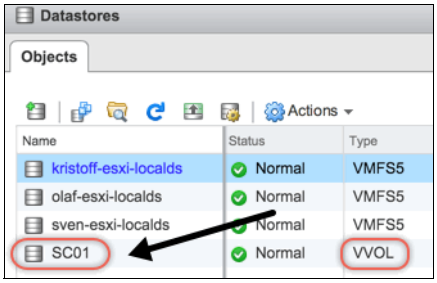

Figure 4-32 The vSphere Web Client VVols datastore

4.3.3 Setup complete

The vSphere administrator now has the essentials that are required to start using Virtual Volumes. At this point, we suggest deploying a standby instance of IBM Spectrum Control Base Edition, for redundancy of virtual machine management. If you want to follow this suggestion, you can find instructions in 4.4, “Configuring additional SCB servers for high availability” on page 39.

For a more detailed explanation of the key tasks involved in managing and maintaining your VVols environment now that it is successfully configured, see Chapter 5, “Managing and maintaining a Virtual Volumes environment” on page 49.

To maximize the potential for Virtual Volumes in driving automation with storage-based policies, see Chapter 7, “Storage Policy-Based Management” on page 83.

If you want a deeper understanding of the storage system concepts and operations involved in an active Virtual Volumes environment, more details are provided in Chapter 8, “Advanced concepts, operations, and functions in the Virtual Volumes environment” on page 93.

4.4 Configuring additional SCB servers for high availability

Multiple IBM Spectrum Control Base Edition instances can be configured to manage the same storage system. By doing so, we ensure the high availability (HA) of the *management* of VMs specific to our VVols environment. The steps involved in setting up an additional standby instance are essentially the same as configuring your initial *active* instance.

The installation and configuration notes can be followed from 4.2, “Configuring IBM Spectrum Control Base Edition” on page 22 through 4.2.12, “Adding a storage space” on page 31, inclusive. For completeness, we cover the required steps again in this section.

Tip: If you created a template of your SCB VM, now would be a great time to ask your vSphere administrator to deploy a second instance for use as your standby SCB server.

4.4.1 Downloading the IBM Spectrum Control Base Edition installation files

Obtain the current SCB installation files:

https://ibm.biz/ibm-scb

Extract them to a temporary folder (for example, `/scb_install`) on your SCB VM, hereafter referred to as *SCB* and the *SCB server*. Example 4-3 shows the typical installation files for SCB after they have been extracted.

Example 4-3 IBM Spectrum Control Base Edition - installation files

```
ibm_spectrum_control-2.2.1-xxxx-x86_64.bin
nginx-1.6.2-1.el6.ngx.x86_64.rpm
postgresql92-9.2.10-1PGDG.rhel6.x86_64.rpm
postgresql92-contrib-9.2.10-1PGDG.rhel6.x86_64.rpm
postgresql92-libs-9.2.10-1PGDG.rhel6.x86_64.rpm
postgresql92-server-9.2.10-1PGDG.rhel6.x86_64.rpm
```

4.4.2 Installing IBM Spectrum Control Base Edition

To install SCB, complete the following steps:

1. Change to the folder on your SCB server where you previously extracted the installation files and start the installation of SCB.
2. On the command line, type `rpm -iv *.rpm`. Sample output is shown in Example 4-4.

*Example 4-4 Example output of command rpm -iv *.rpm*

```
warning: nginx-1.6.2-1.el6.ngx.x86_64.rpm: Header V4 RSA/SHA1 Signature, key ID
7bd9bf62: NOKEY
warning: postgresql92-9.2.1O-1PGDG.rhel6.x86_64.rpm: Header V4 DSA/SHA1
Signature, key ID 442df0f8: NOKEY
warning: uuid-1.6.1-1O.el6.x86_64.rpm: Header V3 RSA/SHA256 Signature, key ID
fd431d51: NOKEY
Preparing packages for installation...
postgresql92-libs-9.2.1O-1PGDG.rhel6
postgresql92-9.2.1O-1PGDG.rhel6
uuid-1.6.1-1O.el6
postgresql92-contrib-9.2.1O-1PGDG.rhel6
postgresql92-server-9.2.1O-1PGDG.rhel6
nginx-1.6.2-1.el6.ngx
```

3. Next, we turn our attention to installing the `.bin` file. At the command line, type `./ibm_spectrum_control-2.2.1-xxxx-x86_64.bin` (note that you may have a later file version).
4. When prompted, enter `1` to accept the agreement. After a minute or two, you will receive an `Installation completed successfully` message, and you can go to the next step.

4.4.3 Changing the IBMSC administration account

To modify the IBMSC user account, on the command line, type `passwd ibmsc` and enter a new password.

4.4.4 Accessing the IBM Spectrum Control Base Edition web interface

To access the SCB web interface, complete the following steps:

1. Open a web browser and connect to your standby SCB server's web interface with the following URL `https://<SCB Server Name>:8443`. You might need to *trust* the certificate before continuing to the logon page, as shown in Figure 4-33.

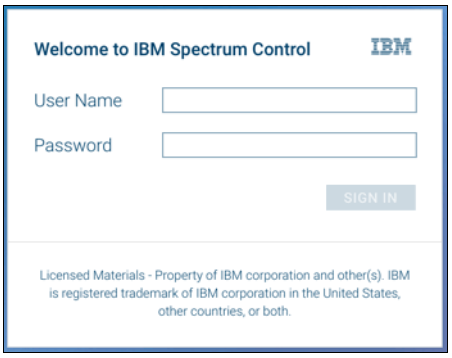

Figure 4-33 SCB web interface logon

2. Log in to the standby SCB web interface with the following default credentials:
   - **Username**: `admin`
   - **Password**: `admin1!`

4.4.5 Changing the default administrator password

To change the password, complete the following steps:

1. Click the **Settings** button and select **Users** in the drop-down menu. You are presented with the Users pane. Highlight the **admin** user account by clicking it. Two small buttons, **Edit User** and **Delete User**, appear to the right, as shown in Figure 4-34.

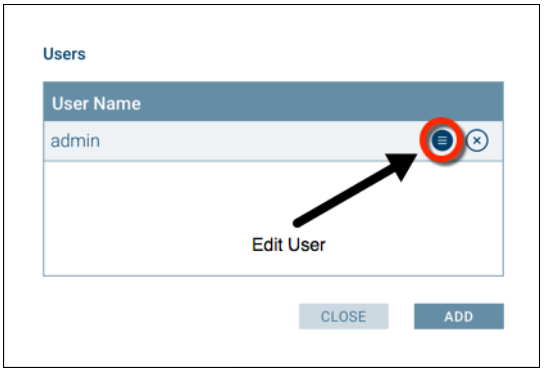

Figure 4-34 SCB Users selection box

2. Click the **Edit User** button to modify the **admin** user properties, as shown in Figure 4-35.

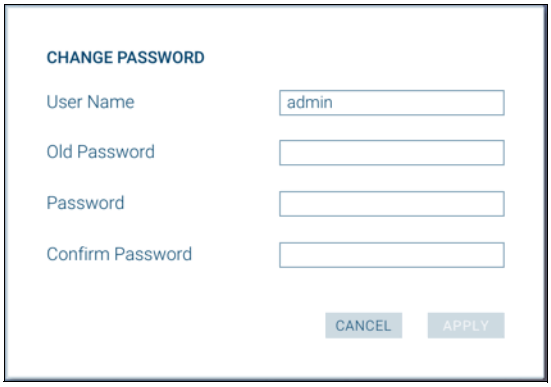

Figure 4-35 SCB CHANGE PASSWORD box

3. Enter the existing password of `admin1!` and then enter a new password for your admin user. Click **APPLY** → **Close**.

4.4.6 Storage Credentials

Earlier in this chapter (Figure 4-4 on page 20), you created a user on the storage system for use by SCB. Recall those details for use in completing the Storage Credentials:

1. Click the **Settings** button and select **Storage Credentials**, as shown in Figure 4-36.

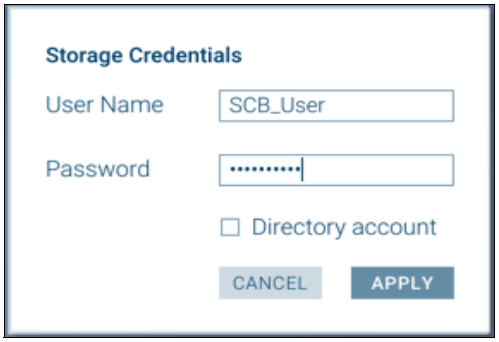

Figure 4-36 SCB Storage Credentials dialog

2. Enter the User Name and Password of the SCB user account created on the storage system.
3. Click **Apply**.

4.4.7 VASA credentials

The VASA credentials are defined in SCB and used by the VMware vSphere Web Client to connect to the SCB server. To configure VASA credentials, complete the following steps:

1. Click the **Settings** button and select **VASA credentials** on the Settings menu. Figure 4-37 shows the VASA Credentials dialog.

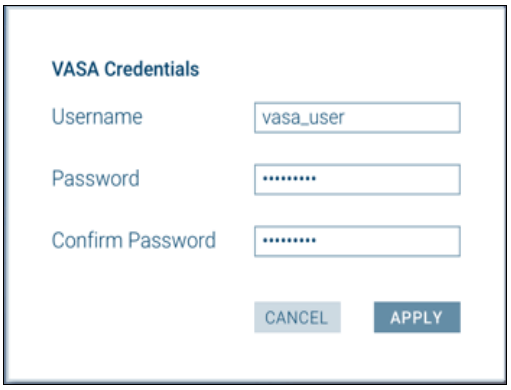

Figure 4-37 SCB VASA Credentials dialog

2. Enter a Username and Password, then click **Apply**.

> **Tip:** The VASA user credentials will be required later when adding the standby SCB storage provider in the vSphere Web Client.

4.4.8 General settings

This step is of particular importance to the configuration of high availability for your VVols environment.

> **Important:** The HA Group name provided to the *standby* SCB instance must be identical to the HA Group name that was set when registering the first, *active* SCB instance.

To configure the general settings, complete the following steps:

1. Click the **Settings** button and select **General Settings** on the Settings menu, as shown in Figure 4-38.

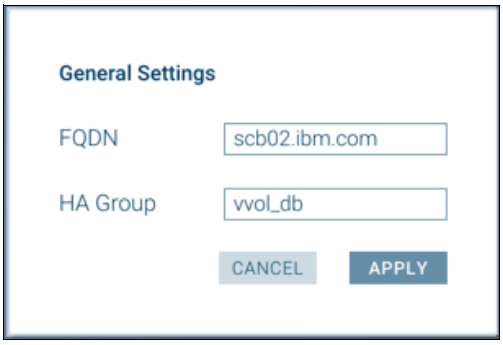

Figure 4-38 SCB General Settings

2. Enter the FQDN of your standby SCB server.
3. Enter the name for the HA Group. This *must* match the HA Group set on the first, *active* SCB instance.
4. Click **APPLY**.

4.4.9 Server certificate

Figure 4-39 shows the Server Certificate dialog for your *standby* SCB server.

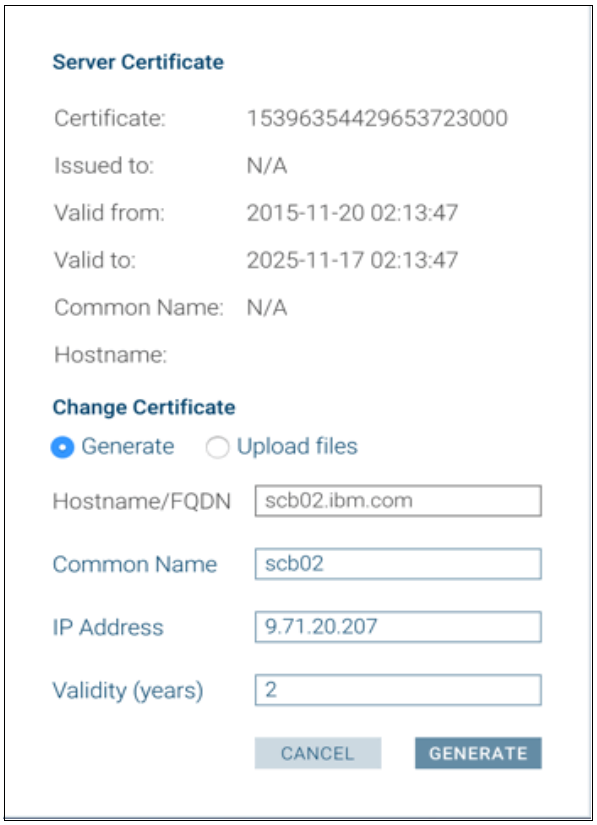

Figure 4-39 SCB Server Certificate dialog

To configure the server certificate, complete the following steps:

1. Click the **Settings** button and select **Server certificate** on the Settings menu.
2. Enter the FQDN of your standby SCB server.
3. Enter a Common Name for your standby SCB server.
4. Enter the IP Address of your standby SCB server.
5. Choose a certificate Validity period.
6. Click **GENERATE**.

4.4.10 Adding a storage system

To add a storage system, complete the following steps:

1. Click the **Add** button on the Storage Systems pane.

2. Enter the Management IP address or Hostname of the storage system, as shown in Figure 4-40. This is the same storage system added to the active SCB instance.

Figure 4-40 SCB Add New Array dialog

3. Click **ADD**. The storage system is added to the Storage Systems pane (Figure 4-41).

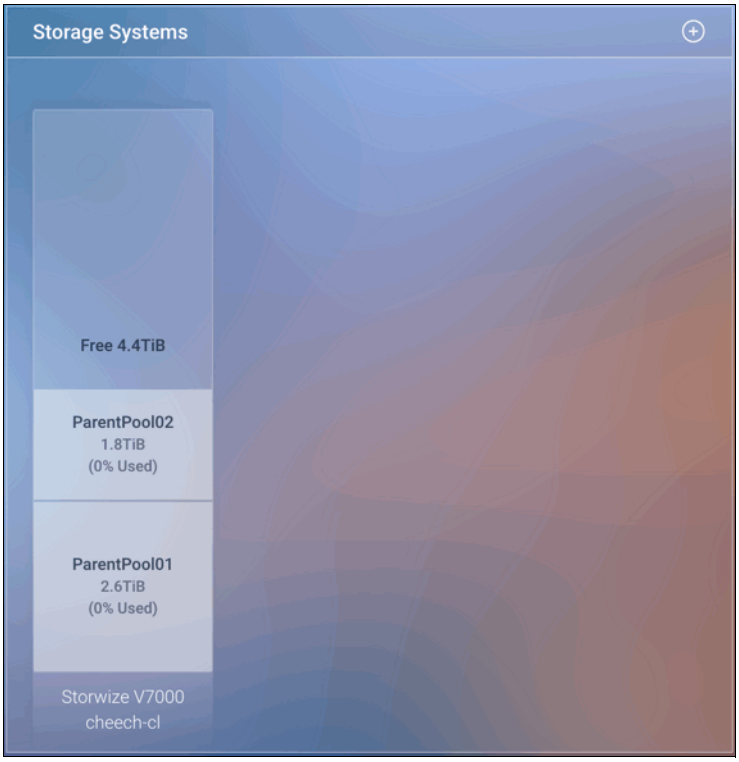

Figure 4-41 SCB Storage Systems

Note: Do not configure the storage spaces or storage services on your *standby* SCB instance. This information is populated from the database stored on the storage system if there is a failover.

4.4.11 Configuring the VMware vSphere Web Client

Adding a standby storage provider to the vSphere Web Client provides the last, but pivotal step in configuring your storage provider HA Group.

To add the standby SCB instance as a storage provider in the vSphere Web Client, complete the following steps:

1. In your vSphere Web Client home window, click **vCenter Inventory Lists**.
2. From the Inventory list on the left, click **vCenter Servers** under the Resources twistie.
3. Select your `vCenter Server` from the pane on the left.
4. Click the **Manage** tab in the right pane.
5. Click **Storage Providers** in the right pane.
6. Click the green Plus sign (+) button.
7. You should now have arrived at the window depicted in Figure 4-42.

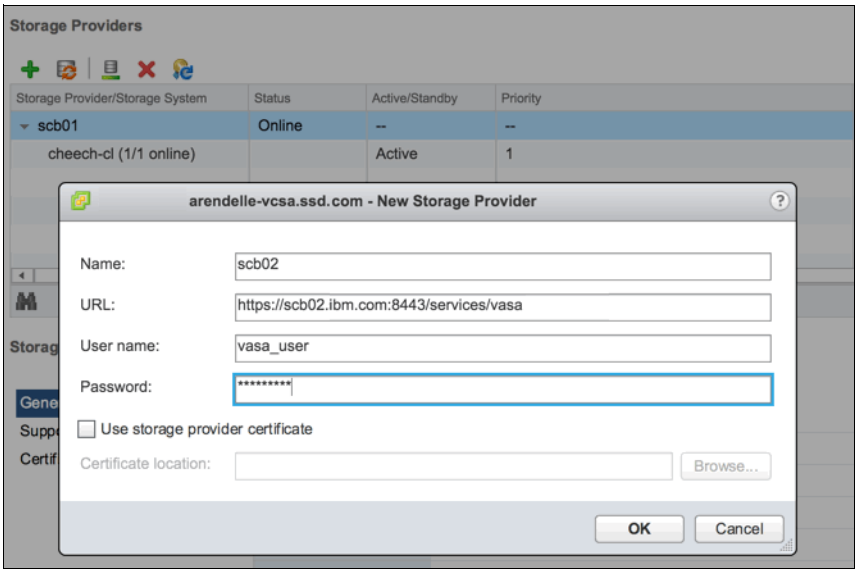

Figure 4-42 vSphere Web Client adding a standby Storage Provider

8. In the New Storage Provider dialog, complete the following steps:
   a. Enter a friendly **Name**.
   b. Enter the **URL** of the standby IBM Spectrum Control Base Edition instance.
   c. Enter the **`User name`** and **`Password`** that we created earlier in “VASA credentials” on page 43.
9. Click **OK** and, when prompted, select **Yes** to accept the security alert.

The SCB instance should now have established a secure SSL connection to the vCenter Web Client as a *standby* storage provider. See Figure 4-43 for reference.

vCenter continuously monitors the availability of the active storage provider. Should an issue arise with the active storage provider where the VASA services have stopped responding, a failover to the standby storage provider is initiated.

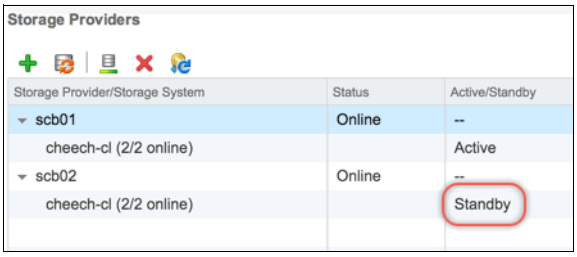

Figure 4-43 vSphere Web Client Storage Providers list

5

Managing and maintaining a Virtual Volumes environment

This chapter contains the information required to manage and maintain your Virtual Volumes (VVols) environment. The content is designed for storage administrators who have already performed the initial configuration, and who want to make changes to their existing environment. If you have not yet configured Virtual Volumes, see Chapter 4, “Installation and configuration” on page 17 for comprehensive instructions.

We begin the chapter by highlighting some high-level workflows, which are fundamental to managing the storage allocated to the Virtual Volumes environment. We then provide detailed instructions for completing the most common VVols-related tasks using the management graphical user interfaces (GUIs) for the storage system and IBM Spectrum Control Base Edition (SCB). For further tasks, or command-line interface (CLI) syntax, see the IBM Knowledge Center for the relevant product.

Although this book is focused on management of the VVols environment at the storage level, we have also included several vSphere Web Client tasks that are pertinent to the storage administrator. For further information about managing virtual machines using the vSphere Web Client, see VMware’s extensive Knowledge Base:

http://kb.vmware.com

This chapter provides information about the following topics:

- 5.1, “Example workflows” on page 50
- 5.2, “Storage system tasks” on page 53
- 5.3, “IBM Spectrum Control Base Edition tasks” on page 55
- 5.4, “VMware vSphere Web Client tasks” on page 66

5.1 Example workflows

Although the Virtual Volumes environment mitigates the need for frequent involvement from the storage administrator, he still maintains overall control of the storage system. The storage administrator is able to allocate or reclaim storage from the VVols environment on demand. We briefly describe some example workflows for these solutions. For comprehensive details about performing each subtask involved in these workflows, follow the references provided.

5.1.1 Allocating more storage to the vSphere administrator

When vSphere administrators request additional storage capacity for virtual machines, they should provide the following information to the storage administrator:

- The capacity requirements of the new storage.
- The storage capabilities required.

To fulfill such requests, the storage administrator can take one of several approaches:

- Expand an existing storage resource (child pool).
- Allocate an additional storage resource (child pool) to an existing storage space.
- Create a new storage space, which includes a new storage service and resource (child pool).

The first two options expand the size of an existing VVol datastore. These options are simpler if the capability options that the vSphere administrator requests can be met by existing datastores. Note that for a VVol datastore, the vSphere administrator will not need to format the new storage after the capacity has been increased.

The third option presents a new storage container to the vSphere infrastructure. This option provides the vSphere administrator with the ability to create a new VVol datastore.

The required actions for each option are listed in the following sections.

Option to expand an existing storage resource (child pool)

This is the simplest approach for allocating additional storage to the VVols environment.

To expand an existing storage resource, the storage administrator should complete the following actions:

1. Resize the storage resource. This allocates additional capacity to a child pool on the storage system. For full steps, see “Resizing a storage resource” on page 60.
2. Inform the vSphere administrator that additional storage has been added to the VVol datastore.

The vSphere administrator should rescan storage providers to update datastore capacity information. For more information, see “Rescanning storage providers” on page 67.

Option to allocate an additional storage resource (child pool) to an existing storage space

In some cases, you might want to create separate storage resources, rather than expanding an existing resource. For example, you might want to offer storage capacity from a separate parent pool.

To allocate an additional storage resource, the storage administrator should complete the following actions:

1. Select the storage space and VVols-enabled service that will be associated with the new storage resource.
2. Create a new storage resource. For full details, see “Creating a storage resource” on page 59.

The vSphere administrator should rescan storage providers to update datastore capacity information. For more information, see “Rescanning storage providers” on page 67.

Option to create a new storage space and add a new storage service and resource (child pool)

This process requires greater involvement from the storage administrator, but might be required if the vSphere administrator specifies that a separate datastore should be provided. This option might also be necessary if new capability options are required by the vSphere administrator, because providing additional capability options could require the creation of a new storage service.

To create a new storage space and add a new storage service and resource, the storage administrator should complete the following actions:

1. Create a new storage space. Ultimately, this will be used as a unique datastore in the vSphere Web Client. For details about this process, see “Creating a storage space” on page 55.
2. Create a new storage service for this space, a service which offers the capabilities that the vSphere administrator has requested. For guidance, see “Creating a storage service” on page 56.
3. Allocate a new storage resource (child pool) from the wanted parent pool. For more information, see “Creating a storage resource” on page 59.
4. At this point, the storage administrator should provide the vSphere administrator with the name of the new storage space.

To use the new storage effectively, the vSphere administrator should complete the following actions:

1. Rescan storage providers to populate information about the new storage space in the vSphere Web Client. For more information, see “Rescanning storage providers” on page 67.
2. Create a new VVol datastore, selecting the new storage space as the backing storage container. For more information, see “Creating a VVol Datastore” on page 68.
3. (Optional) The vSphere administrator might also want to define policies to use additional capabilities that have been offered as part of the new storage space.

5.1.2 Reclaiming storage from the vSphere administrator

If the vSphere administrator ceases to require a particular VVol datastore, the storage administrator can reclaim the capacity that was being used by it. Similarly, if the storage administrator wants to reduce the capacity that is used by the Virtual Volumes environment, they can shrink the amount of capacity that has been allocated.

The vSphere administrator should provide the following information to the storage administrator:

- The name of the backing storage container, which can be removed or reduced in capacity
- The amount of capacity that can be safely reclaimed by the storage administrator (where relevant)

The actions required in each use case are described in the following sections.

Removing a VVol datastore and the associated storage objects

To remove a VVol datastore, the vSphere administrator should complete the following actions:

1. Ensure that any virtual machines (VMs) have been migrated onto another datastore.
2. Unmount the VVol datastore from all hosts. For instructions, see “Unmounting a VVol datastore” on page 70.
3. Inform the storage administrator that the objects associated with this datastore can be safely removed, providing the relevant information about the backing storage container.

The storage administrator should complete the following actions:

1. Remove the storage resource that was previously associated with the VVol datastore. This task deletes a child pool from the storage system, and enables the parent pool to reclaim this quota of storage. For guidance, see “Deleting a storage resource” on page 61.
2. Unless it is required elsewhere, remove the storage service that was associated with the storage resource that was just deleted. See “Deleting a storage service” on page 57.
3. Remove the storage space that corresponded to the VVol datastore. For instructions, see: “Deleting a storage space” on page 56.

Reducing the capacity of a VVol datastore

To reduce the capacity of a VVol datastore, the vSphere administrator should first inform the storage administrator about the amount of storage that can be reclaimed and the name of the backing storage container.

Next, the storage administrator should complete the following actions:

1. Shrink the storage resource that relates to the backing storage container. For more information, see “Resizing a storage resource” on page 60.
2. Inform the vSphere administrator that the capacity reduction is complete.

Finally, the vSphere administrator should rescan storage providers to update datastore capacity information. For more information, see “Rescanning storage providers” on page 67.

5.2 Storage system tasks

The following tasks should be performed using the management GUI for your storage system.

5.2.1 Managing VMware ESXi hosts

Hosts must be running ESXi version 6.0 or higher to use VVols functionality. In addition, the host must be added to the storage system, with the host-type field set to `VVol`. You can also enable VVols on existing hosts by changing the host type to `VVol`. For more information about VVols-enabled hosts, see 3.6, “VVols-enabled host” on page 11.

Adding VVols-enabled hosts

To add a new host to the system, complete the following steps:

1. Select **Hosts** → **Hosts** → **Add Host**.
2. Select the connectivity type for the host.
3. Enter a name for the host.
4. Input the port details.
5. Click **Advanced** to show additional options.
6. Open the **Host Type** drop-down menu. Select **VVol**.
7. Select **Add** to complete the task.

Figure 5-1 shows the Add Host dialog, which is used to complete the steps listed previously.

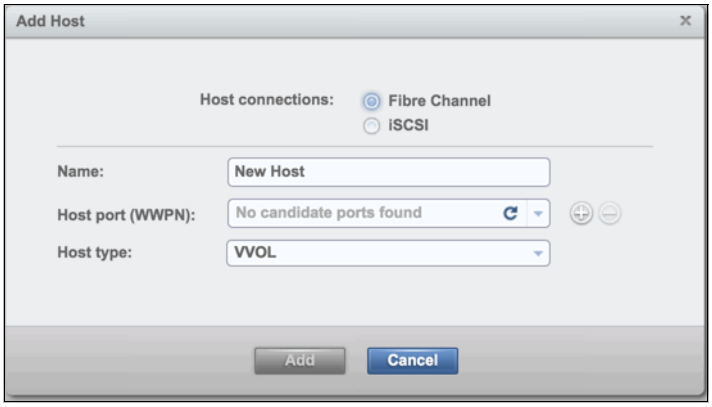

Figure 5-1 Setting the Host Type correctly when adding a New Host to the VVols environment

Modifying existing hosts to enable VVols

You might want to enable VVols for a host that is already registered on the system. Existing hostmap configurations are *unaffected* by enabling your ESXi hosts for VVols. It is possible for the same host to have access to VVols and VMware file system (VMFS) datastores simultaneously.

Important: Hosts of type `VVol` must be configured to access all input/output (I/O) groups, as shown in Figure 5-2.

To enable VVols functionality on an existing host, complete the following steps:

1. Select **Hosts** → **Hosts** and select a host.
2. Select **Actions** → **Properties**. Select the **Overview** tab and click **Edit**.
3. Open the **Host Type** drop-down menu. Select **VVol**.
4. Select **Save** to complete the task.

Figure 5-2 highlights the changes that must be made to enable VVols for an existing host.

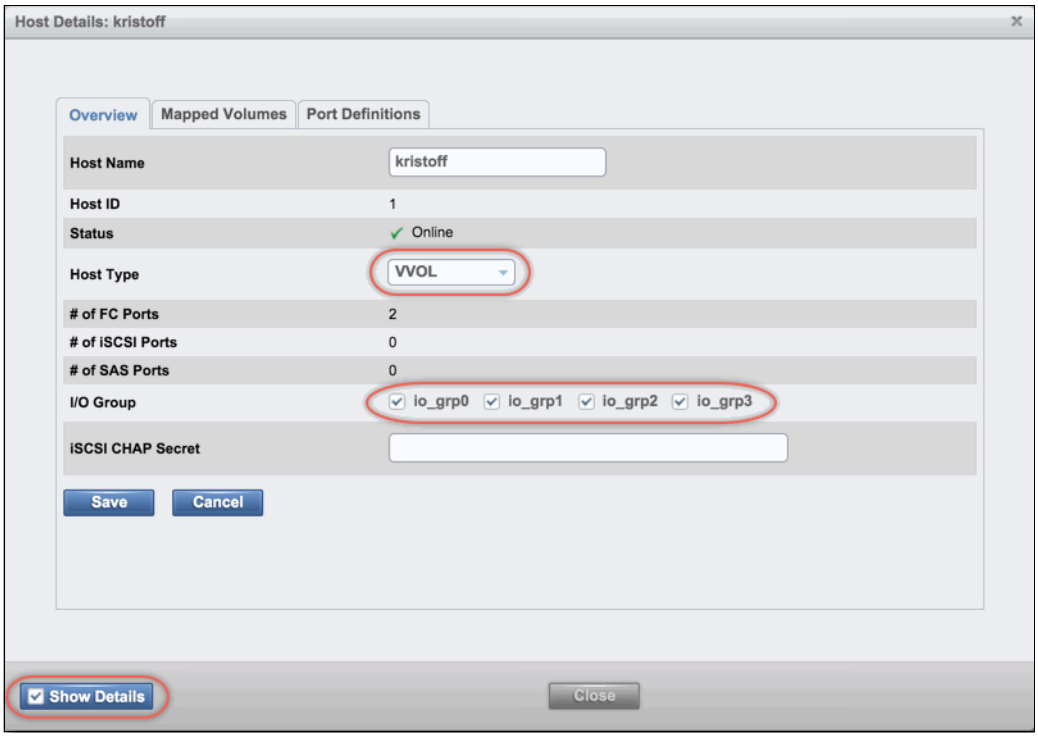

Figure 5-2 Illustrates the Host Type and I/O Group settings required when enabling VVols for a host

5.2.2 Managing the utility volume

Although the database stored on the utility volume is solely managed by IBM Spectrum Control Base Edition, there are two tasks that the storage administrator might want to perform on this volume. For an explanation of the utility volume, see 3.9, “Utility volume” on page 14.

Mirroring the utility volume

We suggest storing a mirrored copy of the utility volume, for redundancy. For more information about this, see 6.1.2, “Mirroring the utility volume” on page 76.

You might have already created a mirrored copy of the utility volume, using the Enable VVols wizard. These steps are shown in 4.1.1, “Enabling VVols” on page 18.

To create a mirrored copy of the utility volume after VVols is already enabled on the system, you must use the CLI. Issue the following command:

```
svctask addvdiskcopy -mdiskgrp <mdiskgrp_id> <volume_id>
```

Important: Additional parameters should not be added to the previous command, because the properties of the utility volume are predefined.

Renaming the utility volume

You might find it easier to track the utility volume by assigning it a user-defined name.

To rename the utility volume, issue the following command:

```
svctask chvdisk -name utility_volume <volume_id>
```

5.3 IBM Spectrum Control Base Edition tasks

The following tasks should be performed using the IBM Spectrum Control Base Edition management GUI. For a brief introduction to this management GUI, see 4.2.5, “An overview of the SCB web interface” on page 24.

Important: IBM Spectrum Control Base Edition tasks should always be run from the active instance. If you are unsure which SCB instance is the active provider, you can check using the vSphere Web Client by completing the following steps:

1. Start at the Home page of the vSphere Web Client.
2. Select **vCenter Inventory Lists** → **vCenter Server** → ***<your vCenter Server instance>***.
3. Select **Manage** → **Storage Providers**.

5.3.1 Managing storage spaces

Storage spaces represent a combination of storage services and storage resources (child pools).

Creating a storage space

To create a new storage space, complete the following steps:

1. Select the **Settings** button and then select **Manage spaces**. The list of storage spaces displays.
2. Click **Add**.
3. Enter the name and description of the storage space that you want to create, and then click **Apply**, as shown in Figure 5-3.

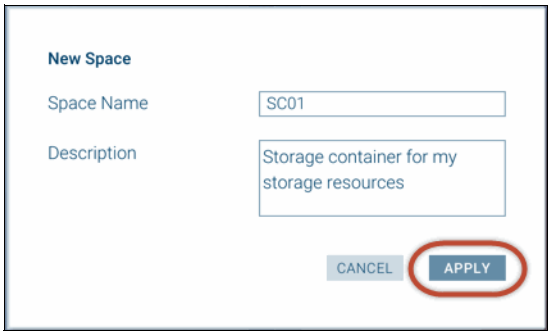

Figure 5-3 Creating a new storage space

4. After a storage space is configured, storage services and storage resources must be added before the object can be used as the backing container for a VVol datastore. For information about creating storage services and storage resources, see “Creating a storage service” on page 56 and “Creating a storage resource” on page 59.
5. If you are creating a new storage space for an instance of IBM Spectrum Control Base Edition that is already registered in the vSphere Web Client, you might need to rescan storage providers to populate information about the new container. For the steps required to rescan a storage provider, see “Rescanning storage providers” on page 67.

Deleting a storage space

Important: Before removing a storage space, ensure that it is not mounted as a datastore (in the vSphere Web Client). On IBM Spectrum Control Base Edition, you must also delete all storage services that are defined on this space. For more information about this, see “Deleting a storage service” on page 57.

To delete a storage space, complete the following steps:

1. Select **Settings** → **Manage spaces**. The list of storage spaces displays.
2. Select the space that you want to remove. Click the **Remove** button, as shown in Figure 5-4.

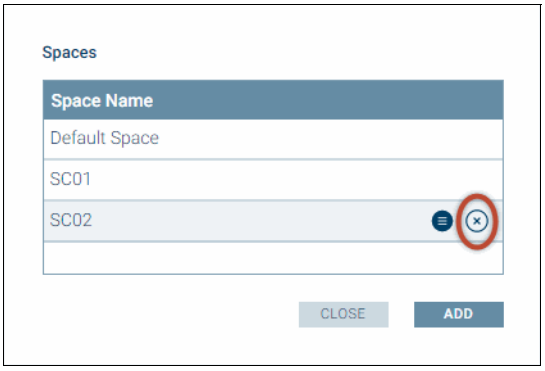

Figure 5-4 Deleting a storage space

3. Click **OK** to delete the space, or **Cancel** to cancel the operation.

5.3.2 Managing storage services

Storage services are always associated with a particular storage space. In the Virtual Volumes environment, they are used to define the set of capabilities that will be associated with a storage resource (child pool or set of child pools) that belongs to the space.

Creating a storage service

To create a storage service:

1. Select a storage space from the Spaces and Services pane.
2. Use the **Add** button to add a storage service to the storage space.
3. Enter a name and description for the storage service.

Remember: The storage service name forms part of the name for the child pool that is created on your storage system, for example `cp_SC01_Service_xxxx`.

4. Define the capabilities for the storage service. For some examples of storage services, and factors to consider when defining storage service capabilities, see 6.2.2, “Defining storage spaces” on page 78.
5. Ensure that you select the **VVol Service** check box.
6. Select **Apply** to create the storage service.

Figure 5-5 shows a VVols-enabled storage service, without storage capabilities explicitly configured.

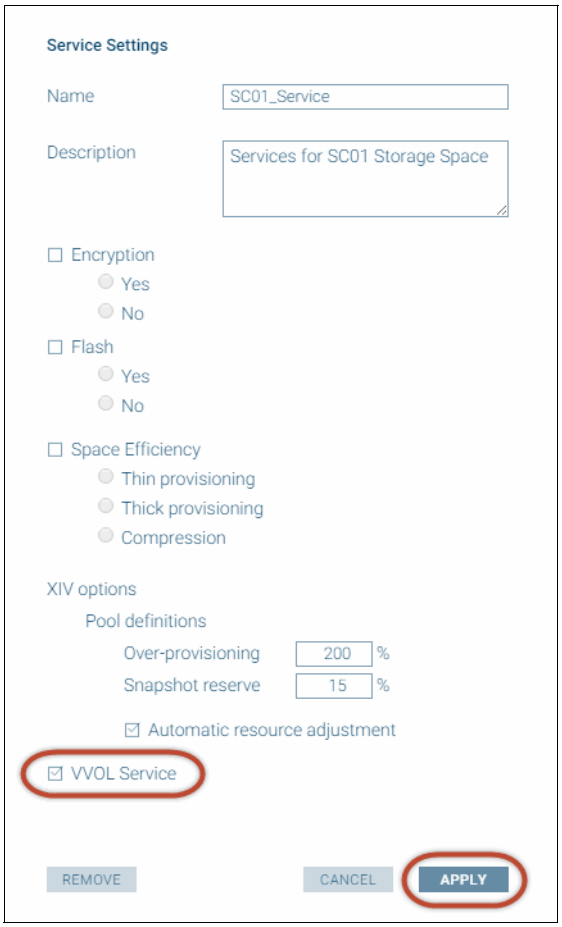

Figure 5-5 Creating an unrestricted storage service for VVols

Deleting a storage service

Important: Before deleting a storage service, ensure that it is not attached to any storage resources. To associate a storage resource with a different storage service, see “Moving storage resources between storage spaces” on page 62. To delete a storage resource, see “Deleting a storage resource” on page 61.

To delete a storage service, complete the following steps:

1. Select the relevant space from the Spaces and Services pane.
2. Hover your cursor over the storage service that you want to delete.

3. Select the **Edit** button and then select **Settings**, as shown in Figure 5-6.

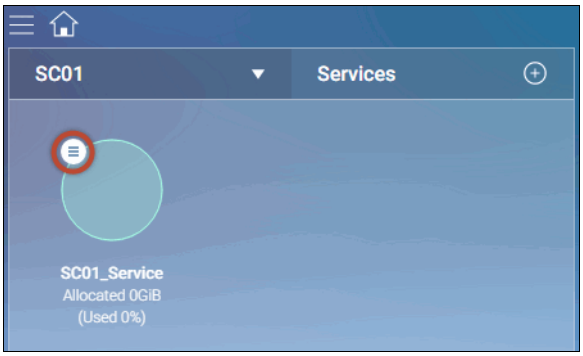

Figure 5-6 Editing a storage service

4. Click **Remove**, as shown in Figure 5-7.

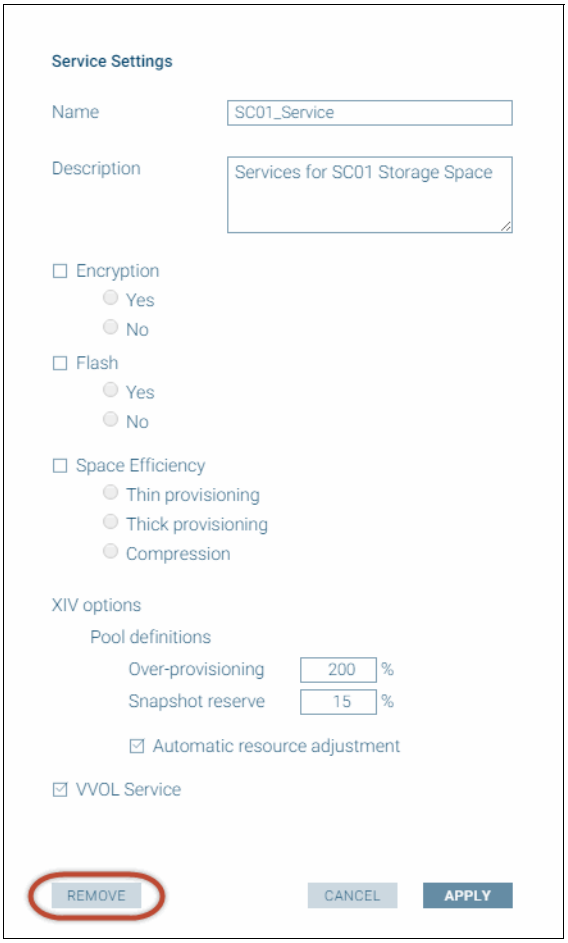

Figure 5-7 Removing a storage service

5. Click **OK** to delete the storage service, or **Cancel** to cancel the operation.

5.3.3 Managing storage resources (child pools)

In the Virtual Volumes environment, storage resources are allocated from the storage system in the form of a *child pool*. A child pool is a logical quota of storage, which the storage administrator allocates from an existing parent pool. Child pools can be resized (either grown or shrunk) at any time.

Creating a storage resource

To create a storage resource, complete the following steps:

1. On the Spaces and Services pane, select the storage space under which you want to define your resource.
2. Select your VVols-enabled storage service. The storage service is highlighted.
3. Hover your cursor over the top of the storage system and select the **Edit** button, and then select **Modify**.
4. From the drop-down menu, select the parent pool that will be used for storage allocation.
5. Enter a size for the new resource.
6. Select **Add** to complete the operation.

   Figure 5-8 illustrates the full process of defining a new resource.

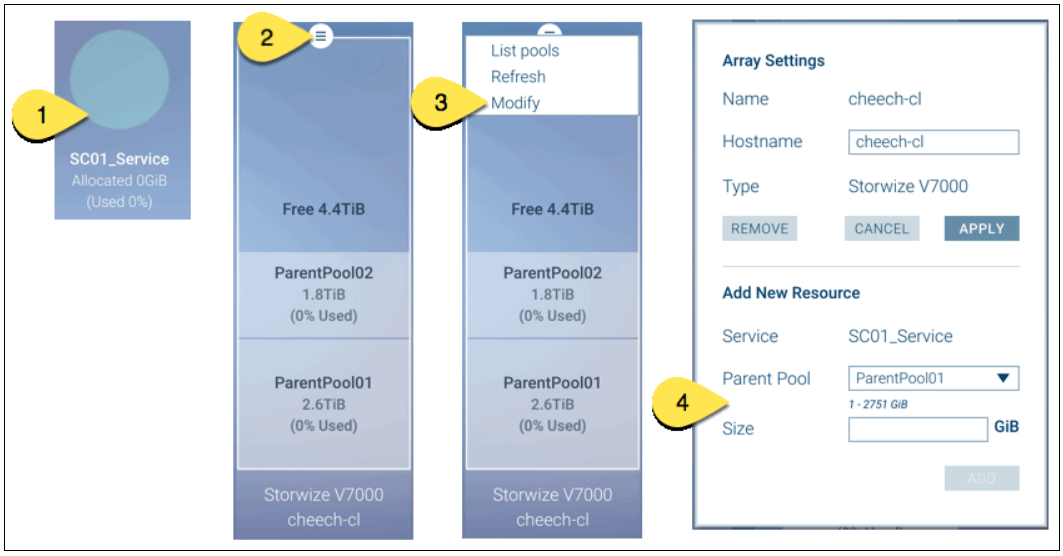

Figure 5-8 Adding a new storage resource (child pool)

7. A child pool is created on the storage system. Figure 5-9 shows the correlation between a storage resource (on SCB) and a child pool (on the storage system).

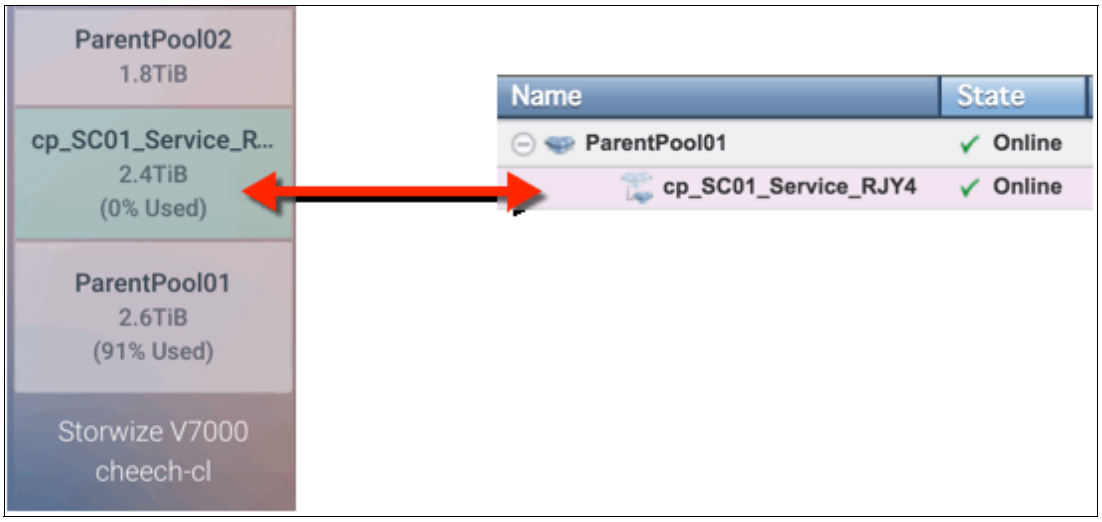

Figure 5-9 Creating a storage resource (child pool) using IBM Spectrum Control Base Edition

Resizing a storage resource

To resize a storage resource, complete the following steps:

1. Hover your cursor over the top of the storage system.
2. Select the **Settings** button and then select **List pools**, as shown in Figure 5-10.

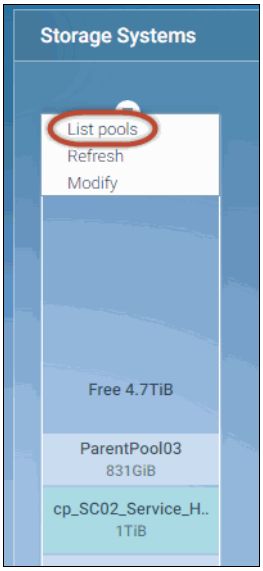

Figure 5-10 List all pools associated with the storage system

3. Select the wanted storage resource and then the **Settings** button, as shown in Figure 5-11.

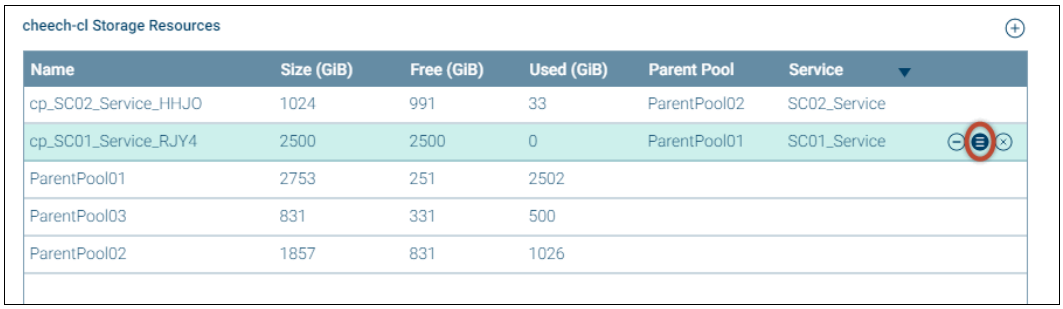

Figure 5-11 Location of the Settings button for a storage resource

4. Edit the **Size** field as required (shown in Figure 5-12).

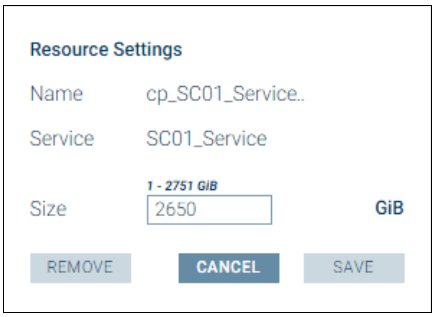

Figure 5-12 Resizing a storage resource

5. Select **SAVE** to complete the operation.

The corresponding child pool on your storage system is resized as requested.

Deleting a storage resource

To delete a storage resource, complete the following steps:

1. Hover your cursor over the top of the storage system.
2. Select the **Settings** button and then click **List pools**, as shown in Figure 5-10 on page 60.
3. Select the wanted storage resource and then click the **Remove** button, as shown in Figure 5-13.

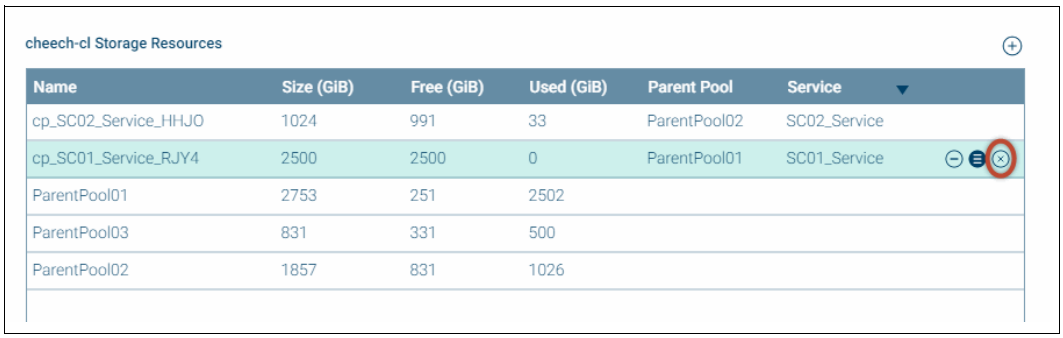

Figure 5-13 Deleting a storage resource

4. Click **OK** to delete the storage resource, or **Cancel** to cancel the operation.

Performing this task deletes a child pool from the storage system. The parent pool reclaims the capacity that was allocated to the child pool. After a child pool has been deleted from the storage system, you can also remove the associated storage service and space from IBM Spectrum Control Base Edition, if wanted. For instructions on completing these steps, see “Deleting a storage service” on page 57 and “Deleting a storage space” on page 56.

Adding storage resources to storage spaces

You might want to add new storage resources to an existing storage space, to increase the capacity of the VVol datastore to which it is mapped. To add another storage resource to your storage space, complete the following steps:

1. From the Spaces and Services pane, select the storage space that you want to expand.
2. Select the VVols-enabled storage service for your chosen storage space. The storage service is highlighted when selected.
3. Create a storage resource (child pool) using the method described in “Creating a storage resource” on page 59.
4. The new storage resource is automatically added to the storage space.

Moving storage resources between storage spaces

You might want to move a storage resource between storage spaces to add or remove capacity from the corresponding VVol datastores. To do so, complete the following steps:

1. On the Spaces and Services pane, use the spaces drop-down menu to select the storage space that you want to remove the storage resource from.
2. Hover your cursor over the VVols-enabled storage service for your chosen storage space, then click the **Edit** button and select **Resources**. Figure 5-14 shows the location of the **Edit** button for a storage service.

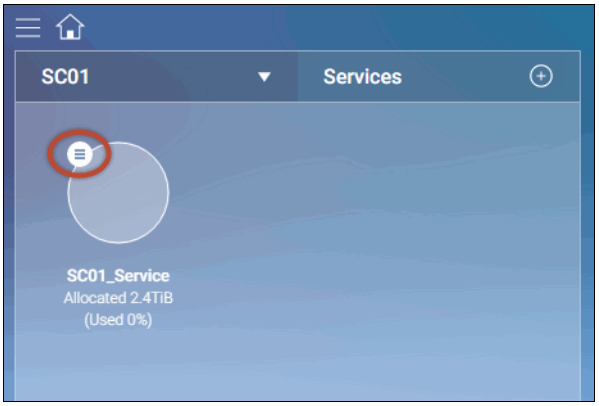

Figure 5-14 Editing a storage service

3. Select the storage resource that you want to relocate and then click the **Resource Detach** button, as shown in Figure 5-15.

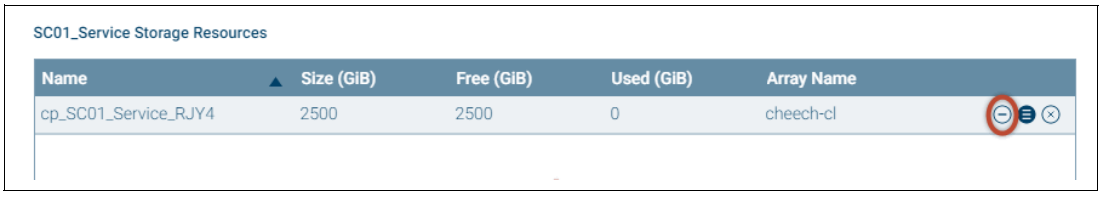

Figure 5-15 Detaching a storage resource from a storage service

4. Click **OK** to detach the storage resource from the storage service, or **Cancel** to cancel the operation.
5. Use the Spaces drop-down menu to select the storage space that you want to allocate the storage resource to. Select the VVols-enabled storage service for this space. The service will be highlighted after selected, as shown in Figure 5-16.

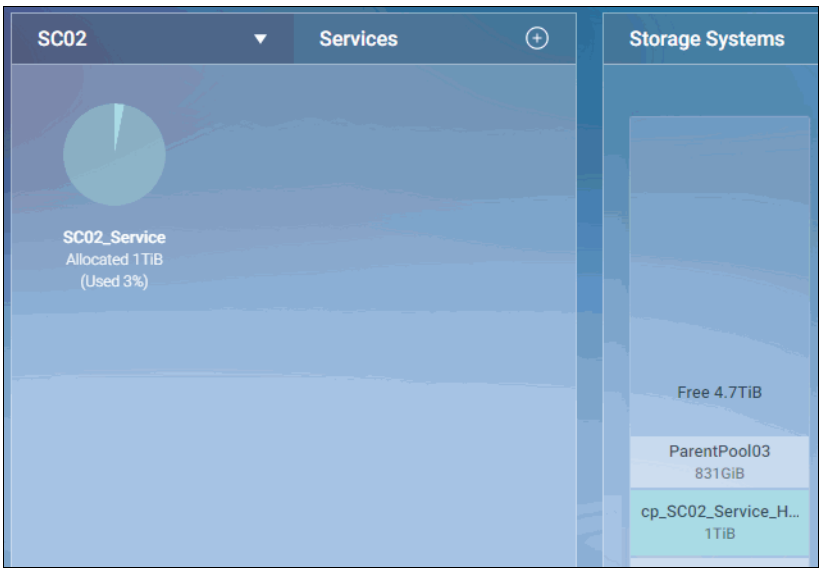

Figure 5-16 Selecting the destination storage space and storage service

6. Hover your cursor over the top of the storage system.
7. Select the **Settings** button and then click **List pools**, as shown in Figure 5-17.

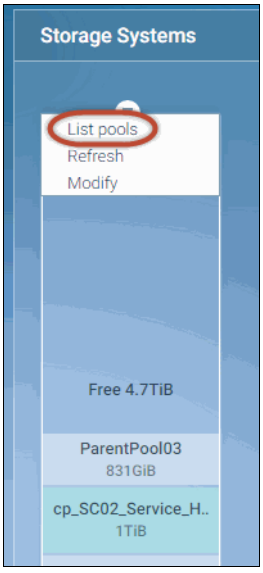

Figure 5-17 List all storage resources (parent and child pools are displayed)

8. Select the storage resource to be relocated and then click the **Resource Attach** button, as shown in Figure 5-18.

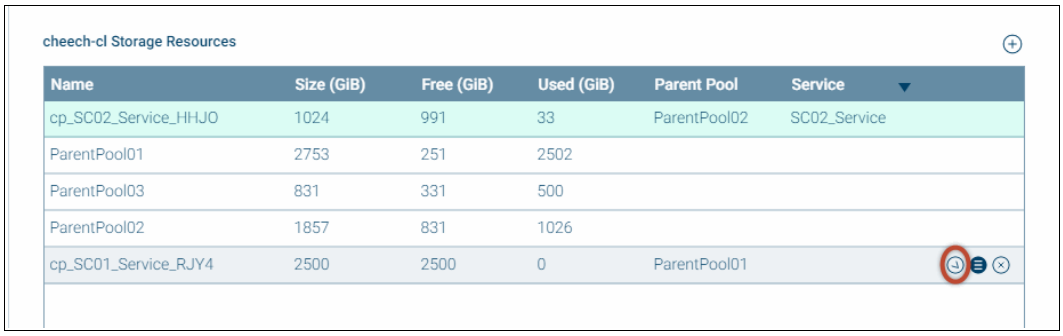

Figure 5-18 Selecting a storage resource and attaching it to a new storage service

9. Click **OK** to attach the storage resource to the new storage service, or **Cancel** to cancel the operation.

The storage resource is now associated with the destination storage space and storage service, as shown in Figure 5-19. The child pool remains unchanged from the perspective of the underlying storage system.

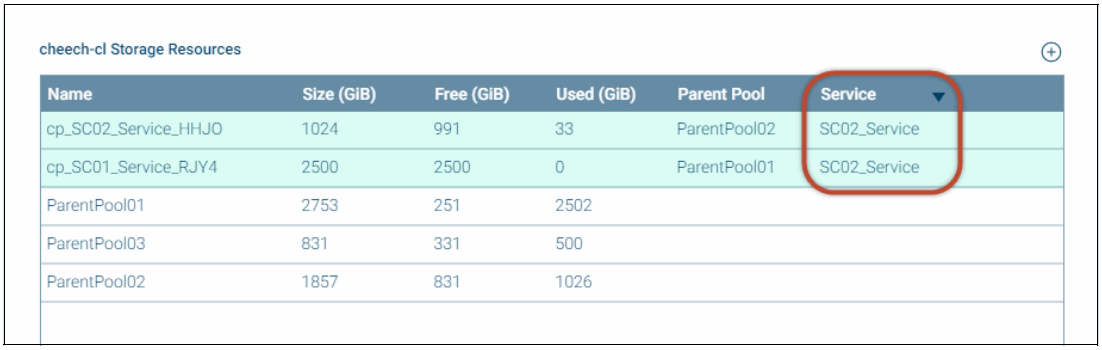

Figure 5-19 Our storage resource is now associated with SCO2_Service

5.3.4 Adding more SCB instances

We strongly advise configuring a standby instance of IBM Spectrum Control Base Edition, for high availability (HA) of virtual machine management. For more information about this suggested practice, and background on HA Groups, see 6.2.1, “High availability” on page 77.

To configure a standby SCB instance, follow the instructions in 4.4, “Configuring additional SCB servers for high availability” on page 39.

When configured, additional SCB instances must be registered as storage providers on the vSphere Web Client. For more information about this step, see “Adding storage providers” on page 67.

5.3.5 Managing Storage Systems

The following sections describe some common storage systems management functions.

Adding more storage systems to SCB

It is possible to manage multiple storage systems using a single instance of IBM Spectrum Control Base Edition. To add another storage system, complete the following steps:

1. Ensure that the relevant steps have been completed on the storage system before attempting to register it with SCB. For more information about the required steps, see 4.1, “Configuring the storage system” on page 18.
2. Navigate to the Storage Systems pane and select the **Add** button.
3. Enter the Hostname of your storage system and select the storage system Type.
4. Click **Add** to complete the operation.

Figure 5-20 shows the details that must be specified to add a new storage system to SCB.

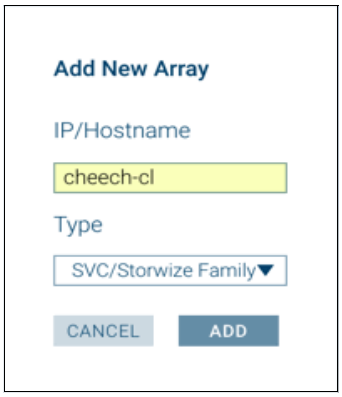

Figure 5-20 SCB Add New Array dialog

Refreshing storage system information about SCB

If you add new objects (such as parent pools) to the storage system, you might want to manually trigger a refresh of the IBM Spectrum Control Base Edition information to update this information immediately. To initiate a refresh, complete the following steps:

1. Hover your cursor over the top of the storage system and select the **Edit** button, and then click **Refresh**, as shown in Figure 5-21.

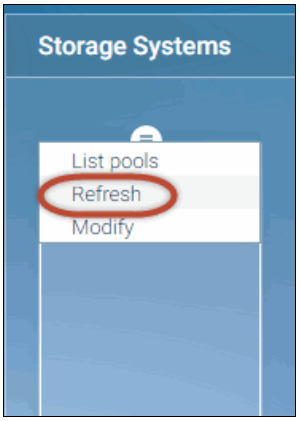

Figure 5-21 Refreshing the storage system information

2. This triggers a repopulation of storage system information in IBM Spectrum Control Base Edition. Note that you might need to manually refresh the GUI before these updates are displayed.

5.4 VMware vSphere Web Client tasks

The following tasks should be performed using the vSphere Web Client. The following sections focus on vSphere features relevant to the storage administrator, and do not aim to provide an extensive list of the VVol operations that can be run from the vSphere Web Client. For information about performing tasks in the vSphere environment, see the following website:

http://kb.vmware.com

5.4.1 Configuring multipathing

Before the release of vSphere 6.0 and VVols, it was considered the leading practice to be using RoundRobin as the preferred I/O Path Selection Policy (PSP). The benefit of RoundRobin over Most Recently Used (MRU) and Fixed PSP is that it rotates through all available paths, enabling the distribution of load across all of the configured paths. If a path fails and then becomes available again, it is automatically included back into the PSP.

With the introduction of protocol endpoints (PE) with vSphere 6.0 and VVols, we still consider the use of RoundRobin to be the leading practice for PSP.

Additional information: Storage systems powered by IBM Spectrum Virtualize support *VMW_SATP_ALUA*. Essentially, Storage Array Type Plug-in (SATP) enables load balancing across multiple paths, in addition to determining which PSP to assign. With asymmetric logical unit access (ALUA), we can optimize paths that have direct access to a volume.

5.4.2 Managing queue depths

With the introduction of VVols in vSphere 6.0, VMware increased the default maximum queue depth on ESXi hosts from 64 to 128 to facilitate the increased number of volumes. The vSphere administrator can adjust these values to suit the Virtual Volumes environment and the loads being applied to the storage system and host bus adapters (HBAs).

There is no magic formula that we can apply here that suits all customer environments. As a general guideline, we aim to share resources equally between all paths from all hosts. Using VMware Distributed Resource Scheduler (DRS), we can ensure that this happens automatically. As such, we advise that you run with the default settings and monitor your resource use accordingly.

Additional information: For more information about Queue Depths, see *IBM SAN Solution Design Best Practices for VMware vSphere ESXi*, SG24-8158 and VMware Knowledge Base article 1267.

5.4.3 Managing storage providers

When configured, IBM Spectrum Control Base Edition must be added to the vSphere Web Client as a storage provider.

Adding storage providers

To add a storage provider, complete the following steps:

1. Ensure that the relevant configuration steps have been performed on your IBM Spectrum Control Base Edition server. For the first SCB instance, see 4.2, “Configuring IBM Spectrum Control Base Edition” on page 22. For standby SCB instances, see 4.4, “Configuring additional SCB servers for high availability” on page 39.

   Remember: You require the VASA credentials that were set during your IBM Spectrum Control Base Edition configuration.

2. Start at the Home page of the vSphere Web Client.
3. Select **vCenter Inventory Lists** → **vCenter Server** and then select your vCenter Server.
4. Select **Manage** → **Storage Providers**.
5. Select the green Plus sign (+) icon to add a new storage provider.
6. Enter a name for your SCB instance.
7. Enter the fully qualified domain name (FQDN) for the SCB instance that you want to register, in addition to the port and service information shown in Example 5-1.

   Example 5-1 Port and service information

   ```
   https://<host.domain>.com:8443/services/vasa
   ```

8. Enter a user name and password. These details should match the *VASA credentials* that were configured on IBM Spectrum Control Base Edition.
9. Select **OK**.
10. Select **Yes** to confirm that you trust the certificate.

Rescanning storage providers

In some rare scenarios, it might be helpful to rescan your storage providers to force an update of IBM Spectrum Control Base Edition information in the vSphere Web Client. For example, rescanning storage providers can help to pull through information about a newly created storage space, if there is a delay in the synchronization of the vSphere Web Client and IBM Spectrum Control Base Edition.

To rescan a storage provider, complete the following steps:

1. Start at the Home page of the vSphere Web Client.
2. Select **vCenter Inventory Lists** → **vCenter Server**, and then select your vCenter Server instance.
3. Select **Manage** → **Storage Providers**.

4. Highlight your active IBM Spectrum Control Base Edition instance and select the **Rescan** icon, as shown in Figure 5-22.

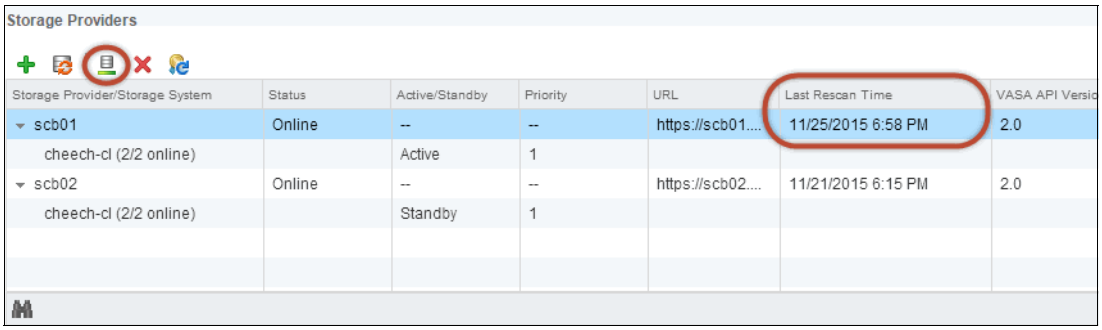

Figure 5-22 Rescanning Storage Providers

5. Confirm that the **Last Rescan Time** field is updated correctly, as shown in Figure 5-22.

5.4.4 Managing VVol datastores

Adding VVol datastores is essential to a Virtual Volumes implementation. Creating the one-to-one mapping between a storage space and a datastore object enables the provisioning of virtual machines on VVol storage.

Creating a VVol Datastore

To create a VVol Datastore, complete the following steps:

1. Start at the Home page of the vSphere Web Client.
2. Select **Storage**.

3. Select your Datacenter, and then select **Actions** → **Storage** → **New Datastore**, as shown in Figure 5-23.

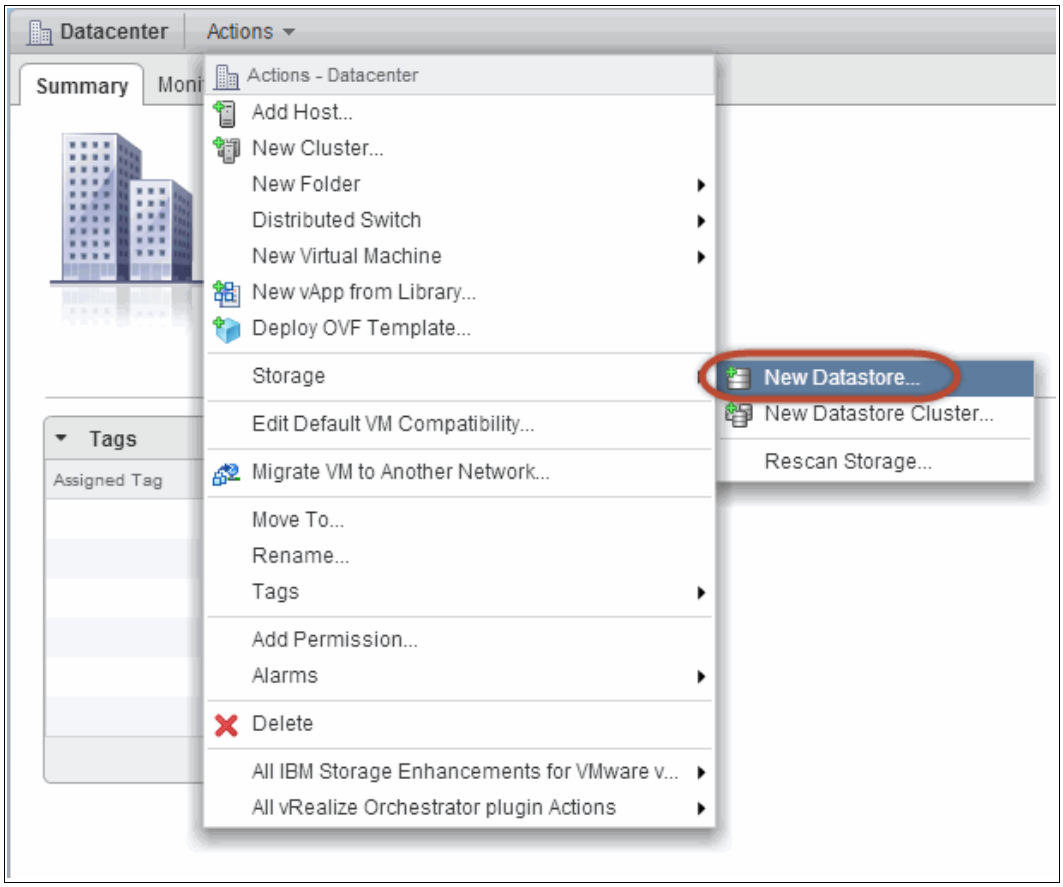

Figure 5-23 Adding a new datastore using the vSphere Web Client

4. Confirm the location of the datastore, and then select **Next**.
5. Select **VVol** as the datastore type, and then select **Next**.
6. Enter a datastore name and select your Storage Space from the list of backing storage containers, and then select **Next**.
7. Select the hosts that will have access to this datastore, and then select **Next**.

> **Requirement:** All of the selected hosts must be registered as VVols-enabled hosts on your storage system. For information about this task, see 5.2.1, “Managing VMware ESXi hosts” on page 53.

8. Confirm that all details are correct, and then select **Finish** to complete the operation.

Figure 5-24 depicts a summary of the settings that are used when creating a VVol datastore.

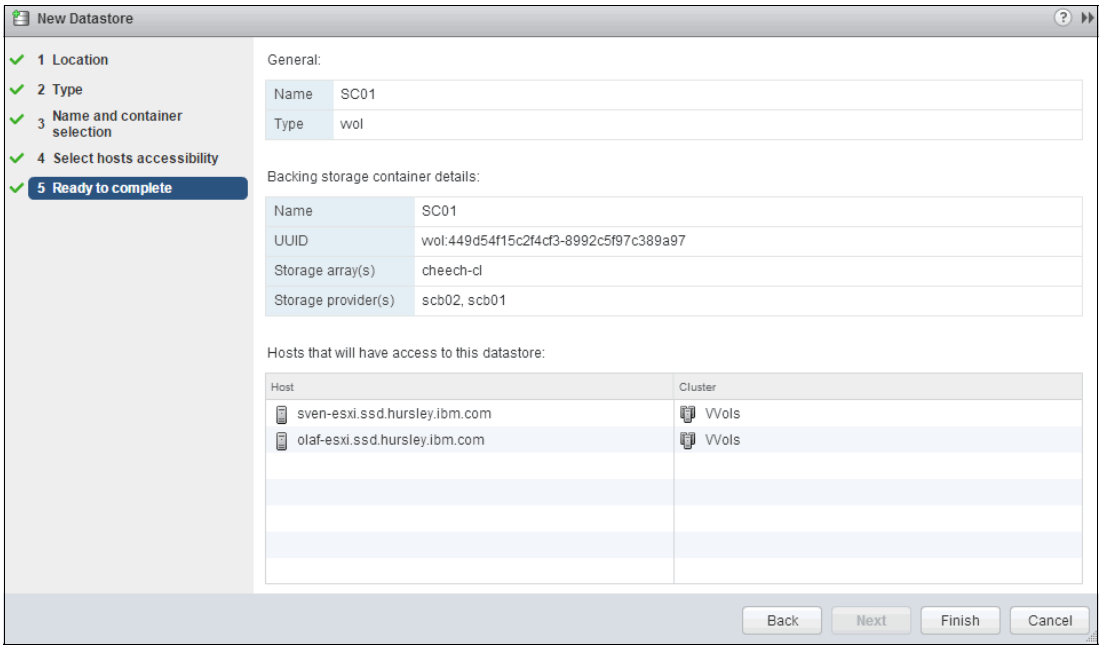

Figure 5-24 Creating a VVol datastore

Unmounting a VVol datastore

Important: Before deleting a VVol datastore, migrate all VMs onto another datastore.

Before deleting the storage-side objects associated with a VVol datastore, the datastore object must be removed from the vSphere Web Client. This process is known as *unmounting a datastore*, and is used to break the association between the datastore object and the storage space behind it. To unmount a VVol datastore, complete the following steps:

1. Start at the Home page of the vSphere Web Client.
2. Select **Storage**.
3. Select the VVol Datastore that you want to delete, and then select **Actions** → **Unmount Datastore**.

4. Select *all* hosts, and then select **OK**, as shown in Figure 5-25.

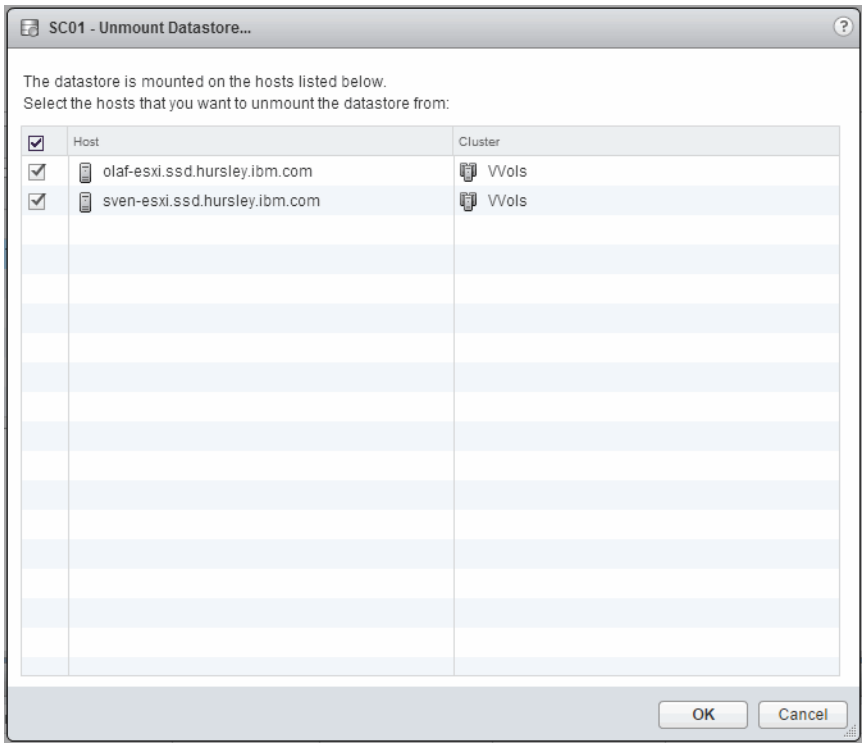

Figure 5-25 Unmounting a VVol datastore

5. If wanted, the child pool previously associated with the datastore can now be deleted, to reclaim this capacity for the parent pool. To complete this task, see “Deleting a storage resource” on page 61.

5.4.5 Working with Storage Policy-Based Management

Storage Policy-Based Management (SPBM) aides provisioning of virtual machine storage and enhances possibilities for automation. For more information about working with SPBM, see Chapter 7, “Storage Policy-Based Management” on page 83.

5.4.6 Migrating virtual machines onto VVol datastores

When VVol datastores are configured on the system, you might want to migrate existing virtual machines onto VVol storage.

The storage administrator *does not* need to perform any tasks on the storage system in this scenario. Instead, the vSphere administrator should use vSphere tasks, such as cold migration or Storage vMotion to relocate the virtual machine storage. This process triggers IBM Spectrum Control Base Edition to run the requisite commands on the storage system without manual interaction from the storage administrator.

Additional information: Migrating VMs with raw device mappings onto VVol datastores:

At the time of writing, Raw Device Mappings (RDMs) are not compatible with VVols. When migrating VMs with attached RDMs onto VVol datastores, you should convert any volumes attached as RDMs into VVols as part of the migration.

Note that cold migration can be used to convert either virtual RDMs (vRDMs) or physical RDMs (pRDMs) into VVols, whereas Storage vMotion can *only* be used to convert vRDMs (and *not* pRDMs) into VVols.

For further information about this topic, see the VMware Knowledge Base:

https://ibm.biz/rdm-compatibility

6

Considerations and best practices

This chapter describes several important factors that must be considered when planning an implementation of Virtual Volumes (VVols). By exploring these topics in more detail, we aim to highlight the decisions that need to be made to maximize the potential of Virtual Volumes and meet the needs of your environment.

This chapter provides information about the following topics:

- 6.1, “Storage system considerations” on page 74
- 6.2, “IBM Spectrum Control Base Edition considerations” on page 77

6.1 Storage system considerations

The following topics should be considered when planning for your VVols implementation.

6.1.1 Defining Parent Pools

Because the storage for VVols is allocated as a child pool, it is important to consider the structure of the parent pools from which these child pools are allocated. In the following sections, we present two contrasting approaches to defining parent pools, and describe how their usage might influence our VVols environment.

Defining parent pools based on drive class

You might want to define parent pools based on the underlying drive class. This approach enables the allocation of child pools from a specific tier of storage. At the vSphere level, these child pools would serve as the backing storage container for distinct VVols datastores. This essentially enables you to determine the class of storage for individual virtual volumes when provisioning virtual machines (VMs).

We could use vSphere policies (for example, Gold, Silver, and Bronze) to easily select the appropriate class of storage when provisioning VMs. See Figure 6-1 for an example of this configuration.

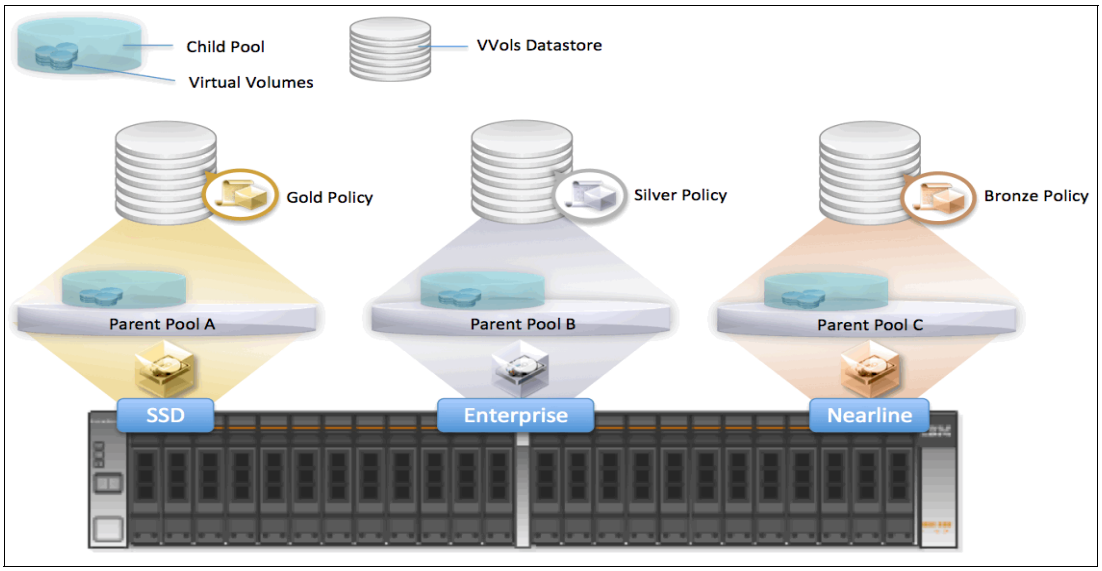

Figure 6-1 Depicts a simplistic VVols environment, where parent and child pools are segregated by drive class

With the introduction of Virtual Volumes, by defining a range of storage services on IBM Spectrum Control Base Edition (SCB), policies can become far more interesting and useful than the simple Gold, Silver, and Bronze model described previously.

Each of the policies described previously could be subdivided further. For example, we could divide our solid-state drive (SSD) parent pool into two distinct child pools. One child pool linked to an encrypted storage service, and the other associated with an unencrypted storage service. This provides the vSphere administrators with the flexibility to provision virtual machines on storage that matches the requirements of the application, on a per-VM basis.

For further information, see Chapter 7, “Storage Policy-Based Management” on page 83.

Using IBM Easy Tier with Virtual Volumes

An alternative to the approach described previously would be to define parent pools with a combination of drive classes, and enable the IBM Easy Tier® feature. By monitoring the heatmap of a volume's extents, Easy Tier can intelligently optimize the usage of storage by automatically migrating these extents onto the most appropriate storage tier.

Because VVols are simply a special kind of volume, Easy Tier can manage their extents in an identical fashion. A "hot" (frequently used) extent of a VVol is promoted to faster storage, such as SSD, and a "cold" (infrequently used) extent of a VVol is moved onto slower drives.

A VVols implementation that takes advantage of Easy Tier can provide greater simplicity for the storage administrator. By defining a child pool within an Easy Tier-enabled parent pool, we enable the storage system to flexibly manage the extents of any VVols created therein.

This removes the requirement for a choice of storage class when the vSphere administrator initially provisions the VM. Such an approach can also minimize the need for Storage vMotion tasks, because Easy Tier eliminates the requirement to manually migrate VVols onto faster or slower storage as the needs of an application change.

Figure 6-2 demonstrates a VVols configuration, based on a single parent pool, with Easy Tier enabled.

Note: Easy Tier also provides benefits within a single-tiered pool. When enabled, Easy Tier automatically balances the load between managed disks (MDisks) to optimize performance.

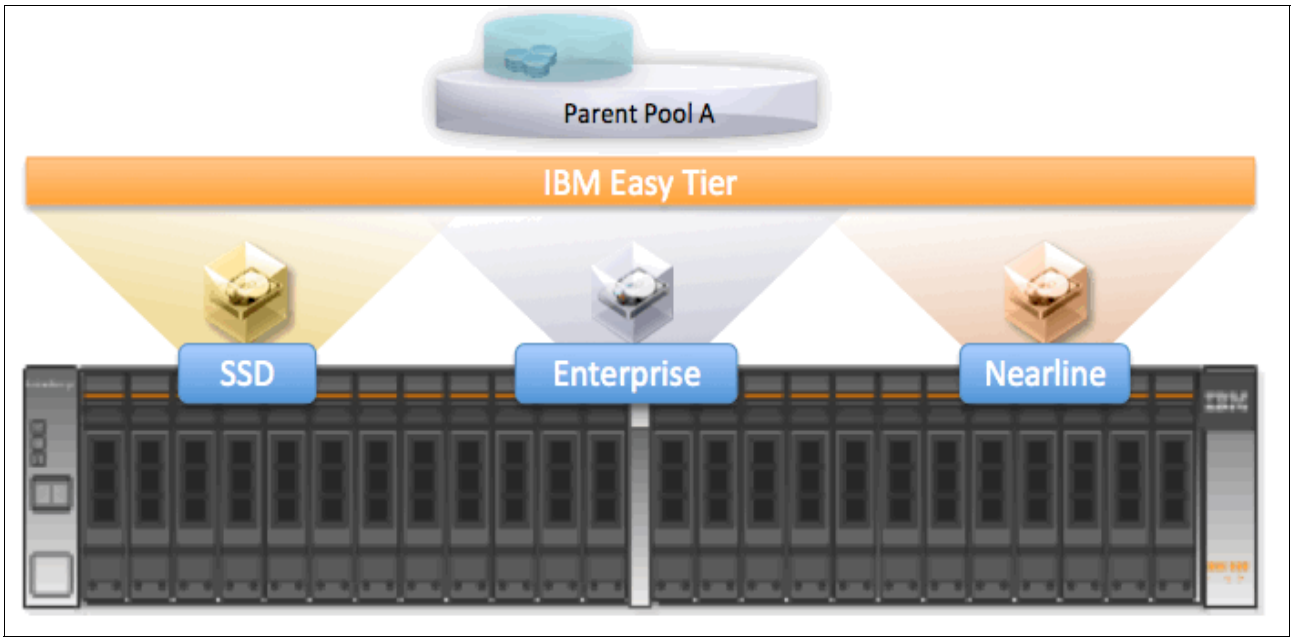

Figure 6-2 The simplified approach to VVol provisioning that can be implemented by enabling Easy Tier

6.1.2 Mirroring the utility volume

A utility volume is created on the storage system as part of the configuration of Virtual Volumes. IBM Spectrum Control Base Edition creates and manages a database on the utility volume. The virtual machine metadata stored in this database is critical to IBM Spectrum Control Base Edition operations and the VVols environment.

Because the availability and integrity of the utility volume is fundamental to the Virtual Volumes environment, we strongly suggest storing a mirrored copy of the volume for redundancy. If possible, store the mirrored copy in a second storage pool that is in a separate failure domain. For example, use a storage pool that is made from MDisks that are presented from different storage systems, or a different input/output (I/O) group.

For instructions on creating a mirrored copy of the utility volume as part of the initial configuration, see 4.1, “Configuring the storage system” on page 18.

If you have already configured Virtual Volumes, but want to retrospectively mirror the utility volume, see 6.2.2, “Defining storage spaces” on page 78.

6.1.3 Performing an upgrade on a storage system with Virtual Volumes enabled

When performing a code upgrade, the storage system intentionally limits the tasks that a user can run. This is a protective measure to ensure that the upgrade is not disrupted. Because managing virtual machines on VVols datastores requires the running of commands on the system, these same restrictions apply in the Virtual Volumes environment.

Therefore, virtual machine management tasks (for example, powering off a virtual machine) fail when run from the vSphere Web Client. Note that any automated services, such as VMware High Availability (HA) and Distributed Resource Scheduler (DRS), are also affected, because they also send system commands using IBM Spectrum Control Base Edition.

Therefore, it is important to plan an upgrade alongside your vSphere administrator. To ensure a smooth upgrade in your VVols environment, consider the following suggestions:

- Plan your upgrade for a time when virtual machine management is not required.
- Ask your vSphere administrator to temporarily disable services, such as VMware HA and DRS, for clusters that use VVols datastores. Inform your vSphere administrator when the upgrade is complete so that these services can be reenabled.
- Warn your vSphere administrator not to run management tasks on VMs stored on VVols datastores during an upgrade, because this results in task failures on the vSphere Web Client.

Tip: After performing an upgrade, it is possible that the vSphere Web Client will mark cold VMs as inaccessible. This simply means that ESXi hosts were unable to start a new binding to these virtual machines (expected during a code upgrade) and should not cause alarm.

To recover management of these virtual machines, the vSphere administrator should remove the affected virtual machines from the inventory and then re-add them.

6.2 IBM Spectrum Control Base Edition considerations

In this section, we describe some IBM Spectrum Control Base Edition considerations.

6.2.1 High availability

To ensure high availability of virtual machine management, two or more instances of IBM Spectrum Control Base Edition can be configured. By pointing multiple SCB instances at the same database, we are able to form a high availability Group (HA Group).

Remember: The database shared by SCB instances in an HA Group is stored on the utility volume. See 3.9, “Utility volume” on page 14.

Figure 6-3 shows the basic concepts of the HA Group.

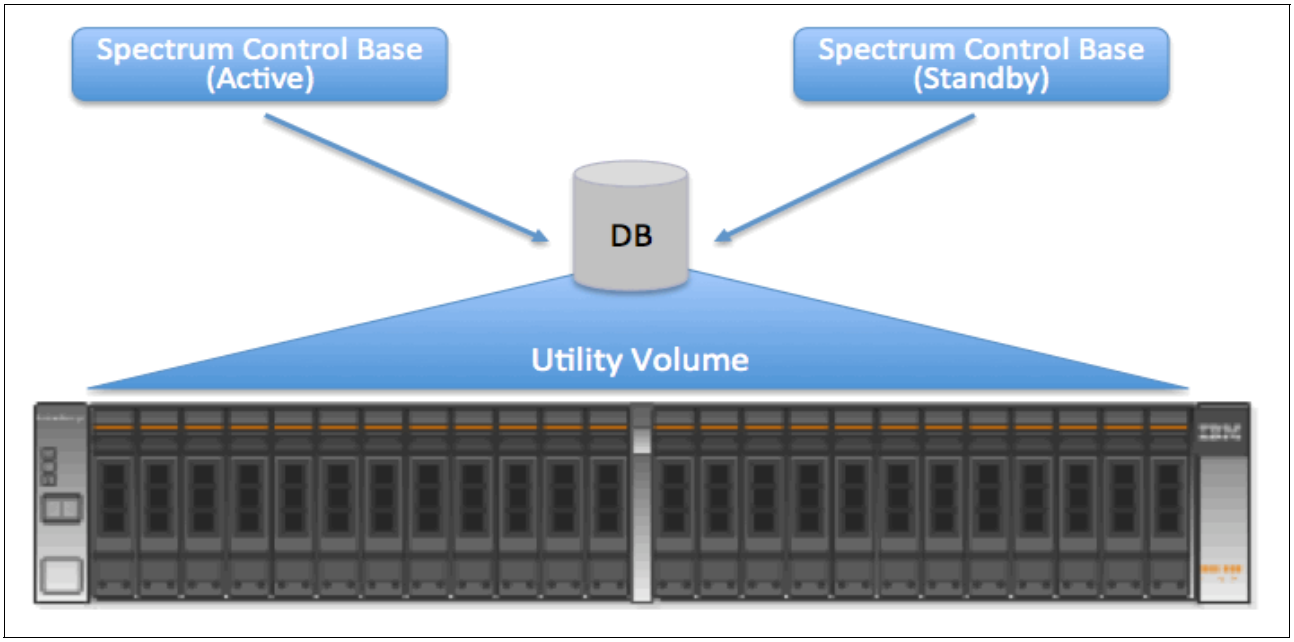

Figure 6-3 The members of the SCB HA group accessing the DB stored on the utility volume, which is created as part of the VVols configuration

Regardless of your configuration or environment, we strongly suggest deploying at least one standby instance of IBM Spectrum Control Base Edition.

IBM Spectrum Control Base Edition high-availability design

SCB’s high-availability design works on an Active/Standby model. In such an implementation, the first instance of SCB registered with vCenter automatically becomes the active provider. If there is an issue with the first instance of IBM Spectrum Control Base Edition (for example, if the virtual machine becomes unavailable, or the SCB services stop responding), vCenter automatically triggers a failover to the standby provider.

This process ensures that there is minimal downtime of virtual machine management. This is particularly important when using VMware features, such as HA and DRS.

Important: The high availability of IBM Spectrum Control Base Edition is not to be confused with VMware's HA feature. When used regarding SCB, the phrase *high availability* refers to the management of virtual machines (for example, powering a VM on or off). It does not refer to the high availability of a virtual machine or application within the confines of vSphere.

Similarly, it is not relevant to a virtual machine's ability to perform I/O to a volume. I/O continues to run using the data path (see Figure 3-2 on page 12) regardless of the control path state (see Figure 1-1 on page 2).

Deploying SCB servers for high availability

To maximize the benefits of the Active/Standby model and avoid single points of failure in your environment, we would also advise you to take the following actions:

- Deploy SCB servers as virtual machines, on ESXi clusters with VMware HA enabled.
- Select shared VMware file system (VMFS) datastores as the storage for the virtual machines that host SCB.

  Important: SCB virtual machines should *not* be deployed onto VVol datastores.

- If VMware DRS is enabled for your ESXi cluster, define VM-VM affinity rules such that your Active and Standby SCB servers are ensured to be hosted on separate physical ESXi servers. This eliminates the risk of your VVols environment being inadvertently exposed to a single point of failure (regarding virtual machine management). For more information about VM-VM affinity rules, see VMware's Knowledge Base:

  http://kb.vmware.com

Configuring Spectrum Control Base instances for High Availability

SCB instances must be combined under the same High Availability Group (HA Group) to use the Active/Standby model. It is essential that both instances are registered with the same HA Group Name, to enable them to access the same database. For specific steps to configure an active IBM Spectrum Control Base Edition instance, see 4.2, "Configuring IBM Spectrum Control Base Edition" on page 22.

For specific steps to configure a standby IBM Spectrum Control Base Edition instance, see 4.4, "Configuring additional SCB servers for high availability" on page 39.

6.2.2 Defining storage spaces

Storage spaces are the logical entities that you present to your vSphere administrator for use as a VVols datastore. Consequently, when defining storage spaces, consider the following questions:

- How many separate VVols datastores will your vSphere administrator need?

  Because storage spaces and VVols datastores are mapped one-to-one, you need to use IBM Spectrum Control Base Edition to define the correct number of spaces for your vSphere environment. At the time of writing, one VVols-enabled storage service can be created per storage space.

  By creating multiple VVols-enabled storage services and storage spaces, you dictate that multiple VVols datastores are required. If you only have a requirement for one or two types of services, you might prefer to implement a simple environment with only one or two VVol datastores.

- What kinds of storage capabilities should you offer to your vSphere administrator?

  It's important to consider whether your environment requires specific storage capabilities for virtual machine provisioning. There might be "must-have" options for the vSphere administrator (for example, encrypted storage). Alternatively, you might want to prevent your vSphere administrator from using certain capabilities. For example, you might want to keep tighter control over which volumes can be compressed.

- How much storage does your vSphere administrator need?

  Although the Virtual Volumes framework transfers some of the administrative tasks from the storage administrator to the vSphere administrator, you still retain complete control over the amount of storage that will be used for Virtual Volumes.

Defining storage services

By defining storage services, you are able to decide which capabilities are offered to the vSphere administrator. The storage services model enables the storage administrator to be as restrictive or open as they want in their approach.

If you want tight control over your storage system, and only want to allow the vSphere administrator to use a very small range of storage capabilities, this is possible. Simply apply restrictions to the relevant options when creating the storage service and allow SCB to provision volumes under the framework you have provided. Figure 6-4 shows an example of a restrictive storage service.

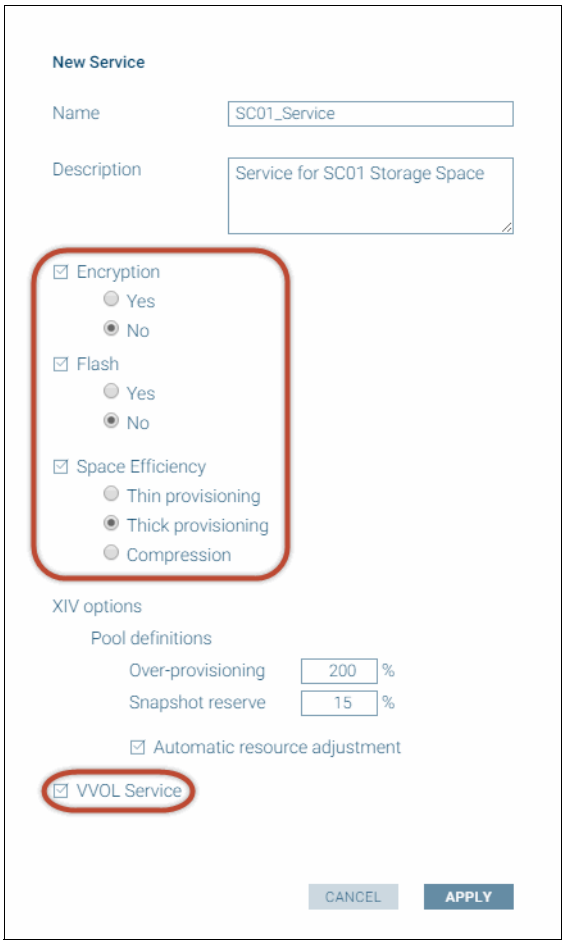

Figure 6-4 The options selected here ensure that any VVols created on the associated child pool are thick-provisioned

Alternatively, you might want to leave the choices relating to storage capability completely open to the vSphere administrator. You can do this by creating generic storage services without restrictions.

Figure 6-5 shows a storage service without capability restrictions.

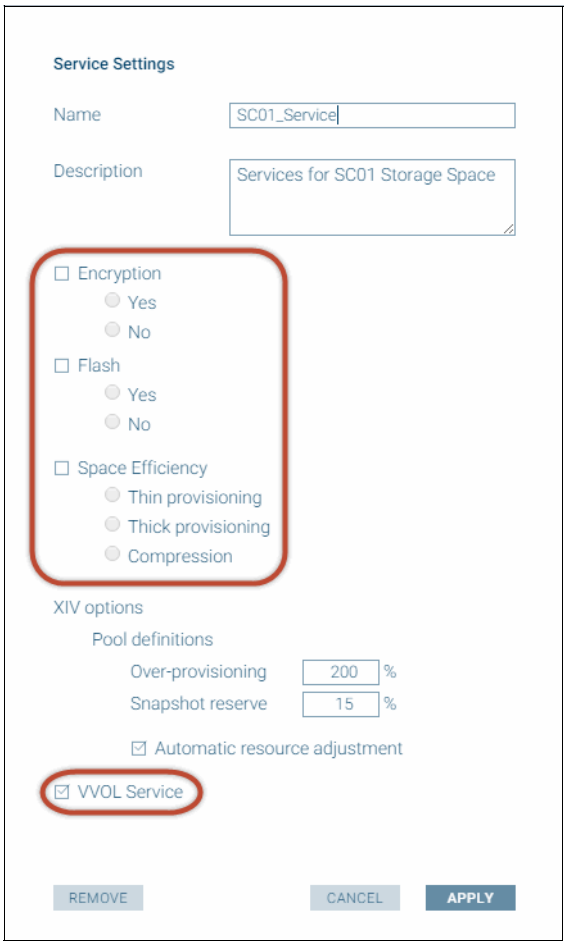

Figure 6-5 By leaving all capability options cleared, you allow the vSphere administrator to choose from a range of options when provisioning VVols

Each storage service is linked to a child pool (or multiple child pools), so even if you offer more extensive capabilities to the vSphere administrator, you never need to give up too much control of the storage system. For example, you can offer a Flash-enabled storage service, but use child pool allocation to limit this service to a small capacity of storage.

After you have decided which range of capabilities you offer to the vSphere administrator and how to divide these capabilities between storage services, you can begin associating a quota of storage to each storage service in the form of a child pool.

Defining child pools

Ultimately, child pools provide the quota of storage that is available to vSphere administrators in the form of a VVols datastore. Therefore, perhaps the most obvious consideration when defining a child pool is its capacity.

Although capacity should be considered, it's useful to note that a child pool can be resized after creation. The size can be increased (subject to available space in the parent pool) or decreased, so if the original choice is inappropriate, there are not long-term implications. Consequently, there isn't really a best practice here: Define the child pool size that you think you need, and then resize if necessary.

Tip: You might initially want to define a small child pool for your test environment, and then gradually increase the size as you become more familiar with the technology.

Similarly, to the question of child pool size, there is no single recommendation on the number of child pools needed for a Virtual Volumes environment. This decision really depends on the needs of your infrastructure. However, it is worth considering the following points when provisioning child pools:

- Child pools can map to storage spaces and storage services one-to-one *or* many-to-one. You can choose whether there are one or more child pools per VVols datastore. For a description of a use-case of a many-to-one relationship, see Chapter 7, "Storage Policy-Based Management" on page 83.
- Child pools can be resized on demand. Provisioning additional child pools is not necessary if you simply want to resize an existing datastore.

6.2.3 Using the IBM Storage Enhancements for VMware vSphere Web Client in a Virtual Volumes environment

In addition to fulfilling the role of vSphere APIs for Storage Awareness (VASA) provider in a VVols environment, IBM Spectrum Control Base Edition also offers an additional plug-in for storage systems running IBM Spectrum Virtualize. This plug-in is known as the IBM Storage Enhancements for VMware vSphere Web Client.

IBM Storage Enhancements for VMware vSphere Web Client

When configured, the IBM Storage Enhancements plug-in enables users to provision volumes on the storage system from within the vSphere Web Client. These volumes can be easily mapped to multiple ESXi hosts and can then be attached as Raw Device Mappings (RDMs), or formatted as VMFS datastores for virtual machine storage.

Figure 6-6 shows volume provisioning from the vSphere Web Client, using the IBM Storage Enhancements plug-in.

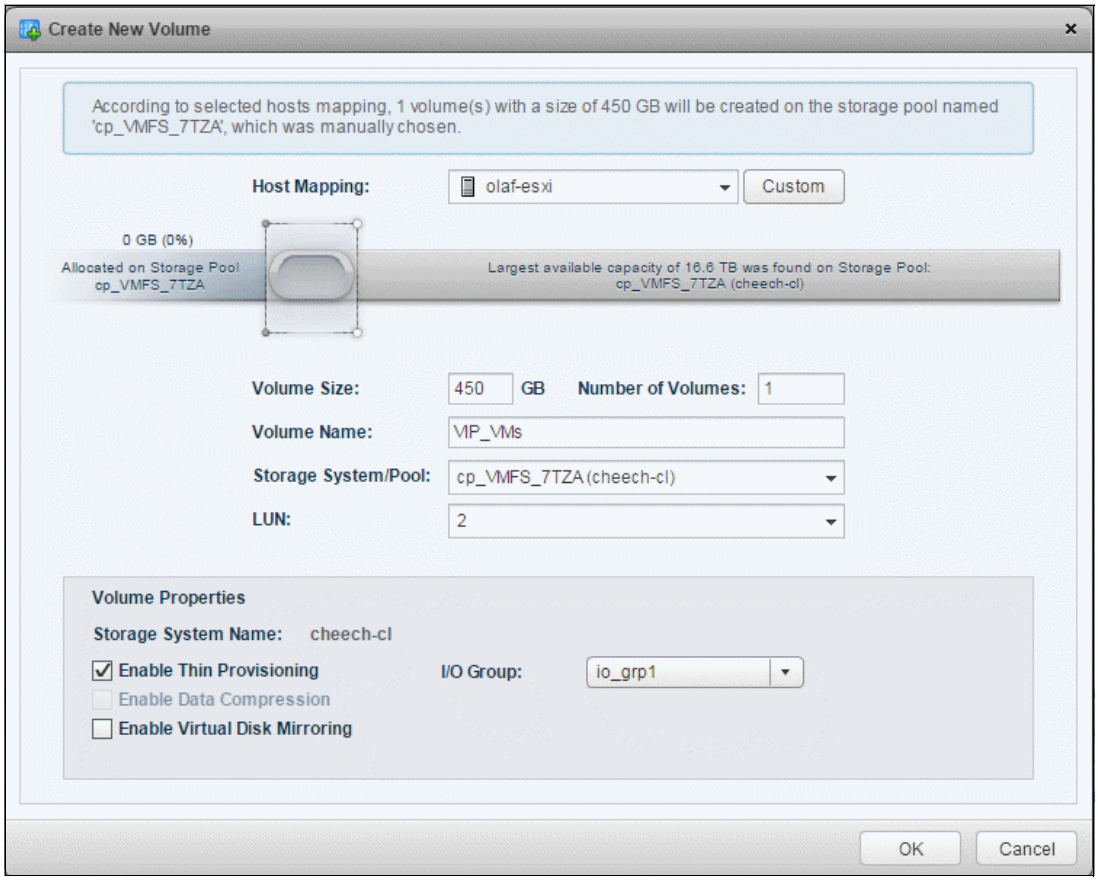

Figure 6-6 IBM Storage Enhancements plug-in: Create New Volume panel

For more information about configuring the IBM Storage Enhancements plug-in, see the IBM Knowledge Center for IBM Spectrum Control Base Edition:

https://ibm.biz/ibm-scb-kc

Using the IBM Storage Enhancements with SCB HA for VVols

You might want to use the IBM Storage Enhancements with your VVols environment. The storage system supports both features.

However, note that the IBM Storage Enhancements plug-in cannot be implemented in Active/Standby mode. Therefore, we do not advise running the IBM Storage Enhancements on the same SCB servers that are used in your VVols environment.

7

Storage Policy-Based Management

Storage Policy-Based Management (SPBM) is an integral component of the Software Defined Datacenter model. In this chapter we show how it can used to automatically select appropriate storage from pre-allocated pools based on application requirements.

This chapter provides information about the following topics:

- 7.1, “Overview of Storage Policy-Based Management” on page 84
- 7.2, “Storage capabilities advertised by SCB” on page 85
- 7.3, “Defining policies in vSphere Web Client” on page 87
- 7.4, “Use case scenarios and examples” on page 88

7.1 Overview of Storage Policy-Based Management

Storage Policy-Based Management (SPBM) is an integral component of the Software Defined Datacenter model. It is used to automatically select appropriate storage from pre-allocated pools based on the specific capabilities and application requirements.

SPBM can increase the flexibility and control of virtual machine storage provisioning, offering a sophisticated level of automation for large vSphere environments.

Figure 7-1 shows an overview of SPBM.

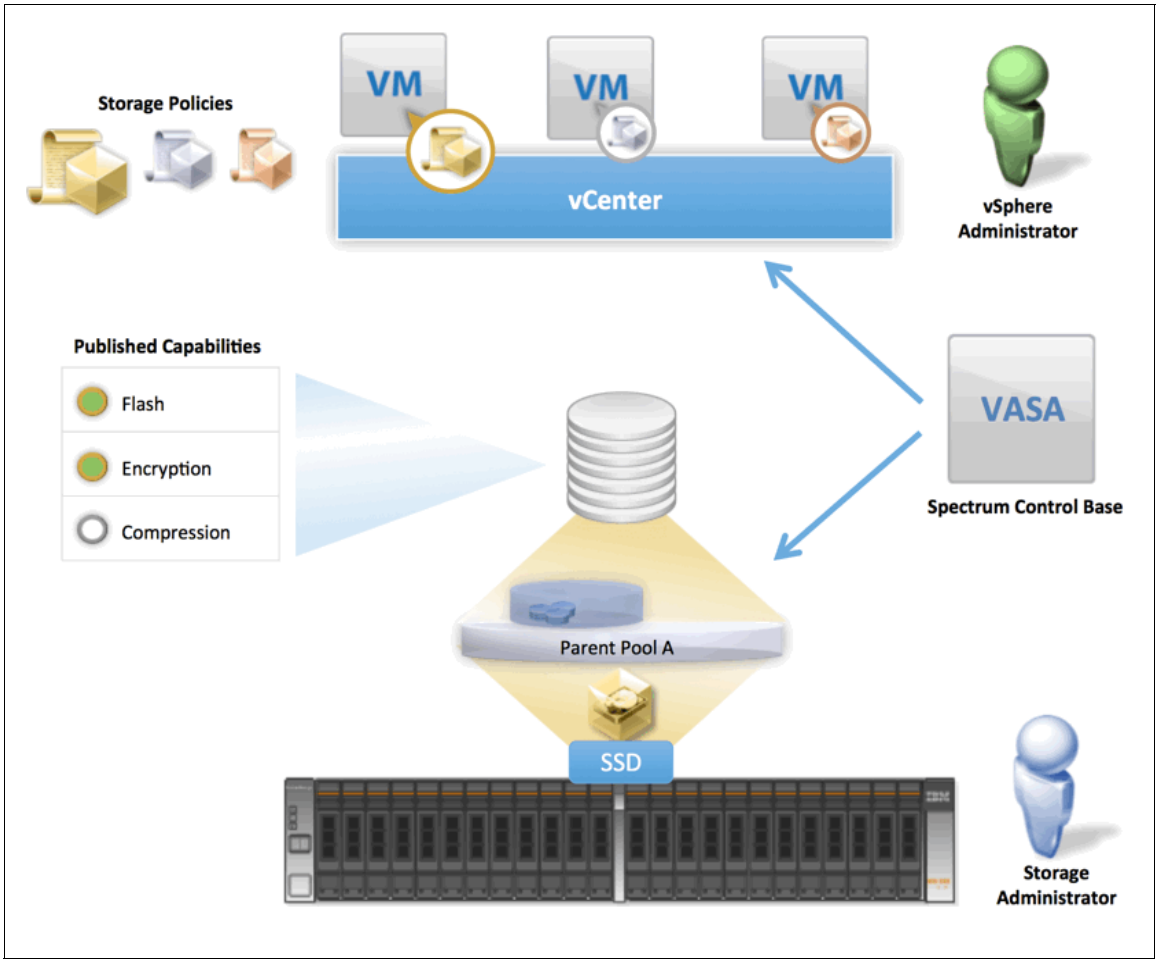

Figure 7-1 An overview of SPBM

There are two areas of focus when configuring SPBM:

- The storage capabilities provided by the storage system that are advertised by IBM Spectrum Control Base Edition (SCB)
- The corresponding policies in vCenter

SPBM gives the vSphere administrator the ability to create policies in the vSphere Web Client. These policies define the storage requirements of an application or virtual machine. When deploying a new virtual machine, the assigned policy filters all available datastores, and displays both compatible and incompatible storage options.

7.2 Storage capabilities advertised by SCB

IBM Spectrum Control Base Edition uses vSphere APIs for Storage Awareness (VASA) 2.0 functionality to advertise the storage capabilities provided by your IBM Spectrum Virtualize storage system. IBM Spectrum Control Base Edition detects the attributes of the available storage, and can offer the storage administrator the granular control of whether to grant or deny the vSphere administrator access to that particular storage resource or capability.

The storage capabilities are defined on the individual storage services within the SCB web interface. Multiple storage services can be created with differing capabilities to offer multiple VVol datastores, each with their own unique attributes. However, bear in mind that each storage space can only contain one VVol-enabled storage service.

Restriction: The storage capabilities are defined when creating a new storage service in IBM Spectrum Control Base Edition, and cannot be altered after creation.

Figure 7-2 shows a VVol storage service in IBM Spectrum Control Base Edition.

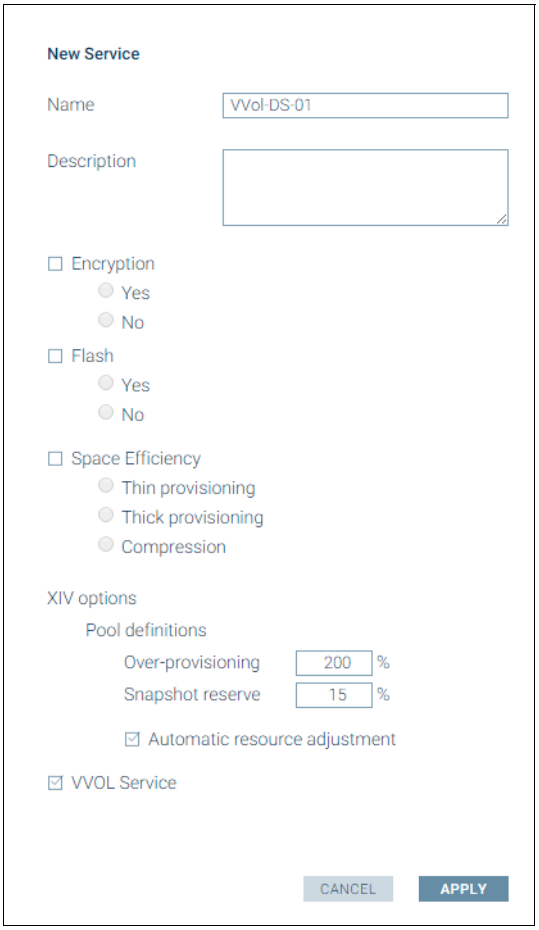

Figure 7-2 VVol Storage Service in IBM Spectrum Control Base Edition

Currently, IBM Spectrum Control Base Edition offers the following capabilities:

- Encryption
- Flash
- Space efficiency (compression)

7.2.1 Encryption

Providing that your IBM Spectrum Virtualize storage system is licensed for Encryption (either by hardware or software), the storage administrator can grant or deny access to the feature by selecting the Encryption check box, and then selecting **Yes** or **No**.

If the storage service has encryption set to **Yes**, all storage resources (child pools) created on the storage system are created with the **`-encrypt yes`** parameter.

Important: Encryption must be licensed, activated, and functional before configuring the feature in IBM Spectrum Control Base Edition.

If the Storage Service has encryption set to **No**, the storage resource/child pool is created with the **`-encrypt no`** parameter.

By leaving the Encryption check box cleared, it can be taken to mean *either*. However, remember that if encryption is enabled in the storage system, any child pools created are encrypted by default.

7.2.2 Flash

If the IBM Spectrum Virtualize storage system contains any SSDs or flash memory, the storage administrator can grant or deny the vSphere administrator access to these resources by selecting the Flash check box, and selecting **Yes** or **No**. This can be particularly useful with large flash parent pools, because the storage administrator might only want to allocate a small proportion of this expensive storage to the vSphere administrator.

7.2.3 Space efficiency (compression)

After enabling the IBM Real-time Compression™ (RtC) feature within the IBM Spectrum Virtualize storage system, the storage administrator can grant or deny access to this feature. If compression has not first been activated on the storage system, any creation of compressed VVols fails. This is to prevent an inquisitive vSphere administrator from creating compressed volumes in the storage system and affecting system performance without the storage administrator being aware.

Tip: Monitor the performance statistics of your storage system before and after enabling compression, because this feature could use a significant amount of system resources.

To enable compression on the storage service, select the SpaceEfficiency check box, and select the Compression radio button. Also, ensure that the **`lssystem`** command on the storage system returns `compression_active = yes`.

Example 7-1 shows the output of the **`lssystem`** command indicating that compression is enabled.

Example 7-1 Output of the lssystem command

```
IBM_2145:vvol-cg8:superuser>lssystem | grep compression_active
compression_active yes
```

If the **`lssystem`** command shows `compression_active no`, ensure that at least one compressed volume exists in the storage system and rerun the **`lssystem`** command.

7.3 Defining policies in vSphere Web Client

After defining the storage capabilities in IBM Spectrum Control Base Edition, the corresponding storage policies are required in vCenter. To configure these policies, from the Home action menu, select **Policies and Profiles**.

Figure 7-3 shows the Configure Storage Policies dialog in vCenter.

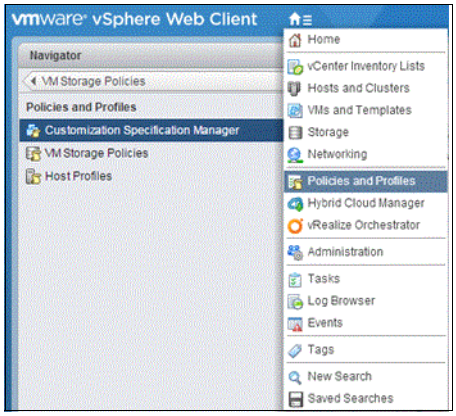

Figure 7-3 Configure Storage Policies in vCenter

The storage policies are essentially a selection of filters that correspond to the storage capabilities that were defined in the IBM Spectrum Control Base Edition interface. SPBM can also apply to existing VMFS datastores using VASA 1.0. Because of this, there are different profile options available to VMFS and VVol configurations. For VVol policies, ensure that the VASA 2.0 ruleset is selected. This ruleset is named `com.ibm.storageprofile.policy`, as shown in Figure 7-4.

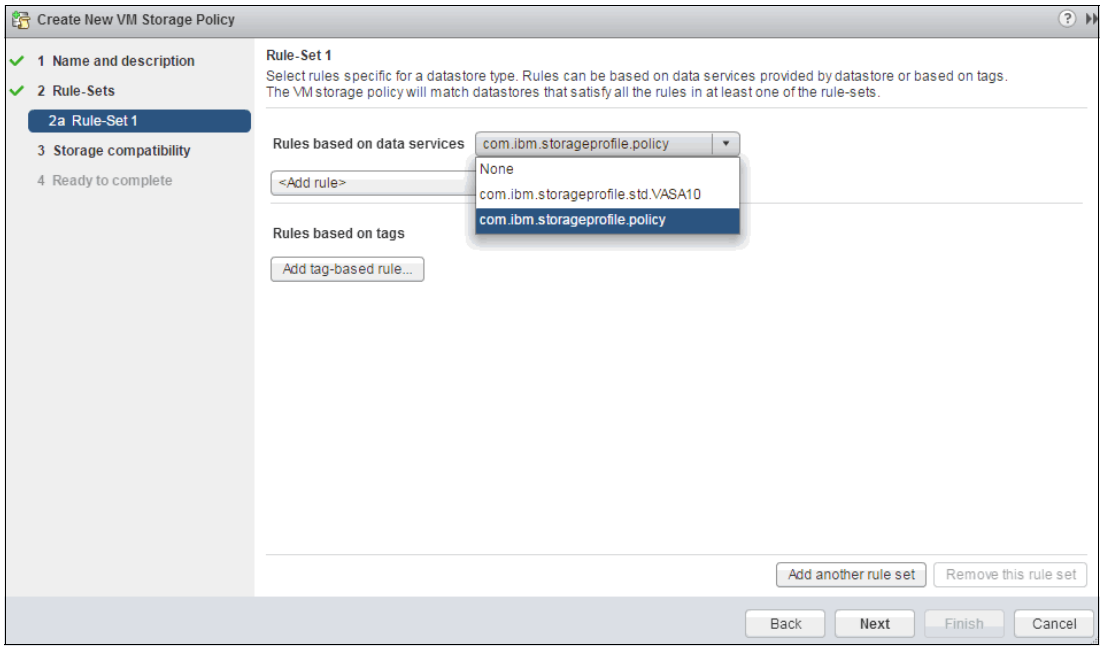

Figure 7-4 Create New VM Storage Policy

After selecting the VASA 2.0 ruleset, the following three attributes are available:

- Encryption
- Flash
- Compression

A policy can consist of one or more capabilities, with each capability defined as a rule. When adding a rule, the options are either *Enabled* or *Disabled*.

Figure 7-5 shows how to add rules to a service policy.

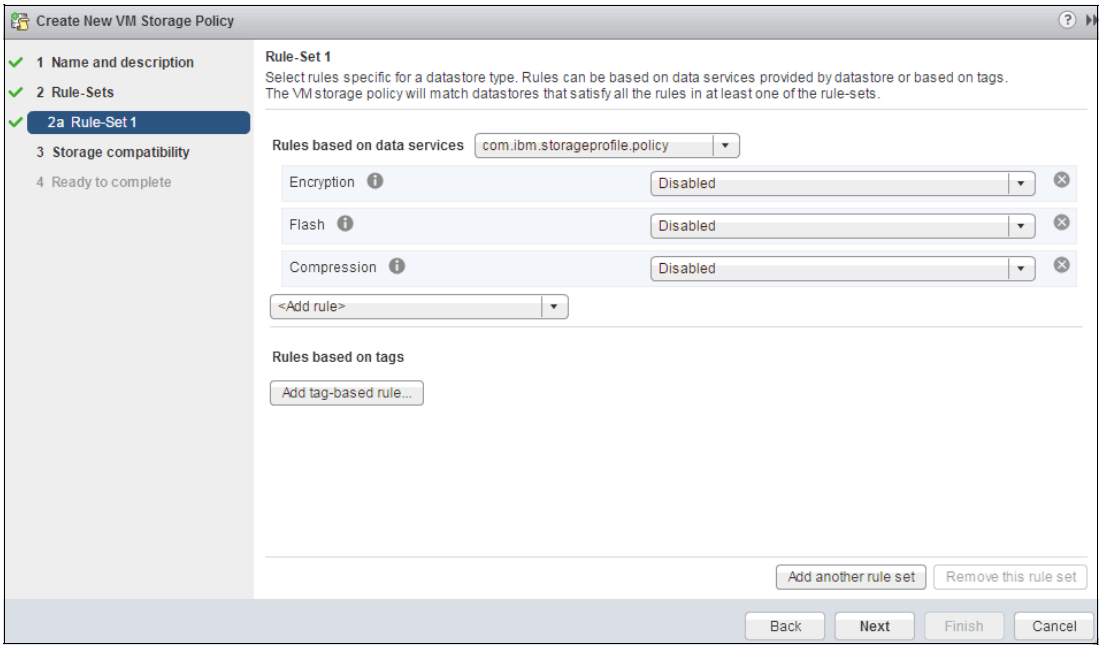

Figure 7-5 Creating a Storage Policy

7.4 Use case scenarios and examples

This section describes some use case scenarios and some examples.

7.4.1 Scenario 1

Consider creating a storage policy, with a single flash-based rule set to Enabled.

When deploying a virtual machine aligned to this policy, only VVol datastores backed by Flash/SSD MDisks are shown as compatible. Any VVol datastores that are made up of traditional spinning disks are shown as incompatible.

7.4.2 Scenario 2

Consider a further example when creating a policy with both Flash and Encryption rules, both set to Enabled.

When deploying a virtual machine aligned to this policy, only VVol datastores made from both Encrypted child pools in Flash/SSD MDisks are listed as compatible. All other datastores are shown as incompatible.

7.4.3 Scenario 3

Taking this one step further, consider a financial application that stores large amounts of confidential, text-based, archived financial information. A policy could be created specifying Encryption Enabled, Compression Enabled, Flash Disabled.

When deploying a virtual machine associated with this policy, only datastores backed by Enterprise/Nearline spinning disks in encrypted arrays, from storage systems with compression enabled are displayed as compatible. All other datastores are listed as incompatible.

When considering a virtual machine with multiple data disks, the storage requirements can differ between disks, Therefore, each disk could be assigned different policies. By configuring multiple policies with different properties, the associated virtual volumes would automatically be created in the appropriate VVol datastores, with the required characteristics.

Figure 7-6 shows default storage selection.

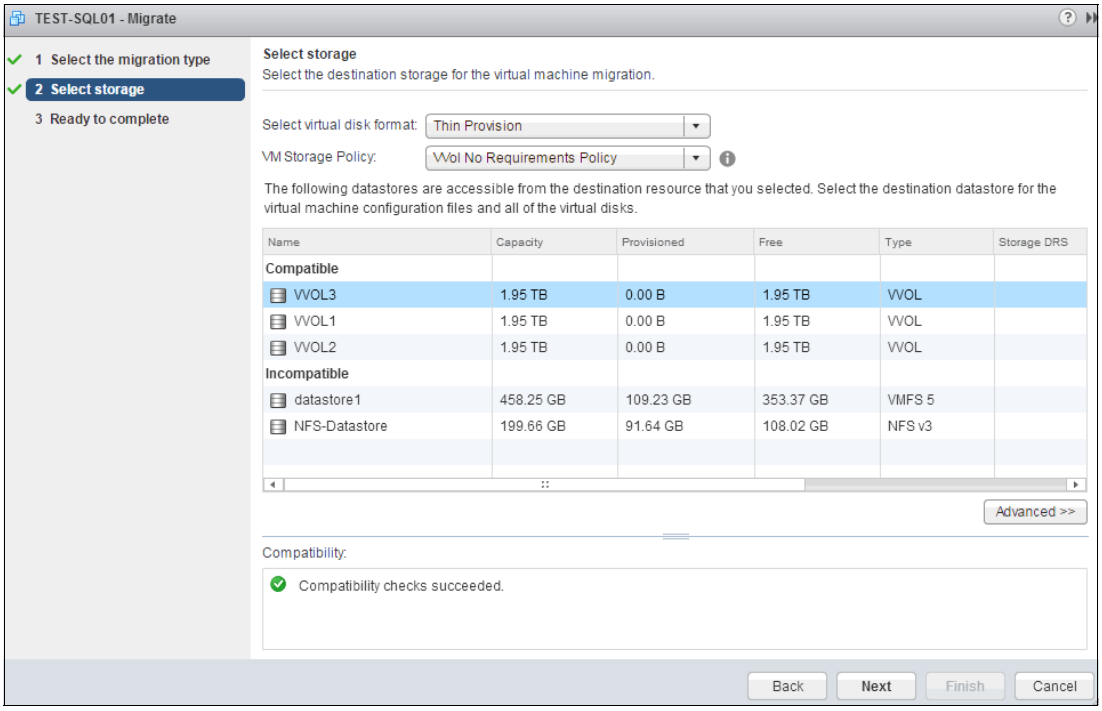

Figure 7-6 Default storage selection

Figure 7-7 shows storage selection based on policy.

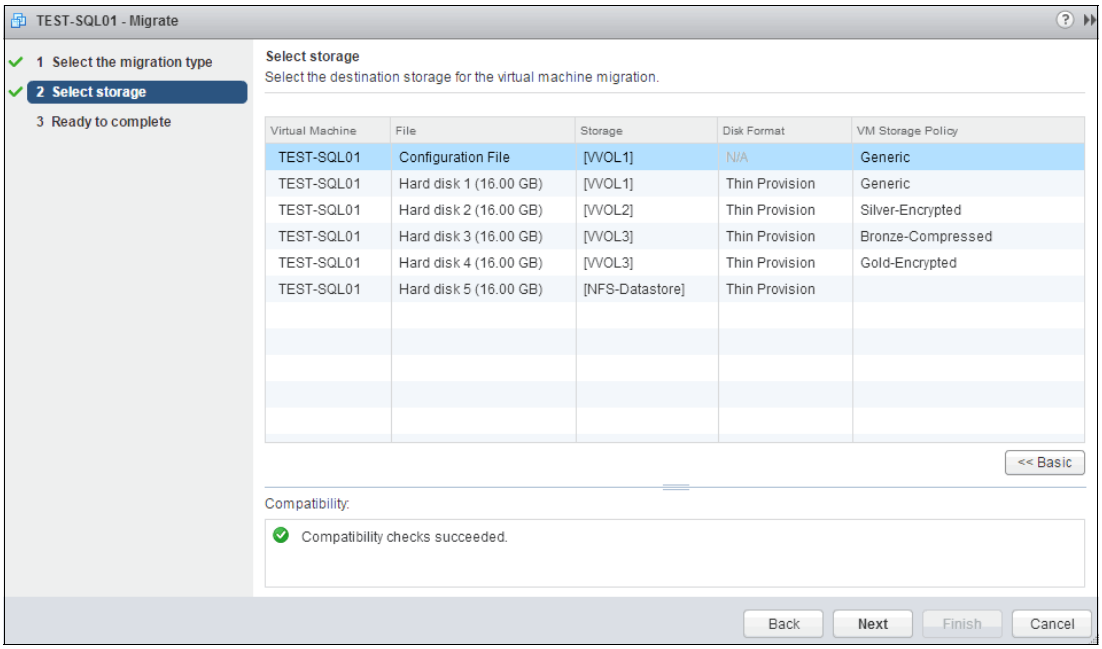

Figure 7-7 Storage selection dialogue

An ideal use case for this would be applied when using virtual machine templates. A vSphere administrator might rely on virtual machine templates for different deployment application requirements. Similarly, storage policies can be associated either to the template or individual volumes within the template. When using these templates, any policy associations propagate down to any cloned virtual machines.

Therefore, vSphere administrators do not need to retroactively change the policies associated to each volume after the virtual machine has been deployed. This can greatly improve the time to deploy and configure large numbers of virtual machines.

Consider the following examples where a virtual machine would have different volumes with specific storage requirements.

Example 7-2 shows database servers, such as Microsoft Windows SQL.

Example 7-2 Database servers

```
Config Home (VM metadata) – Generic / Silver tier
C:\ OS Volume – Generic / Silver tier
D:\ Software Binaries – Compressed / Bronze tier
E:\ Database – Flash / Gold tier
F:\ Backups – Encryption / Bronze tier
```

Example 7-3 shows file servers with document archives and sensitive data.

Example 7-3 File servers

```
Config Home (VM metadata) – Generic / Silver tier
C:\ OS Volume – Generic / Silver tier
D:\ Document Archive File share – Compressed / Bronze tier
E:\ Financial File shares (confidential) – Encryption / Silver tier
F:\ VIP File shares (confidential) – Encryption / Gold tier
```

Example 7-4 shows a web server that is powering a mobile app that a development team is creating.

Example 7-4 Web server

```
Test & Application Development Guest VM
Config Home (VM metadata) – Compressed / Bronze tier
C:\ OS volume – Compressed / Bronze tier
D:\ Application – Compressed / Bronze tier
```

Because the application is in the early stages of deployment, the storage requirements are minimal. Therefore, the app is currently assigned the smallest footprint possible by enabling compression on the slowest/cheapest tier available.

Perhaps one member of the development team demonstrates the app to his friends, who then share it on a popular social network. All of a sudden, 1000s of users are now logging in to the application and the storage requirements have rapidly changed. By changing the storage policy associated to the virtual machine, you can now ensure optimal response time and availability to the user base.

7.4.4 Configuring storage services with multiple child pools

The Storage administrator has the option of provisioning multiple child pools to a single storage service. In this scenario, although all child pools associated to a single storage service must be on the same storage system, they can be from different parent pools or tiers.

In this scenario, a single VVol datastore (for example, named VVol-Hybrid in vCenter) can be made from both a child pool from a Flash parent pool *and* a child pool from traditional spinning disks. When provisioning a virtual machine in vCenter with a `Flash=Enabled` policy, the VVol-Hybrid datastore is shown as compatible. In this case, when the associated VVols are created on the storage system, they are created in the child pool from the Flash/SSD parent pool.

Conversely, if another virtual machine was to be deployed and assigned a `Flash=Disabled` policy, the same VVol-Hybrid datastore is still listed as compatible. This time, the associated VVols are created in the child pool made from the traditional Enterprise/Nearline parent pool.

Remember: In this hybrid datastore scenario, the vSphere administrator only has visibility of the total capacity of the datastore, with no indication of the usage/capacity limits of the individual pools within.

Similarly, consider a storage service has a capacity of 2 terabytes (TB) made up of multiple pools from different storage tiers, for example, 500 gigabytes (GB) Flash, 500 GB 15,000 revolutions per minute (RPM), 500 GB 10,000 RPM, 500 GB Nearline 7200 RPM. In this complicated configuration, the vSphere administrator assumes that there is 2 TB of free space in this datastore.

However, the size of any VVol in this datastore is limited to 500 GB, because a single volume on the storage system cannot span multiple child pools. Particular attention should be given to space-efficient volumes, which grow in size as new data is written.

8

Advanced concepts, operations, and functions in the Virtual Volumes environment

This chapter contains information about the internal storage system mechanisms that are required to support the Virtual Volumes (VVols) environment. Strictly speaking, the storage administrator does not *need* to be aware of all of the details contained within this chapter, because the entire environment is maintained automatically by the system after the initial configuration has been performed. However, there are a few items described here that are pivotal to understanding VVols within the context of your storage system:

- 8.1, “Types of VVols” on page 94
- 8.2, “Number of VVols” on page 94
- 8.3, “IBM FlashCopy functionality in a VVols environment” on page 95

For interested readers, we also consider how the storage system responds to the following vSphere tasks in the VVols environment:

- 8.4, “Clone operations” on page 96
- 8.5, “Snapshot operations” on page 98
- 8.6, “Cold migration and storage vMotion in a Virtual Volumes environment” on page 103
- 8.7, “vMotion in a Virtual Volumes environment” on page 104

8.1 Types of VVols

From the perspective of the vSphere infrastructure, there are five different types of VVol. Each of them provides a specific function depending on their role within the vSphere environment:

- Config VVol
  - Contains VM configuration files, logs, and so on.
  - Created on the storage system as a thin-provisioned VVol with 4 gigabyte (GB) capacity.
- Data VVol
  - Equivalent to a conventional virtual machine disk (VMDK), holds a VM's system data.
  - Created on the storage system per the selected storage policy. Capacity corresponds to size of the virtual machine's (VM's) hard disk.
- Swap VVol
  - Virtual machine memory swap file.
  - Created on the storage system as a thick-provisioned VVol. The capacity correlates to the amount of memory allocated to the VM.
- Memory VVol
  - Captured as part of some snapshot operations.
  - Created on the storage system as a thick-provisioned VVol. Capacity directly correlates to the amount of memory allocated to the VM.
- Other VVol
  - vSphere solution-specific object.

All virtual machines are created with a Config VVol and a Data VVol. When powered on, a Swap VVol is created for a virtual machine. This means that any running VM consists of at least three VVols. Because each of a virtual machine's disks corresponds to an individual VVol, adding another hard drive to a VM creates a unique Data VVol on the storage system.

The usage for a Memory VVol is described in the following sections, when we consider virtual machine snapshots. Other VVols are rarely seen in the conventional environment. Therefore, a description of such objects is not covered in this IBM Redbooks publication.

8.2 Number of VVols

It is important to remember that, from the point of view of the storage system, a VVol is just a special type of volume, with certain unique properties. Consequently, you should note that the total number of volumes provisioned on the storage system is inclusive of the number of VVols. This means that VVols should be factored into any metrics concerning maximum limits with regards to volumes.

Because the storage administrator does not need to interact with individual VVols, they are filtered out of the Volumes panel on the management graphical user interface (GUI), for convenience. However, you can use the GUI to monitor the overall number of VVols in a child pool by navigating to **Pools** → **Volumes by Pool**.

Figure 8-1 shows the volume count for a child pool allocated to the VVols environment.

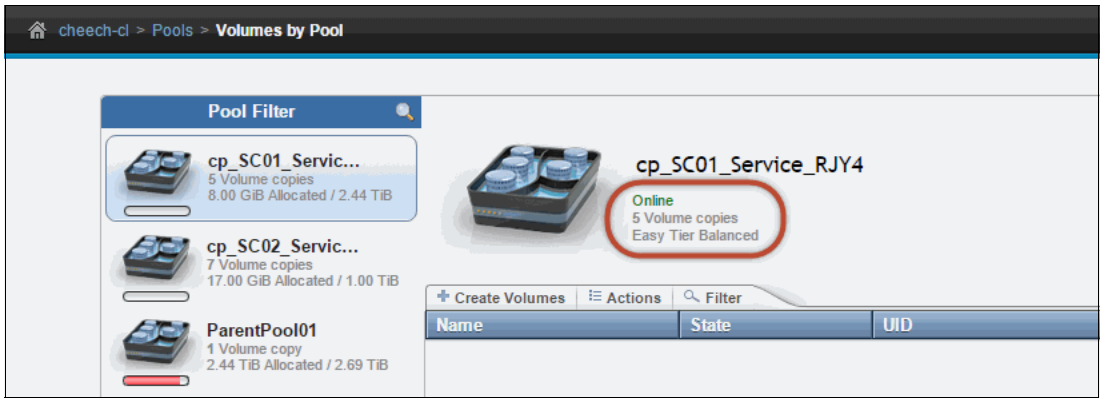

Figure 8-1 Monitoring the number of VVols in a child pool

8.3 IBM FlashCopy functionality in a VVols environment

The capability to take a snapshot or clone of a virtual machine has been available in the vSphere environment for several years. Before the introduction of Virtual Volumes, the disks generated by these operations were contained within a VMFS datastore, and the relationships between them were managed entirely by the vSphere infrastructure.

By making virtual machine disks (individual VVols) available at the storage level, the Virtual Volumes feature enables the storage system to handle snapshot and clone operations using IBM FlashCopy.

Important: As the use of clone and snapshot operations is ubiquitous in vSphere environments, we stipulate that the appropriate IBM FlashCopy license for your system is configured before implementing Virtual Volumes.

There are several benefits that can be derived from offloading these tasks to the storage system, including the following benefits:

- Processor (CPU) use required for snapshot and clone tasks is transferred to the storage system, freeing up VMware ESXi host resources.
- Rapid availability of clones, regardless of data disk size.
- Increased recommended lifecycle and length of chain for snapshots: No requirement to remove a snapshot after 72 hours or keep chains to 2 - 3 snapshots.

Later in this chapter, we consider how FlashCopy achieves the last two benefits by analyzing some specific use-cases.

8.3.1 A basic introduction to IBM FlashCopy

Before considering specific VVols-related tasks, it is important to understand the basic function that FlashCopy provides. By enabling users to take a point-in-time copy of a volume, FlashCopy offers a powerful solution for several critical and challenging business requirements. With FlashCopy, both the source and target volumes involved in the operation are available for host input/output (I/O) as soon as the relationship is started. Crucially, both volumes remain accessible throughout the copy process.

In the Virtual Volumes environment, the source volume in a FlashCopy relationship could be the Data VVol, for example, the guest operating system (OS) for a VM. This entails that a VVols-enabled ESXi host can have access to a full, independent image of this Data VVol (the FlashCopy target) almost instantaneously, without affecting I/O operations to the source.

Consequently, we can immediately envision scenarios where FlashCopy can provide benefits in the VVols environment. For a further explanation of particular VVols use-cases, see the following examples.

Additional information: FlashCopy delivers these powerful capabilities by implementing a bitmap to track the differences between the source and target volumes involved in a FlashCopy relationship. This bitmap is used to record whether each grain (64 kilobyte (KB) or 256 KB unit of data) is common to the source and target volumes.

On starting a FlashCopy mapping, this bitmap is initialized. In the majority of cases, every grain in the bitmap is marked as *unsplit*, meaning that both volumes contain an identical set of data (from the host point of view).

An incoming host write to either the source or target volume, results in updates to the bitmap, indicating that the data for the corresponding *grain* is now *split*. That is, the data on the source and target is non-identical for a particular unit of data. An incoming read operation to the target results in the system consulting the bitmap to establish whether the relevant grain should be read from the source (in the case of an unsplit grain) or the target (for a split grain).

When required, an indirection layer enables read operations to be redirected to the source without disrupting the host or application. From the host point of view, this enables us to create the impression that the target volume presents a complete version of the data almost instantaneously, without quiescing the I/O for a long time, or taking the source volume offline. For further information about IBM FlashCopy, see *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574.

8.4 Clone operations

In the vSphere environment, a virtual machine can be cloned while it is running or powered-off. It is also possible to clone a virtual machine from a template. FlashCopy is used to perform these tasks if the virtual machine is on a VVol datastore. For the purposes of this example, we consider a simple clone operation of a powered-off virtual machine.

Note: These steps vary slightly if the virtual machine is running, but the core FlashCopy function is unchanged by this factor.

We start with the following VVols:

- `Source_Config_VVol` (containing virtual machine metadata)
- `Source_Data_VVol` (containing the OS)

When the user initiates a clone operation from within the vSphere Web Client, IBM Spectrum Control Base Edition (SCB) runs several commands on the storage system in rapid succession.

SCB runs the following commands:

1. Create a Config VVol for the clone target. This stores metadata for the new VM.
2. Create a Data VVol for the clone target (`Target_Data_VVol`). This contains the OS for the new VM.
3. Create bindings between the VVols and a protocol endpoint (PE). This enables an ESXi host to access the newly created VVols for I/O operations.
4. Create a FlashCopy mapping of type `clone` between `Source_Data_VVol` and `Target_Data_VVol`.
5. Prepare and start the FlashCopy mapping.

At this point, IBM Spectrum Control Base Edition can inform vSphere that the new virtual machine is available for use. Because the target of a clone operation is available to the host upon the FlashCopy relationship *starting* (rather than on *completion*). The new virtual machine is available almost instantaneously. Figure 8-2 provides an overview of the clone operation for a virtual machine's Data VVol.

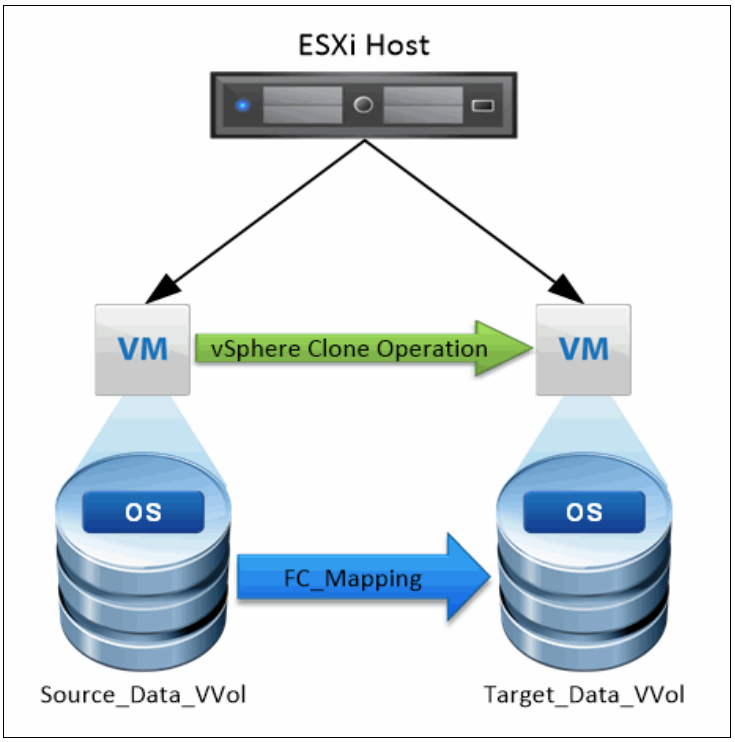

Figure 8-2 FlashCopy clones the Data VVol for a source virtual machine. Note that the ESXi host has full read/write access to both Data VVols as soon as the mapping is started.

As far as the vSphere administrator is concerned, he has a new stand-alone virtual machine as soon as the FlashCopy relationship has been started. Because of this, the required resource use (in terms of the speed of the availability of the target) is not increased significantly by the size of our `Source_Data_VVol`. Consequently, a VM with very large additional disks can be cloned extremely rapidly when using FlashCopy.

Additional information

Behind the scenes, triggering a clone operation causes the storage system to begin transferring data from `Source_Data_VVol` to `Target_Data_VVol` using a process known as *background copy*. Setting a background copy rate of greater than zero results in a complete copy of the source volume being created on the storage system.

Because the clone command triggered by SCB also specifies the `-autodelete` flag, the FlashCopy mapping is deleted after the background copy process is complete. This results in a stand-alone copy of the data and, in this case, a stand-alone virtual machine.

8.5 Snapshot operations

In the vSphere environment, users can take a snapshot of a virtual machine. They can revert to this point-in-time image on demand. For virtual machines based on VVol datastores, these tasks are again offloaded to the storage system, where FlashCopy is used to meet the requirements of the vSphere environment.

8.5.1 Taking a snapshot

The process for taking a snapshot of a virtual machine is almost identical for a powered-on or powered-off virtual machine. We indicate the differences for a live VM where relevant.

We start with the following VVols:

- `Base_Config_VVol` (containing virtual machine metadata)
- `Base_Data_VVol` (containing the OS)
- `Swap_VVol` (only present if the VM is powered on)

When the user takes a virtual machine snapshot using the vSphere Web Client, SCB runs several commands on the storage system in rapid succession:

1. Create a Data VVol, from now on called `Snapshot_Data_VVol`. This serves as the point-in-time image of `Base_Data_VVol`.
2. Create a FlashCopy mapping of type `snapshot` between `Base_Data_VVol` and `Snapshot_Data_VVol`.
3. Prepare and start the FlashCopy mapping.

Note: If the user chooses to snapshot the virtual machine's memory, SCB also performs the following actions:

1. Creates a new Memory VVol.
2. Creates a binding between the Memory VVol and a protocol endpoint. This enables an ESXi host to write the snapshot memory-state to this VVol.

Taking a snapshot of a virtual machine does *not* result in any downtime of the VM. When the relationship is started, the source volume is briefly “frozen” to initialize the FlashCopy bitmap, and then the I/O on the source volume is enabled to resume. The target volume (`Snapshot_Data_VVol`) serves as a point-in-time image of the `Base_Data_VVol`, and is never directly accessed for I/O operations by the ESXi host, as shown in Figure 8-3.

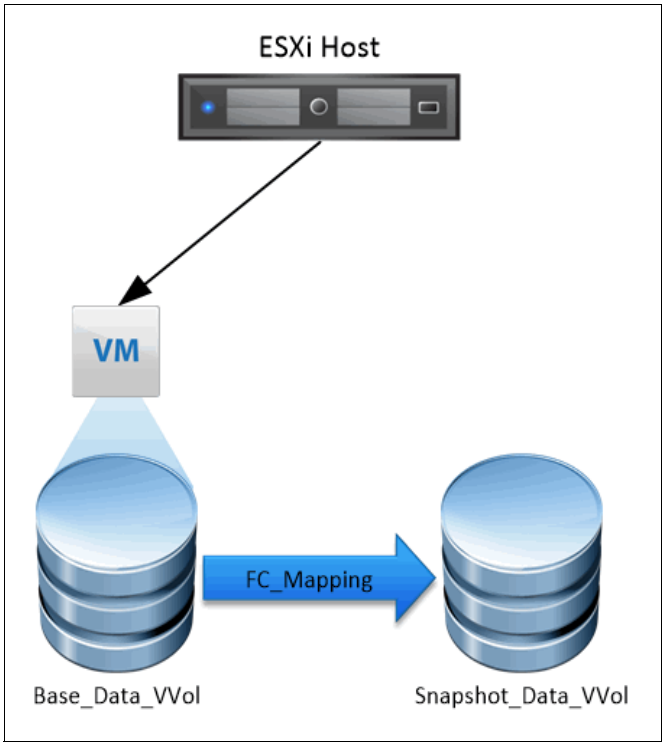

Figure 8-3 In a snapshot operation, the ESXi host only requires access to the Base_Data_VVol for I/O operations

Additional information

The storage system *does not* create a full copy of the `Base_Data_VVol` if the FlashCopy mapping is a snapshot. Unlike a clone, a snapshot FlashCopy relationship is created with a background copy rate of `zero`. Data is only copied onto the target (`Snapshot_Data_VVol`) if changes are made to the source (`Base_Data_VVol`).

For example, we might update the OS hosted on `Base_Data_VVol` after the snapshot is taken. For each host write associated with this update, FlashCopy checks the bitmap to see whether the corresponding grain has already been copied to the target.

Assuming that it has not (as is the case in our snapshot example), the original data for this grain is copied onto `Snapshot_Data_VVol`. The incoming data can then be written to `Base_Data_VVol`.

This approach ensures that a snapshot volume only ever contains the *changes* and never occupies more space than it needs to. See Figure 8-4.

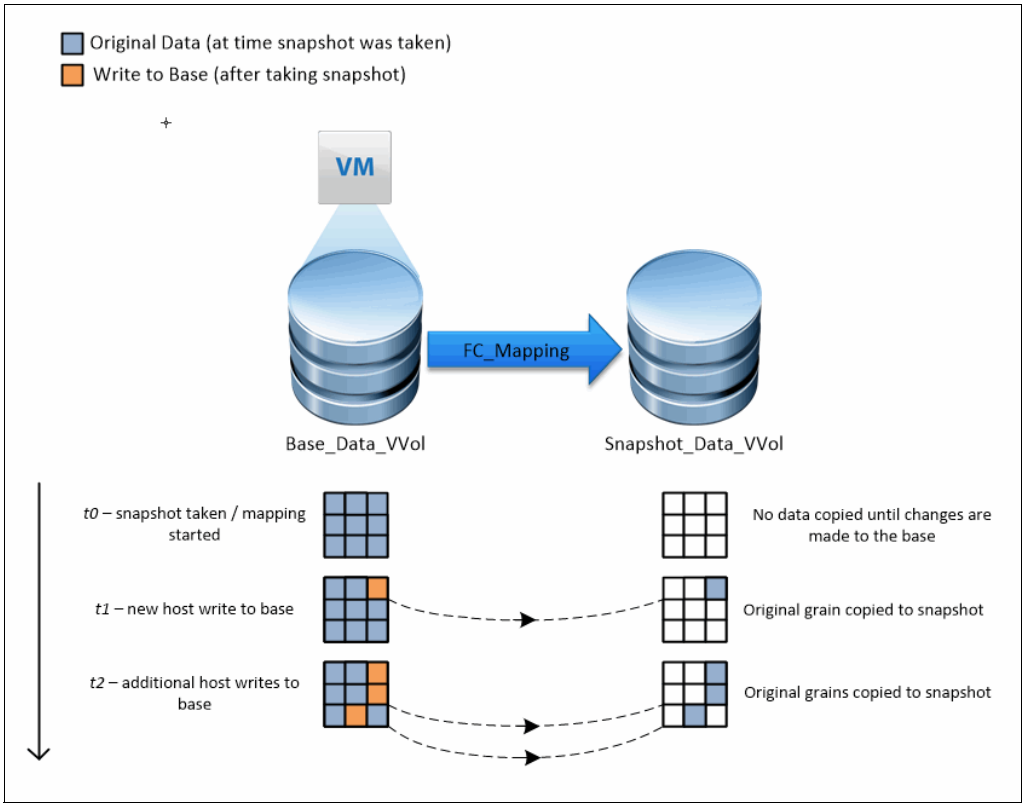

Figure 8-4 Movement of data in a snapshot FlashCopy relationship

Note here that the new data is always written to the source, and only original data (the data as it was when the snapshot was created) is copied onto the target. This contrasts with the current VMware file system (VMFS) model, where the new disk created upon taking a snapshot always becomes the running point for the virtual machine.

A snapshot disk (for example, `Snapshot_Data_VVol`) is created as a thin-provisioned VVol with a very small real size. This ensures that the snapshot VVol contains the smallest amount of storage possible when first created (*t0* in Figure 8-4). Setting the autoexpand parameter to `on` ensures that the snapshot can grow as required (when changes are made to the source), if there is capacity available in the child pool.

8.5.2 Reverting to a snapshot

Outside of the context of the Virtual Volumes environment, it is already possible to use FlashCopy to restore a volume to an earlier state, using a technology known as *Reverse FlashCopy*. For more information about this, see the section about using Reverse FlashCopy to restore a volume in *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services*, SG24-7574.

In the Virtual Volumes environment, a revert to snapshot operation run by the vSphere administrator uses similar FlashCopy functionality.

Consider a simple example of a VM with a single snapshot. We start with the following VVols:

- `Base_Data_VVol` (containing the OS). This is the live version of the data disk, which is accessible to the host.
- `Snapshot_Data_VVol`. This is the point-in-time image of `Base_Data_VVol`, taken at a known *good state* (for example, before applying an update).

The FlashCopy relationship is `Base_Data_VVol` to `Snapshot_Data_VVol`.

Suppose we apply an update to the VM's guest OS and this causes an issue with an important application. Consequently, we decide that we want to roll back to our snapshot. Initiating the revert to snapshot operation from the vSphere Web Client causes IBM Spectrum Control Base Edition to start a Reverse FlashCopy relationship from `Snapshot_Data_VVol` to `Base_Data_VVol`.

This gives the host the impression that the `Base_Data_VVol` is identical to `Snapshot_Data_VVol`. The ESXi host can resume I/O operations to `Base_Data_VVol` as soon as the FlashCopy restore operation is initiated.. VM operations can resume from our known *good state* with minimal delay.

Important: The vSphere infrastructure also triggers the creation of a temporary snapshot as part of a revert to snapshot operation. This is a protective measure specific to the ESXi task, and is not pertinent to the FlashCopy functionality used to restore the VM. For simplicity, we omit these steps from the example described here.

Importantly, FlashCopy restore functionality enables the storage system to ensure that the ESXi only ever needs to perform I/O operations to the `Base_Data_VVol`. The ESXi host never directly accesses `Snapshot_Data_VVol`. This enables us to maintain a consistent point-in-time image, so that the vSphere administrator can revert to the same snapshot again later. A snapshot can be reverted to as many times as required.

Additional information

To understand how a FlashCopy restore works, it is essential to grasp the concept of a *dependency chain*.

Because the full copy of data in a FlashCopy relationship only exists on the source volume when the mapping is started, it can be said that the target depends on the source. In the previous example, starting the snapshot mapping (`Base_Data_VVol` to `Snapshot_Data_VVol`) gives us the dependency chain `Base_Data_VVol(p)` to `Snapshot_Data_VVol(p)`.

Starting the restore operation (`Snapshot_Data_VVol` to `Base_Data_VVol`), introduces a secondary instantiation of `Base_Data_VVol` and gives us the following chain:

`Base_Data_VVol(s)` → `Snapshot_Data_VVol(p)` → `Base_Data_VVol(p)`

`Base_Data_VVol(s)` must remain at the base of the dependency chain, because the base image of the data is here and `Snapshot_Data_VVol` depends on it. The primary instance, `Base_Data_VVol(p)`, services host I/O directed at `Base_Data_VVol`. The indirection layer ensures that a host read to a grain that has not yet been restored is redirected to `Snapshot_Data_VVol`, giving the impression that the restore is immediate.

In the background, the system copies grains from `Snapshot_Data_VVol` onto `Base_Data_VVol`, until the restore is complete. The restore operation only copies the grains that are needed to restore `Base_Data_VVol` to the state that it was in when our snapshot was taken (only the changes that have occurred since taking the snapshot are copied).

After this operation completes, we return to our original FlashCopy dependency chain:

`Base_Data_VVol(p)` → `Snapshot_Data_VVol(p)`

This means that we can give the host the impression of an immediate restore, while preserving the consistency of `Snapshot_Data_VVol` for future restore operations.

At the time of writing, the FlashCopy restore functionality used for Virtual Volumes is not available to the general user. In addition to the simple scenario described previously, the new FlashCopy restore functionality is unique in its ability to deal with the following use cases:

- Request to perform a new restore when another restore is already in progress
- Request to perform a restore when the VVol to be restored is also the target in a FlashCopy relationship

This ensures that the vSphere administrator can perform a revert to snapshot operation on demand, on any VM. The functionality required to perform a FlashCopy restore in the complex use cases mentioned previously is beyond the intended scope of this IBM Redbooks publication.

8.5.3 Deleting snapshots

In the Virtual Volumes environment, each snapshot (delta) is maintained as a unique VVol. A FlashCopy relationship is used to track the differences between the base VVol and a snapshot. It is important to remember that the host only ever reads from the base VVol, which always contains the latest data and serves as the running point for the virtual machine. The snapshot VVol only contains the original data for the grains that have been changed on the base VVol. The host never directly performs I/O operations to the snapshot disk.

Consequently, a delete snapshot operation simply causes IBM Spectrum Control Base Edition to delete the snapshot VVols from the storage system. This process is completely nondisruptive for I/O operations.

Additional information

In the VMFS environment, taking a snapshot causes the creation of a *delta* file. This delta file always becomes the running point for VM. That is, after we've taken a snapshot, write I/O is serviced by the delta disk (snapshot disk), although the base disk services read I/O. Consequently, in the VMFS environment, deleting a snapshot (or snapshot chain) results in the data contained on the delta disks being rolled back into the base disk.

These operations can be time-consuming if the delta disk becomes too large, or if the chain of snapshots is too long. Consequently, VMware's official preferred practice for VMFS snapshot management states that no single snapshot should be used for more than 24 - 72 hours, and that only 2 - 3 snapshots should be used in a single chain.

Deleting very large snapshots (or large chains of snapshots) is non-problematic in the VVols environment. Because the running point of the VM remains constant (the base VVol), deleting the snapshots simply discards the unused VVols. This entails that the leading practices for snapshots, regarding lifecycle and length of chain, do not apply in the VVols environment.

Although this leading practice is no longer applicable, you might still want to delete snapshots that are no longer required to free up system resources. VVols that serve as managed snapshots only occupy a tiny amount of space when first created. However, each new write to the running point of the VM (the base VVol) causes the original data to be copied onto the managed snapshot VVol. Consequently, deleting old snapshots can be economical, because the amount of data contained by these VVols is likely to be more substantial.

8.6 Cold migration and storage vMotion in a Virtual Volumes environment

The vSphere administrator can transfer the storage resources for a virtual machine between datastores. If the virtual machine is powered-off, this operation is known as a *cold migration*. If the virtual machine is running, the operation is known as a *Storage vMotion*. The Virtual Volumes environment introduces support for the following migration paths for both operations:

- VVol Datastore → VVol Datastore
- VMFS Datastore → VVol Datastore
- VVol Datastore → VMFS Datastore

Tip: The ability to migrate VMs between VVol and VMFS datastores (bidirectionally) ensures that a mixed VVol and VMFS environment is possible.

As far as the storage system is concerned, the basic process for each of these operations is the same:

1. Create new VVols (or `.vmdks`) on the destination datastore.
2. Establish a data path between the new VVols and the ESXi host, by creating new subsidiary volume mappings.
3. Copy the data from the source VVols (or `.vmdks`) to the target VVols (or `.vmdks`).
4. When the copy is complete and the VM is solely dependent on the target VVols (or `.vmdks`), remove the source VVols (or `.vmdks`) from the storage system.

Although the fundamental principles of all of these storage migrations are the same, the copy method selected by the ESXi host in step 3 differs, depending on the specifics of the operation (for example, the type and location of the source and destination datastores).

Where possible, the ESXi host uses vSphere APIs for Storage Awareness (VASA) or vStorage APIs for Array Integration (VAAI) to perform the copy operation to optimize performance. These primitives translate into FlashCopy and XCOPY operations on the storage system. If these optimized copy methods are not available (for example, when the source and target datastores are on separate storage systems), the ESXi host automatically falls back to traditional methods.

8.7 vMotion in a Virtual Volumes environment

A vMotion operation is used to transfer a running virtual machine from one compute resource (ESXi host) to another. Because a vMotion operation only changes the compute resource (and not the storage resource), the effect on the storage system is relatively minimal. To complete the operation, the following operations are run by IBM Spectrum Control Base Edition:

1. Create a new Swap VVol.
2. Create a subsidiary volume mapping between the new Swap VVol and the target ESXi host.
3. Create subsidiary volume mappings between the config and data VVols and the target ESXi host. This establishes the data path between the VVols and the new compute resource.
4. After the vSphere infrastructure has transferred the compute resource, remove the redundant subsidiary volume mappings (the mappings between the source ESXi host and the VVols).
5. Remove the additional Swap VVol, because this is no longer required.

At this point, the VVols on the storage system are identical to those that existed before the vMotion operation was initiated. Only the data path has been altered.

Redbooks

Configuring VMware Virtual Volumes for Systems Powered by IBM

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8328-00

ISBN 0738441414

Printed in U.S.A.